Overlapping Fund Generations and the Self-Reinforcing Growth of Private Credit

Erica Xuewei Jiang[†] David X. Xu[‡]

ABSTRACT

We study how the intermediation structure of private credit shapes lender incentives and creates a feedback loop between lending and fundraising. Lenders (GPs) operate overlapping generations of funds, allowing them to use capital from a young fund to lend to a credit-impaired borrower held by a seasoned fund. This rescue lending preserves the seasoned fund's reported return without resolving the underlying insolvency, increasing the GP's performance-based compensation, future fundraising, and thus its capacity for further (rescue) lending. Linking U.S. private credit transactions to loan valuations, fund returns, and fundraising, we find that impaired borrowers disproportionately obtain new credit from existing lenders, despite no subsequent recovery in credit conditions. Rescue occurs when the GP has both a well-performing seasoned fund and deployable capital from a young fund. This combination of GP incentives and capacity also predicts more favorable subsequent seasoned-fund returns, and these returns attract future LP commitments. We develop a quantitative model to analyze how lending and fundraising reinforce each other through this mechanism, amplifying both lending expansion and credit-risk accumulation.

[†]UCLA Anderson and NBER. Email: erica.jiang@anderson.ucla.edu.
[‡]Southern Methodist University. Email: davidxu@smu.edu.

1 Introduction

Corporate lending has significantly shifted toward nonbank private-credit lenders, whose assets under management grew from less than $200 billion before the global financial crisis to nearly $2 trillion by 2024.[1] These lenders typically operate through closed-end, finite-life funds: they raise capital from institutional investors and lend it to companies, particularly those with substantial credit risk.[2] As they grow, their lending decisions increasingly shape where corporate credit flows and the risk profile of marginal borrowers. The resulting credit losses ultimately fall on the pension funds, insurers, and other investors that supply their capital; a 2026 proposal to facilitate private assets in defined-contribution plans could extend this exposure more directly to retail retirement savers.[3] Understanding how these lenders intermediate between institutional investors and borrower companies is therefore central to assessing their effects on credit allocation and aggregate risk.

We study how this intermediation structure shapes lender incentives, giving rise to a novel mechanism through which lending and fundraising reinforce each other and contribute to the growth of private credit. Central to the mechanism is new credit that lenders extend to their existing borrowers with impaired credit quality, a practice we refer to as "rescue lending." This practice does not necessarily eliminate the underlying credit risk, but lenders nonetheless have incentives to rescue because they operate overlapping generations of funds and receive compensation based on fund performance. By allowing impaired borrowers to continue debt servicing, such lending preserves reported fund performance, which not only increases fee income from existing funds, but also facilitates subsequent fundraising and provides more capital for lending. Thus, this mechanism links the capital supplied by young funds to the performance of old funds, creating a feedback loop that can amplify both lender growth and accumulated credit risk.

We test the feedback loop by linking data on private-credit transactions, borrower condition, fund performance, and fundraising. Its first step—identifying rescue lending—requires observing

[1] For estimates of the size and growth of the private-credit market, see Davydiuk, Marchuk and Rosen (2024a), Jang, Kim and Sufi (2025b), and Jang, Kim, Sufi and Chen (2025a).

[2] Private-credit loans outstanding total roughly $1.4 trillion, equal to one-third of below-investment-grade debt outside bank loans (Board of Governors of the Federal Reserve System, 2026).

[3] See the proposal here: (U.S. Department of Labor, 2026).

both the lender on a new loan and the borrower's condition before origination. PitchBook provides a broad database of U.S. private-credit transactions that identifies borrowers and lenders. Because these borrowers are generally private and unrated, we measure their condition using quarterly BDC filings, which report the fair value and cost of debt held by BDCs. Building on Davydiuk, Erel, Jiang and Marchuk (2025), who use BDC-reported loan valuations to proxy for borrower credit quality, we aggregate fair value and cost across each borrower's BDC-held debt and classify the borrower as impaired when that debt is valued below 90 cents per dollar of cost.

We then link the two sources. For each transaction, we use the most recent BDC valuation reported strictly before origination and within the prior year. Because the link is at the borrower rather than the loan level, the new loan and its lender need not appear in a BDC portfolio. The resulting sample therefore extends beyond BDC loans: among debt instruments issued to borrowers with a BDC valuation in the same year, 42 percent are never themselves reported in a BDC portfolio. Pre-loan valuation coverage rises over time, reaching 61 percent of private-credit debt instruments and 57 percent of debt volume by 2024. We add fund-level reported returns and manager fundraising histories to study the rest of the mechanism, tracing rescue lending and borrower outcomes through older-fund performance and subsequent fundraising to successor-fund lending capacity.

Using these data, we document a sequence of facts consistent with self-reinforcing growth. We begin by showing that the more deeply impaired a borrower's existing debt, the more likely its new loan comes from an existing lender. Among borrowers receiving a new loan, the incumbent's share is below 30 percent when existing debt is valued near or above cost, 39 percent when it is valued between 80 and 85 cents per dollar of cost, and 43 percent when it is valued below 80 cents. The pattern scales with the incumbent's stake: the deeper the markdown, the larger the loss that keeping the borrower current defers—and only for the incumbent does a new loan also protect an existing claim.

This compositional pattern could still reflect differences across impaired borrowers rather than a relationship-specific lending pattern: impaired borrowers may differ in their industries, lender networks, and deal structures. We therefore recast originations as a matching problem, forming

potential pairs of active borrowers and lenders in each year and asking which pairs transact. Comparing potential lenders that face the same borrower in the same year, an existing relationship is associated with a 27-percentage-point higher probability of a new loan. This association is a further 8.8 percentage points stronger when the borrower is impaired—an increase of roughly one-third. Lender-year fixed effects absorb each lender's annual capacity and risk appetite. The association between an existing relationship and a new loan is therefore stronger among impaired borrowers even after accounting for borrower-year demand and lender-year supply.

Does the incumbent's tilt reflect informed selection? In the classic account of relationship lending, the inside lender's information advantage lets it identify which distressed borrowers are worth supporting, so its continuation credit should flow to borrowers poised to recover. We track the subsequent valuation of the borrower's pre-existing debt—the new loan itself does not enter the outcome, so rescue cannot mechanically improve it—and find no trace of such selection: over the next one to three years, the existing debt of rescued borrowers does not appreciate relative to that of other impaired borrowers, and it is 17 percentage points more likely to be on non-accrual after one year. The evidence is therefore inconsistent with advantageous selection into rescue lending.

We next examine what drives managers to direct capital to impaired incumbents. This mechanism requires two conditions to coincide: a young fund provides lending capacity financed by investors in the new fund, while an above-hurdle seasoned fund aligns the manager with investors in the existing fund. Once the hurdle is met, both the manager and the existing fund's investors benefit from preserving the older fund's return; the manager also benefits when that track record supports future fundraising. The mechanism therefore predicts rescue specifically when lending capacity overlaps with an above-hurdle seasoned fund.

The data follow this prediction. Managers without a young fund essentially never rescue. Among managers with a young fund, the rescue rate is 9.9 percent when an above-hurdle seasoned fund is also on the roster—roughly twice the rate otherwise. With manager and year fixed effects and the number of impaired borrowers available to rescue held fixed, the interaction between young-fund capacity and above-hurdle status is 2.6 percentage points, about 40 percent of the mean

rescue rate, while the young-fund effect alone falls to zero. The same interaction does not predict lending to new borrowers. Rescue therefore concentrates not in periods of broad deployment, but when lending capacity overlaps with an older track record worth preserving.

Young-fund status directly captures fresh lending capacity, but it is an equilibrium outcome, also reflecting the manager's demand for capital. We therefore use investor mandates as a more exogenous, investor-side source of funding-supply variation. Institutional investors announce mandates to allocate capital to private credit. For each manager-year, we count mandates announced over the prior one or three years by investors that first committed to the manager's funds at least five years earlier. This predetermined relationship network maps investors' allocation decisions into manager-specific funding exposure without being shaped by current fundraising. The exposure predicts subsequent fund closings, whereas future mandates and mandates targeting other asset classes do not. A one-standard-deviation increase in mandate exposure raises rescue by 2.5 to 2.6 percentage points specifically when the manager has an above-hurdle seasoned fund. The realized capacity captured by a young fund and the investor-side funding shift captured by mandates therefore yield the same pattern: funding supply increases rescue when the manager is aligned with investors in the existing fund to protect its performance.

Rescue lending, in turn, preserves reported fund performance. Estimating this relation directly by regressing fund IRR on realized rescue activity would understate the effect because rescue occurs precisely when portfolio conditions deteriorate. We therefore combine the preceding first-stage evidence with a reduced-form test. The first stage establishes that funding capacity interacting with above-hurdle status increases rescue; the reduced form asks whether the same interaction predicts subsequent changes in seasoned-fund IRR. Controlling for manager and year effects and initial-performance deciles, above-hurdle funds whose manager also has a young fund report IRR changes that are 0.9 to 1.4 percentage points more favorable over the next one to three years. Substituting investor-mandate exposure for young-fund status yields the same pattern. The design draws variation from manager-level funding capacity rather than the condition of any individual borrower. Taken together, the first-stage and reduced-form estimates indicate that rescue induced

by incentive-relevant funding capacity improves reported IRR under both capacity measures.

Finally, stronger reported performance predicts the capital available to the next fund generation. Among private-credit funds raised between 2000 and 2017, a higher final reported IRR makes a successor fund more likely, and each percentage point of IRR is associated with 1.6 to 1.9 percent greater successor-fund growth—roughly twice our estimate for private-equity funds, the asset class in which the flow–performance relation is best established (Kaplan and Schoar, 2005; Chung, Sensoy, Stern and Weisbach, 2012). At the investor level, an existing investor's decision to recommit responds to the manager's old-fund IRR specifically, not its young funds' (Hochberg, Ljungqvist and Vissing-Jørgensen, 2014). These relations are predictive rather than causal, but they establish that the track record managers have an incentive to preserve also predicts their future lending capacity.

Together, the findings trace the self-reinforcing loop created when the same manager allocates credit to borrowers and raises capital from investors. A younger fund provides lending capacity, while an above-hurdle seasoned fund aligns the manager with investors in the existing fund to preserve its performance. Rescue preserves reported performance; reported performance predicts capital raised for the next fund generation; and that capital renews the capacity for lending and future rescue. The loop therefore links market expansion to risk accumulation.

We then formalize the loop in a model of overlapping fund generations and calibrate it to our estimates. At every date a manager operates a young fund, whose committed capital finances ordinary loans and rescue loans, and a seasoned fund, whose reported return determines the manager's carried interest when the fund is above its hurdle and, through investors' flow–performance response, the capital the next fund raises. Each rescue dollar keeps a dollar of impaired exposure current and out of the seasoned fund's report, at a cost that rises with rescue intensity. The funding response has two parts: a class-level part, through which stronger reported returns draw capital into private credit as a whole, and a relative part, through which a manager's report draws capital away from rivals. The two funding loadings come from the pension-commitment and successor-fund responses, and three structural parameters—the rescue-cost curvature, the deployment wedge, and the

loading of unrecognized distress on the reported return—are estimated by requiring simulated data to reproduce the capacity × above-hurdle interactions in rescue and in subsequent seasoned-fund IRR, recognized distress, and young funds' deployment ratio.

The calibrated model reproduces the sign and magnitude of the two interactions and delivers three conclusions. First, rescue is cheap and pervasive: managers keep 57 percent of newly distressed exposure current each period, so concealed distress is 8 percent of the seasoned book against 6 percent recognized, and reported returns understate distress by more than half. Second, the fund-level flow–performance response is mostly reallocation. The class-level funding response is one-fifth of the response a single manager obtains, so four-fifths of the capital a manager attracts by outperforming comes from rivals, and the feedback through rescue, reported return, and fundraising amplifies one manager's capital response by 29 percent but the industry's by only 5 percent. Third, the mechanism's aggregate importance lies less in how much capital it attracts than in what that capital sustains: about half of the exposure that rescue conceals passes into the next fund generation's book, where keeping it current requires the next young fund's capital, so reported returns are a lagging measure of credit quality whose accuracy depends on continued inflows.

Self-reinforcing growth changes the nature of private-credit risk. Closed-end capital commitments, high equity funding, and limited maturity mismatch make private-credit funds less exposed than banks to runs and forced liquidation (Matvos, Piskorski and Seru, 2026; Chernenko and Scharfstein, 2025). Lower run risk, however, does not eliminate feedback risk. While fundraising continues, capital remains available to support impaired borrowers, delay loss recognition, and sustain reported performance and intermediary scale even as impairment accumulates. A slowdown in investor commitments can reverse the loop: rescue capacity contracts, latent losses surface, reported performance weakens, and fundraising becomes still harder. Vulnerability shifts from runnable liabilities and regulatory capital toward interim valuation, performance-sensitive capital flows, and the timing of credit resolution. The incidence of the deferred losses gives this shift its policy weight: private-credit capital comes largely from pension funds and insurers, so losses that surface late ultimately reach retirees and insurance policyholders—and because rescue finances an

older fund's borrowers with a younger fund's capital, the burden tilts toward the investor cohorts that committed after the losses were embedded, precisely the cohorts attracted by preserved track records. The 2026 proposal to admit private assets into defined-contribution plans would extend this exposure directly to household retirement savers. Because banks and institutional investors also remain connected to private-credit vehicles, these dynamics need not stay contained within the funds themselves.

Existing explanations of private-credit growth emphasize shifts in borrower demand and investor supply. One account attributes the shift from banks to private lenders to tighter bank regulation and related supervisory costs (Chernenko, Erel and Prilmeier, 2022; Davydiuk et al., 2024a; Cai and Haque, 2024; Jang et al., 2025a; Chernenko, Ialenti and Scharfstein, 2026), as well as to private lenders' speed and contractual flexibility (Cai and Haque, 2024). A complementary account emphasizes investor supply, as low interest rates drew institutional capital toward illiquid assets offering high reported returns (Block, Jang, Kaplan and Schulze, 2024; International Monetary Fund, 2024). In these accounts, external conditions generate private-credit growth. Our findings show that the story also runs in reverse: lending decisions can shape reported fund performance and future fundraising, allowing lending to help generate the capital that finances subsequent lending. Private credit can therefore grow partly from within.

The mechanism also links delegated asset management to credit allocation. Models and evidence show that investor capital follows reported performance (Berk and Green, 2004; Kaplan and Schoar, 2005), and that investor discipline governs how direct lenders raise and deploy capital (Davydiuk, Marchuk and Rosen, 2024b). Here, that flow–performance relation reaches back into the treatment of impaired borrowers: fundraising-relevant fund states predict how a manager treats its impaired incumbents, while reported returns predict the manager's future lending capacity.

The comparison with bank evergreening clarifies what is distinctive about private credit. Existing work emphasizes loss recognition and regulatory capital as the reasons weak banks continue lending to weak borrowers (Peek and Rosengren, 2005; Caballero, Hoshi and Kashyap, 2008; Tantri, 2021). Private-credit funds mute these channels, but their reliance on future fundraising shifts the

relevant state variable from bank capital to the manager's track record and ability to raise the next fund. This shift can make rescue expansionary during an inflow cycle and expose accumulated impairment when that cycle turns.

2 Institutional Background and Conceptual Framework

2.1 Private-Credit Fund Organization

Private credit is supplied through closed-end finite-life private funds, business development companies (BDCs), and evergreen vehicles (Ivashina, 2026). We focus on finite-life funds, the predominant organizational form, in which institutional investors commit capital for seven to ten years without redemption rights (Board of Governors of the Federal Reserve System, 2026). Our empirical exercise also uses regulatory filings of BDCs, which provide the debt valuations used to measure borrower credit quality.

A finite-life fund is a partnership between a general partner (GP), which selects and manages investments, and limited partners (LPs), which supply most of the capital. Pension funds account for roughly half of LP capital commitments in our data; Appendix Figure B.1 shows that pensions, insurers, and endowments together account for more than four-fifths. The GP typically receives an annual management fee of about 1.5 percent of the fund's committed (or deployed) capital and carried interest of about 15 percent of profits after the fund performance exceeds a contractual hurdle (Erel, Flanagan and Weisbach, 2024; Institutional Limited Partners Association, 2019).

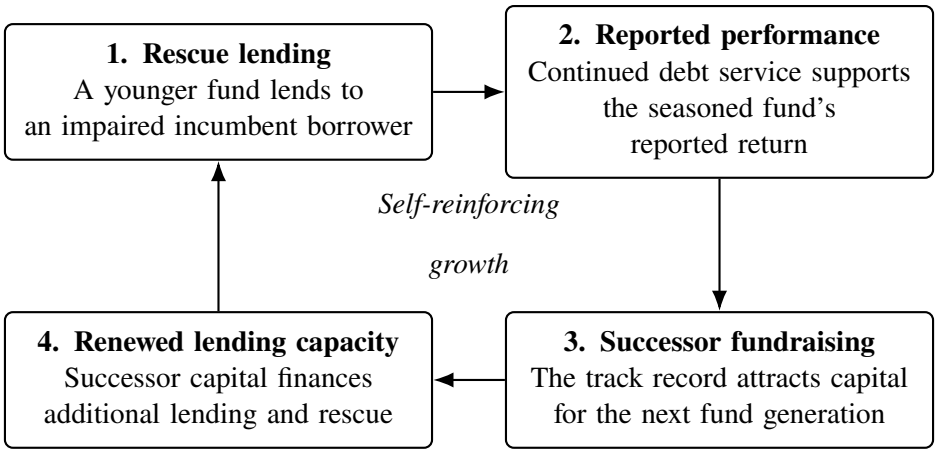
Fundraising for a successor typically begins before an earlier fund is liquidated, so a GP often operates overlapping generations of funds. Each fund remains legally separate, has its own LPs, and reports its own performance. Overlapping funds may also invest in the same company. ILPA defines this practice as *cross-fund investing*: a current fund participates in financing a company held by an earlier fund (Institutional Limited Partners Association, 2019). Common management across these separate funds and investor cohorts provides the institutional setting for the framework below.

2.2 Self-Reinforcing Growth across Fund Generations

The framework asks how common management of overlapping funds connects lending and fundraising. A seasoned fund supplies the track record used to raise a successor, while a younger fund supplies capacity for new loans. The GP benefits from the seasoned fund's performance through carried interest when the fund is above its hurdle and through income from successor funds when a stronger track record attracts successor capital. Therefore, with both a well-performing seasoned fund and a young fund under management, the GP may have both the incentive and the capacity to improve the reported performance of a seasoned fund.

Consider a borrower held by the seasoned fund with impaired credit quality. A new loan from the young fund may allow the borrower to continue servicing its existing debt, supporting the seasoned fund's reported performance without necessarily resolving the underlying credit risk. This financing may be an efficient workout if it helps the firm meet temporary liquidity needs and avoids costly default. If it instead delays the recognition of fundamental insolvency, the older fund and its investors receive the performance benefit while the younger fund's investors bear the new exposure. The GP shares the benefit when the seasoned fund is above its hurdle (through carry) or if its performance attracts LP capital (through future income). Alignment of interests within the older fund can therefore create a conflict across fund generations.

These incentives generate the feedback loop below. Rescue lending can support the seasoned fund's reported performance; that performance helps attract successor capital; and successor capital provides the capacity to lend, including to other impaired portfolio borrowers.



The framework yields four empirical predictions. First, cross-fund rescue requires the GP to use a young fund to lend again to an impaired borrower held by a seasoned fund; new credit to impaired borrowers should therefore disproportionately come from existing lenders. Second, rescue should concentrate when a GP has both lending capacity and a seasoned fund whose reported performance is valuable to preserve. Third, the same combination of fund generations that lead to rescue lending should support the seasoned fund's subsequent reported performance. Finally, stronger reported performance should predict successor-fund formation and size, renewing the GP's lending capacity. The empirical analysis follows this order. Section 7 then formalizes the loop as a model of overlapping fund generations and calibrates it to the estimates.

This feedback makes market expansion and credit-risk accumulation joint outcomes. Continued fundraising supplies capital that can keep impaired borrowers current, while the reported performance supported by that lending attracts additional capital. If fundraising slows, the mechanism can operate in reverse: rescue capacity contracts, unresolved losses become harder to defer, and weaker reported performance further reduces the capital available for lending.

3 Data and Measurement

Our empirical analysis combines four datasets: private credit transactions, BDC portfolio holdings, private credit GPs and funds, and LP commitments and mandates. Each dataset supplies the information needed for a different part of the feedback loop. We describe the sources and

principal measures below. Appendix A provides the detailed sample screens, linking procedures, and variable construction, and Table 1 reports summary statistics.

3.1 Data Sources

Private credit transactions. Our sample of private credit loans comes from PitchBook corporate debts database. The data identify the borrower, origination date, loan terms, purpose, and participating lenders at the debt tranche level. These fields allow us to construct each borrower's lending relationships and determine whether a new loan comes from an existing or a new lender. We define a loan deal as private credit if it satisfies four conditions: (a) the deal is not syndicated and does not include bonds; (b) at least one debt tranche is a term loan or classified as direct lending; (c) all term loan tranches have non-missing lender information; and (d) at least one lender is a nonbank institution (i.e., not commercial banks, investment banks, and merchant banks). The sample covers U.S. dollar-denominated, voluntary-purpose loans (i.e., excluding purposes such as buyouts, acquisitions, and recapitalization) to U.S. borrowers from 2010 through 2025. Figure 2 summarizes loan pricing over the sample period.

BDC portfolio holdings. We use Pitchbook's quarterly portfolio holdings of BDCs to measure borrower credit quality. Private credit borrowers are generally private and unrated, and regulatory disclosures of BDCs provide a unique window to measure the quality of a borrower's existing debt before a new loan and to track the same debt afterward. This dataset provides holding-level information on investment type, fair values, investment costs, outstanding principal, and non-accrual status. We link these holdings to private credit borrowers via PitchBook's Company ID and use them to construct borrower-level measures of credit impairment.

GPs and funds. PitchBook's fund records report each fund's manager (GP), strategy, vintage, size, sequence, and net internal rate of return (IRR). Fund vintages and manager identifiers allow us to measure overlapping fund generations; reported IRRs measure the track record available to investors; and fund sequences and sizes identify successor fundraising. Figure 3 plots aggregate private-credit fundraising and reported fund returns. Our fund dataset includes all debt funds (e.g.,

direct lending, mezzanine, distressed debt) and private equity funds (e.g., buyout, growth/expansion, restructuring).

LPs, commitments, and mandates. PitchBook also reports LP commitments to individual funds and these institutional investors' announced mandates to allocate capital to private funds. Commitments identify the investors behind each GP and the duration of those relationships. Mandates provide a measure of investor-side changes in the funding available to GPs. We complement PitchBook with the Public Plans Database (PPD), which provides an annual panel of U.S. public pensions with each pension's characteristics and asset class-level allocation targets. Appendix Figure B.1 describes the composition of investors in private-credit funds.

3.2 Measurement

Impaired borrowers. Identifying rescue lending requires observing both the lender on a new loan and the borrower's condition before origination. We link PitchBook private credit transactions to BDC portfolio holdings through a common PitchBook borrower identifier. For each transaction, we use the borrower's most recent BDC debt valuation reported strictly before origination and within the preceding four quarters. Extending the approach in Davydiuk et al. (2025), for each BDC's holding position we first compute the ratio of fair value to investment cost (*FV/Cost*) and then aggregate to the borrower-quarter level using cost-weighted averages across BDC positions.[4] We define a borrower as *impaired* if it has an FV/Cost ratio below 0.90 before the new loan.

The link is at the borrower rather than the loan level: some BDC must hold the borrower's debt, but neither the new loan nor its lender must appear in a BDC portfolio. Pre-origination coverage rises from 21 percent of private-credit debt instruments in 2010 to 61 percent in 2024, motivating the 2010 starting date for the valuation-conditioned analysis (Appendix Figure B.2, Panel A). By 2024, covered instruments represent 57 percent of private-credit debt volume (Appendix A.4). Moreover,

[4]Specifically, we exclude equity, credit commitments (revolvers and delayed-draw term loans), as well as positions with non-positive cost and newly originated loans, which mechanically trade near par. We use cost rather than principal because private credit loans are typically held at cost unless they are impaired; principal would introduce noise from original issue discounts (OID) and other contractual features that are not informative about current credit quality. We winsorize loan-level FV/cost at the 1st and 99th percentiles to reduce the impact of outliers.

42 percent of debt instruments issued to borrowers with a same-year BDC valuation never appear in BDC holdings (Appendix Figure B.2, Panel B). The linked sample therefore uses BDC valuations to measure the condition of borrowers in a substantially broader set of private-credit transactions.

Existing lenders and rescue lending. A lender is an *existing lender* in year (t) if it participated in the borrower's most recent lending year before (t). A *rescue loan* is new financing from an existing lender to an impaired borrower. This definition describes an observable lending decision without assigning its welfare consequences. A rescue loan may be an efficient workout; we use *evergreening* only for the stronger interpretation that continuation occurs without commensurate economic recovery.

Fund status. We use fund age to measure a manager's capacity to extend new credit and reported performance to measure its incentive to protect an existing track record. A *young fund* is one to four years old and represents capital available for deployment. A *seasoned fund* is at least five years old. We classify a seasoned fund as *above hurdle* when its reported net IRR is at least 10 percent, where carried interest is plausibly in the money. Because contractual hurdles are not observed, the 10 percent cutoff is an empirical proxy rather than a fund-specific contractual threshold. These measures identify manager-years in which lending capacity overlaps with a seasoned track record worth preserving.

Fundraising and investor exposure. We measure successor fundraising using the sequence and size of funds raised by the same manager. As an investor-side measure of funding conditions, we count private-credit mandates announced over the preceding one or three years by LPs whose first commitment to the manager occurred at least five years earlier. Restricting the measure to longstanding LP relationships prevents current fundraising from determining which investors enter the manager's exposure measure. Appendix Table B.1 shows that mandate exposure predicts subsequent fund closings, whereas future mandates and mandates outside private credit do not.

Table 1 describes the three principal estimation samples. The borrower-year panel contains

27,161 observations from 2010 through 2024. Nine percent of borrower-years are impaired, 6 percent include financing from an existing lender, and 15 percent include financing from a new lender. The manager-year panel contains 4,534 observations from 2010 through 2025. Rescue lending occurs in 3 percent of manager-years, compared with new-borrower lending in 43 percent, and 76 percent of manager-years include a young fund. The seasoned-fund panel contains 6,397 observations; mean reported net IRR is 9.65 percent, and 40 percent of observations are above the 10 percent hurdle proxy.

4 Rescue Lending to Impaired Borrowers

The self-reinforcing growth framework begins with cross-fund rescue: a manager uses capital raised for a younger fund to extend new credit to an existing borrower whose debt in a seasoned fund has become impaired (a "credit-impaired" borrower). We begin the empirical analysis by examining this rescue-lending behavior. We first document the raw relation between borrower impairment and existing-lender financing and then compare potential lenders facing the same borrower in the same year. Finally, we test the advantageous-selection explanation—that existing lenders reserve rescue for impaired borrowers whose prospects are better than their valuations suggest.

4.1 Repeat Financing and Borrower Impairment

Cross-fund rescue is feasible only when a manager operates a young fund and a seasoned fund at the same time. Figure 1 shows that this configuration is the norm rather than the exception. Panel A traces the fund lifecycles of four representative managers, one from each quartile of private-credit capital raised. Each raises a successor every two to three years, so a fund in its investment period almost always coexists with a predecessor still in its harvest period. Panel B shows that the overlap is pervasive and growing: managers with four or more active funds held about two-thirds of private-credit capital in 2010 and nearly 90 percent by 2025, whereas managers with a single active fund held less than 3 percent. When a borrower held by a seasoned fund becomes impaired, its manager therefore almost always has a younger fund with capital to deploy.

Figures 2–4 show how fundraising, loan pricing, reported returns, and repeat lending evolved as private credit expanded. Annual fundraising rises from about $50 billion in 2010 to more than $400 billion in recent vintages, increasing managers' capacity to originate loans (Figure 3). Capital inflows, however, also intensify competition for investment opportunities and should depress returns (Gompers and Lerner, 2000; Pástor and Stambaugh, 2012). Consistent with greater competition for borrowers, average spreads on new loans fall from about 800 to slightly above 500 basis points (Figure 2). Despite this spread compression, mean reported net IRR remains between roughly 8 and 12 percent (Figure 3). Stable reported returns therefore should not be interpreted as evidence that investment opportunities remained abundant relative to the capital entering the market.

Over the same period, the share of borrowers receiving a new loan from their most recent prior lender rises from about 5 percent to more than 25 percent (Figure 4, Panel A). The most immediate interpretation is profit-maximizing relationship lending. As competition for new borrowers intensifies, relationship-specific information and borrower switching costs may make repeat lending more attractive to the younger fund. Under this interpretation, managers lend again because the existing relationship raises the expected return on the new loan, not because that loan protects the seasoned fund's existing exposure.

The cross-sectional pattern is difficult to explain by relationship benefits to the new loan alone. Figure 4, Panel B, shows that repeat financing becomes more prevalent as the value of a borrower's existing debt deteriorates. Existing lenders provide new financing to 43 percent of borrowers whose debt is valued below 80 cents per dollar of cost and 39 percent of those valued between 80 and 85 cents, compared with 26 to 29 percent for borrowers valued near cost. Relationship-specific information or switching costs may increase the expected return on a new loan, but neither provides an immediate reason for repeat lending to concentrate among the most impaired borrowers. These are precisely the borrowers for whom a new loan can also protect the lender's existing claim. The raw pattern leaves two alternative explanations. First, impaired borrowers may differ in credit demand, industry, or available lenders, while lenders with greater capacity may lend more generally. Second, existing lenders may use private information to identify impaired borrowers whose prospects are

better than their valuations suggest. The next two subsections examine these explanations in turn.

4.2 Existing Relationships and New Lending

We first test whether borrower and lender heterogeneity explains the raw pattern. Borrower-year comparisons hold fixed the borrower's condition, credit demand, industry, and general attractiveness in that year, while lender-year comparisons hold fixed each lender's annual capacity and risk appetite. Within this design, we ask whether the borrower's existing lender is especially likely to provide new credit when the borrower is impaired. We construct the comparison set in each year by forming the Cartesian product of active borrowers and active lenders. We retain borrower-years with a pre-origination FV/Cost ratio and deals involving no more than five lenders, yielding approximately 3.4 million potential borrower–lender pairs. We estimate:

$$\text{Match}_{ijt} = \beta_1 \,\text{Existing}_{ijt} + \beta_2 \,\text{Existing}_{ijt} \times \text{Impaired}_{it} + \mu_{it} + \nu_{jt} + \varepsilon_{ijt}, \tag{1}$$

where Match_{ijt} indicates that borrower i and lender j complete a transaction in year t, Existing_{ijt} identifies the borrower's most recent lender, and Impaired_{it} indicates that the borrower's pre-origination FV/Cost ratio is below 0.90. The terms μ_{it} and ν_{jt} denote borrower-year and lender-year fixed effects, respectively. The coefficient β_2 measures whether an existing relationship is more strongly associated with a new loan when the borrower is impaired, conditional on these borrower-year and lender-year components. The design does not treat impairment as exogenous and should not be read as the causal effect of becoming impaired.

An existing relationship is more strongly associated with a new loan when the borrower is impaired. With borrower, lender, and year fixed effects, an existing relationship is associated with a 27.1-percentage-point higher annual match probability for a non-impaired borrower, and this association is 8.6 percentage points stronger for an impaired borrower (Table 2, column 1). The more stringent borrower-year and lender-year comparison produces nearly identical estimates: 26.9 and 8.8 percentage points, respectively (column 2). The existing-relationship coefficient is

therefore about one-third larger among impaired borrowers.

The result is similar when borrower condition is measured continuously. In columns 3 and 4, *Impairment* equals $1 - \text{FV}/\text{Cost}$, bounded at zero for debt valued above par. A ten-percentage-point decline below par is associated with a 2.7-percentage-point stronger relation between an existing relationship and a new loan. Existing lenders therefore become increasingly likely, relative to other potential lenders, to provide the new loan as impairment deepens; the result is not an artifact of the 90-cent cutoff.

Taken together, the matching estimates show that impaired borrowers disproportionately receive new credit from their existing lenders even when potential lenders are compared for the same borrower in the same year. The raw pattern therefore does not reflect greater borrower demand or general lender capacity; it is a relationship-specific pattern in which impairment is associated with a greater propensity for the existing lender to provide new credit.

4.3 No Advantageous Selection into Rescue

We next test the second alternative: existing lenders use private information to select impaired borrowers whose prospects are better than their valuations suggest. If this advantageous-selection account explains repeat financing, the existing debt of borrowers selected by existing lenders should subsequently outperform that of other impaired borrowers. We test this prediction using the later performance of debt already outstanding before the new loan. Excluding the new loan is important: a newly issued loan valued near par could otherwise raise the borrower's portfolio-level FV/Cost ratio mechanically, even if the existing debt did not improve.

For borrower i, year t, and horizon $h \in \{1, 2, 3\}$, we estimate

$$\begin{aligned} Y_{i,t+h} = \beta_1 \,\text{Impaired}_{it} + \beta_2 \,\text{Impaired}_{it} \times \text{BorrowExisting}_{it} + \beta_3 \,\text{Impaired}_{it} \times \text{BorrowNew}_{it} \\ + \gamma_1 \,\text{BorrowExisting}_{it} + \gamma_2 \,\text{BorrowNew}_{it} + \gamma_3 \,\text{Maturity}_{it} + \delta_{s(i)t} + \varepsilon_{it}. \end{aligned} \tag{2}$$

The outcome $Y_{i,t+h}$ is either the change in the preexisting debt's FV/Cost ratio from $t-1$ to $t+h$ or an indicator that the debt is on non-accrual at $t+h$. $\text{BorrowExisting}_{it}$ and BorrowNew_{it} indicate new

financing from the incumbent and from another lender, respectively. The specification controls for the remaining maturity of the existing debt and includes industry-year fixed effects; standard errors are clustered by borrower. Advantageous selection predicts a positive β_2 for changes in FV/Cost and a negative β_2 for non-accrual.

The existing debt of rescued borrowers does not subsequently outperform. The interaction between *Impaired* and *Borrow from Existing* is negative for changes in FV/Cost at all three horizons: −0.023, −0.050, and −0.133, respectively (Table 3, columns 1–3). After one year, the existing debt of rescued borrowers is also 17.0 percentage points more likely to be on non-accrual (column 4). By contrast, impaired borrowers financed by a new lender are 8.8 percentage points less likely to be on non-accrual, producing a 25.8-percentage-point difference between incumbent and outside financing.

The appendix addresses two features of this comparison. Appendix Table B.2 restricts the sample to borrower-years financed by exactly one lender type. Within this sample, incumbent financing is associated with a 19.4-percentage-point lower change in FV/Cost after three years and a 28.8-percentage-point higher probability of non-accrual after one year. Appendix Table B.3 separately reports whether the preexisting debt disappears from BDC holdings. Because disappearance may reflect repayment, refinancing, sale, write-off, or lost coverage, we do not classify attrition as recovery.

Taken together, the results establish rescue lending as a relationship-specific response to borrower impairment. As a borrower's condition deteriorates, its existing lender becomes increasingly likely, relative to other potential lenders, to provide new credit, even after accounting for borrower-year demand and lender-year capacity. The preexisting debt of rescued borrowers does not subsequently outperform, making the evidence inconsistent with advantageous selection based on lender private information.[5]

[5]Note that this outcome comparison is a selection test, not an estimate of rescue's effect on borrower outcomes: rescued borrowers might have performed worse without additional financing, and some rescue loans may be efficient workouts.

5 Manager Incentives and Rescue Lending

This section studies how managers' incentives and lending capacity shape rescue lending. Newly raised capital supplies the capacity to extend rescue credit. Managers' incentive to rescue depends on the performance state of the seasoned fund. When that fund is above the hurdle, the manager is aligned with its investors in preserving its reported return: the investors benefit from higher performance, while the manager retains carried interest and protects the track record used in future fundraising. Rescue should therefore be most common when lending capacity coincides with an above-hurdle seasoned fund. By contrast, if relationship-specific advantages, switching frictions, or efficient workouts make repeat lending profitable for the younger fund on its own, rescue may depend on capacity and available borrowers but should not depend on the seasoned fund's performance state.

5.1 Young-Fund Capacity and Rescue Lending

We begin by examining whether the availability of recently raised capital predicts rescue lending, particularly when the manager also operates an above-hurdle seasoned fund. We proxy for recently raised capital with an indicator for whether the manager operates a fund aged one to four years, when committed capital is typically still being deployed.

Figure 5 shows the pattern of rescue lending across lenders. Among managers with a young fund, rescue occurs in roughly 4 percent of manager-years throughout the below-hurdle range, jumps to 14 percent just above the hurdle, and settles at 8 percent at higher IRRs. In sharp contrast, managers without a young fund record no rescue at any performance level.

Table 4, Panel A reports the underlying cells, including the 216 manager-years without a young fund — none with a rescue loan — and the 5.8 percent rescue rate among managers with a young fund but no seasoned reporting fund, the omitted state in the regressions below. Two features of this raw pattern discipline the design that follows: rescue requires a young fund, and, given one, it responds to the seasoned fund's performance state at the hurdle.

The raw comparison may reflect persistent differences in managers' lending strategies or changes

in the number of borrowers available to rescue. We therefore compare the same manager in years with different combinations of young-fund availability and seasoned-fund performance, while controlling for its rescue opportunities. For manager g in year t, we estimate

$$\begin{aligned}\text{Rescue}_{gt} = \beta_1 \,\text{Capacity}_{gt} \times \text{AboveHurdle}_{gt} + \beta_2 \,\text{Capacity}_{gt} \\ + \beta_3 \,\text{AboveHurdle}_{gt} + \beta_4 \,\text{BelowHurdle}_{gt} + \gamma' X_{gt} + \alpha_g + \tau_t + \varepsilon_{gt}.\end{aligned} \tag{3}$$

The outcome indicates that the manager makes at least one rescue loan during the year. In columns 1 and 2 of Panel B, Capacity_{gt} indicates a young fund. AboveHurdle_{gt} and BelowHurdle_{gt} classify seasoned funds using the 10-percent reported net IRR threshold, with no seasoned reporting fund as the omitted state. The controls X_{gt} include the numbers of impaired and total incumbent borrowers entering the year and cumulative capital raised. Manager fixed effects absorb persistent lending strategies, while year fixed effects absorb aggregate credit and fundraising conditions. The coefficient β_1 therefore asks whether the same manager is more likely to rescue when lending capacity overlaps with an above-hurdle seasoned fund, holding fixed the number of borrowers available to rescue.

Rescue increases when young-fund capacity overlaps with an above-hurdle seasoned fund. The unconditional interaction is 5.1 percentage points (Table 4, Panel B, column 1). With manager and year fixed effects and the opportunity-set controls, the interaction is 2.6 percentage points, or about 40 percent of the 6.8-percent mean rescue rate (column 2), while the young-fund main effect falls to zero. Capacity alone therefore does not explain rescue. Rescue increases specifically when the same manager also has a seasoned fund whose reported performance is valuable to preserve, consistent with the cross-fund incentive.

5.2 Investor Mandates and Rescue Lending

A young fund measures realized lending capacity, but its presence is an equilibrium outcome that reflects both investor funding supply and the manager's demand for capital. Investor mandates provide a complementary source of variation originating with investors and less directly tied to the

manager's current demand for capital. Institutional investors announce programs to allocate capital to private credit, and PitchBook records 683 such announcements between 2008 and 2025. The mandate design asks whether variation in funding opportunities originating with investors produces the same incentive-specific rescue pattern as young-fund status.

We measure a manager's exposure using mandates announced by its long-standing investors—LPs that first committed to the manager at least five years earlier. For each manager-year, we count these investors' private-credit mandates over the prior year or prior three years and standardize the count. Fixing the relationship network before the mandate announcement prevents current fundraising from determining which investors generate the exposure. A one-standard-deviation increase in mandate exposure predicts a 2.1- to 2.4-percentage-point increase in the probability that the manager closes a debt fund. Future private-credit mandates and mandates over the same lookback windows targeting other asset classes do not predict current fund closings (Appendix Table B.1). These results support interpreting mandate exposure as an investor-side measure of funding conditions.

Mandate exposure produces the same interaction as young-fund status. A one-standard-deviation increase in exposure is associated with 2.5 to 2.6 percentage points more rescue when the manager has an above-hurdle seasoned fund (Table 4, Panel B, columns 3 and 4). Exposure alone predicts, if anything, less rescue. Easier investor funding conditions therefore predict rescue only when the manager also has a seasoned track record worth preserving.

5.3 New-Borrower Lending Placebo

The interaction between capacity and seasoned-fund performance could reflect broad capital deployment rather than an incentive specific to rescue. We test this alternative using lending to new borrowers, where a new loan cannot protect an existing claim. Young-fund capacity is associated with 10.5 percentage points more new-borrower lending in the specification with controls and fixed effects (Appendix Table B.4), but none of the capacity-times-above-hurdle interactions is statistically distinguishable from zero. Capacity therefore increases ordinary deployment, whereas

the seasoned fund's performance state predicts how capacity is directed toward impaired existing borrowers. This distinction is inconsistent with a capacity-only explanation of rescue lending.

Taken together, the results show that managers direct lending capacity toward impaired existing borrowers specifically when they also have a seasoned fund whose reported performance is valuable to preserve. The pattern appears with both realized young-fund capacity and investor-side mandate exposure, but not in lending to new borrowers. Rescue thus reflects the joint presence of capacity in one fund generation and performance incentives in another, rather than general capital deployment.

6 Reported Performance and Successor Fundraising

Finally, we examine whether rescue lending affects reported performance and future fundraising. The framework suggests that rescue may preserve a seasoned fund's reported return by keeping an impaired borrower current. If investors reward the resulting track record, the manager may raise more successor capital. That capital could renew its capacity to support impaired borrowers, delaying recognition of accumulated credit risk and making private-credit growth self-reinforcing. We examine each link in turn.

6.1 Rescue Lending and Reported Performance

We first ask whether the combination of lending capacity and an above-hurdle seasoned fund shown in Section 5 to increase rescue also predicts more favorable changes in the seasoned fund's reported IRR. A direct regression of reported IRR on realized rescue would compare rescue years with other years even though managers rescue precisely when portfolio conditions deteriorate. This negative selection would bias the estimate against finding that rescue supports performance. We therefore trace the relation through the combination of lending capacity and seasoned-fund performance that predicts rescue. Section 5 shows that rescue increases when lending capacity coincides with an above-hurdle seasoned fund; this is the first stage. We now ask whether the same combination predicts a more favorable subsequent change in the seasoned fund's reported IRR; this is the reduced form. Using the same capacity–above-hurdle interaction in both tests avoids

conditioning directly on rescue activity that is selected on portfolio distress.

For seasoned fund f, managed by manager g, we estimate

$$\begin{aligned}\Delta\text{IRR}_{f,t\to t+h} = \beta_1\,\text{Capacity}_{gt} \times \text{AboveHurdle}_{ft} + \beta_2\,\text{Capacity}_{gt} \\ + \beta_3\,\text{AboveHurdle}_{ft} + \gamma' X_{ft} + \kappa_{d(f,t)} + \alpha_g + \tau_t + \varepsilon_{ft},\end{aligned} \tag{4}$$

where $h \in \{1, 2, 3\}$. Capacity is measured by a young fund in columns 1–3 and related-LP mandate exposure in columns 4–6. AboveHurdle_{ft} indicates that the seasoned fund's reported net IRR is at least 10 percent, and X_{ft} contains fund size and age. Current-IRR-decile fixed effects $\kappa_{d(f,t)}$ compare funds beginning in the same performance range, manager fixed effects compare funds within the same organization, and year fixed effects absorb common movements in reported returns. The coefficient β_1 therefore measures whether an above-hurdle fund's IRR evolves differently when its manager has greater lending capacity, conditional on its initial performance.

Above-hurdle seasoned funds report more favorable subsequent returns when the manager has greater lending capacity. When capacity is measured by a young fund, their reported IRR changes are 0.887, 1.140, and 1.363 percentage points more favorable over the next one, two, and three years, respectively (Table 5, columns 1–3). Related-LP mandate exposure yields the same pattern: the corresponding estimates are 0.172, 0.251, and 0.339 percentage points per standard deviation (columns 4–6).

Read together, the first-stage and reduced-form estimates indicate that rescue induced by incentive-relevant lending capacity improves reported IRR under both capacity measures. This comparison is informative because manager-level funding conditions are not determined by the condition of any single borrower. The estimates do not identify a per-loan treatment effect: loan-to-fund mappings are unavailable, and the same capacity–above-hurdle interaction could affect reported IRR through valuation or exit decisions in addition to rescue. The evidence instead identifies rescue as the lending component of a broader performance-preservation response.

6.2 Asset-Class Performance and Investor Allocations

Performance preservation renews lending capacity only if investors reward reported returns with capital. We examine this response at two levels: the institutional investors that supply private credit's capital, and the manager raising a successor fund. The two levels answer different questions. Investor-level evidence shows whether reported performance draws capital into private credit as an asset class; manager-level evidence shows whether one manager's record draws capital away from another's. Only the first margin expands aggregate lending, a distinction that Section 7 makes precise. We begin with investors.

We study U.S. public pensions, which account for roughly half of the LP capital committed to private-credit funds in our data. The Public Plans Database (PPD) provides an annual panel of state and local plans with their allocation targets and characteristics. We match 161 of its investment pools, holding 92 percent of PPD assets, to their PitchBook commitment histories. We measure investor demand through policy targets and dated commitments rather than through reported allocation weights, because most plans do not itemize private credit in the weights they report. Appendix A.3 details the match and the sample construction.

Table 6 follows the allocation process from the top down. Panel A asks whether pensions revise their strategic allocation toward private markets after private assets outperform public ones. The unit is a pension-year. The dependent variable is the change from year t to $t + 1$ in the pension's private-markets policy target: the sum of its private equity, private credit, and real estate target weights as a share of plan assets. The regressor is the lagged private-minus-public composite return over the window ending at the pension's fiscal year end, and pension fixed effects allow each pension its own trend in the target. A ten-percentage-point higher private-minus-public return raises the next year's private-markets target by 0.17 percentage points of plan assets (column 1). Among pensions whose written policies itemize every component target—the strict sample of column 3—the response is 0.39 percentage points. Adding the pension's net cash flow, funded ratio, and size leaves the estimates unchanged (columns 2 and 4), so the revision does not proxy for liquidity pressure, funding status, or scale. The outcome is a policy decision rather than a capital flow, and

the magnitudes are correspondingly modest. The point is directional: strategic allocations respond to relative performance in the direction the framework requires.

Panel B asks whether new commitments implement the response at the asset-class level. The unit is a pension × asset class × commitment year for four private asset classes—private equity, venture capital, private credit, and real estate—over 2009–2024. Commitments are dated by the pension's decision. The regressor is PitchBook's trailing four-quarter benchmark return for the class through the end of the prior year, a market index rather than the pension's own return. Pension × asset class fixed effects absorb each pension's standing allocation, and pension × year fixed effects absorb its total commitment activity in the year. The coefficient therefore measures whether a pension tilts a given year's commitments toward the class that has just outperformed the others. A ten-percentage-point higher lagged class return raises the probability that the pension commits to the class by 1.5 percentage points, against a base rate of 50 percent (column 1); raises the class's share of the pension's commitments by 0.9 percentage points (column 2); and raises the number of commitments by about 5 percent (column 3, Poisson).

Columns 4–6 let the slope differ by class. The response is concentrated in private credit and real estate, the private equity and venture capital slopes are indistinguishable from zero, and equality of the four slopes is rejected in every column. The class-specific private-credit slope is less stable than the pooled response, however. It depends heavily on 2009, the only year with a negative lagged private-credit return, and it does not survive replacing the return with its deviation from the four-class mean, whereas the real-estate slope does (Appendix A.3). We therefore read the private-credit evidence from the pooled response and the target revisions rather than from the class-specific slope alone. Appendix Table B.6 shows that the commitment response is stronger, on the share margin, for pensions under less net-outflow pressure and does not vary with plan size, so it is not confined to the largest plans.

Two features of this evidence matter for the mechanism. First, the signal is an asset-class benchmark rather than any manager's own record, so the response is capital moving into private credit as a class when the class's reported performance improves, over and above any reallocation

among managers. This is the margin on which preserved reported returns across many managers translate into aggregate fundraising. Second, the response operates through commitments, the closest observable counterpart to the capital that funds new lending, and it appears on the extensive, share, and count margins alike. The investor side of the loop therefore responds to reported class performance. The next subsection turns to the manager-level margin on which a fund's own track record determines its successor.

6.3 Reported Performance and Successor Fundraising

Table 7 follows private-credit funds of at least $50 million with vintages from 2000 through 2017. A successor is defined as a fund of the same strategy raised by the same manager within eight years, and the explanatory variable is the prior fund's final reported net IRR. Higher final IRR predicts whether a successor is raised: the probit coefficients are 0.010 and 0.011 without and with vintage fixed effects, respectively (columns 1 and 2). Conditional on raising a successor, each percentage point of prior-fund IRR is associated with 1.6 to 1.9 percent greater growth in fund size (columns 3 and 4). The corresponding estimate for private-equity funds is 0.9 percent per IRR point (Appendix Table B.5), about half the private-credit sensitivity.

These flow–performance estimates are predictive rather than causal. Unobserved time-varying manager quality or investment opportunities may affect both reported returns and fundraising, even though the fund-level specifications include strategy and vintage fixed effects. For the feedback mechanism, the relevant fact is that reported performance predicts both whether the manager raises again and the scale of the successor fund. A larger successor, in turn, expands the capital available for future lending.

Together, the pension-level and manager-level evidence shows why rescue has consequences beyond the seasoned fund. The manager-level conditions associated with rescue also predict improvements in that fund's reported IRR, and investors allocate more capital to managers with stronger reported returns. Rescue can therefore affect not only how existing credit problems appear in reported performance but also the resources the manager controls in the next fund generation.

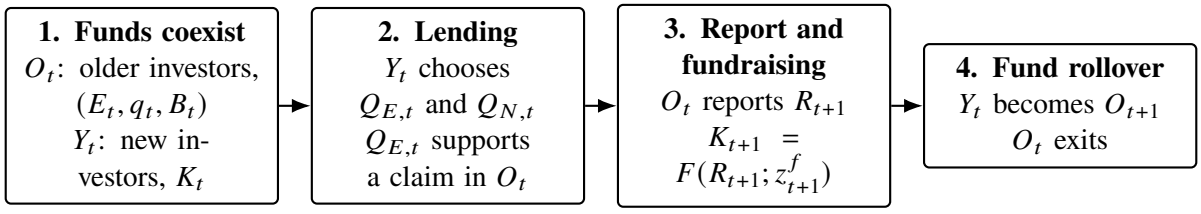
By converting a protected track record into additional lending capacity, the fundraising response completes the link from the treatment of existing credit exposures to subsequent market growth. Section 7 formalizes these links and quantifies how strongly they reinforce one another.

7 A Model of Self-Reinforcing Growth

Sections 4–6 document each link of the loop in Section 2.2: a young fund's capital finances rescue, rescue supports a seasoned fund's reported return, and that return attracts capital for the next generation. This section formalizes the loop and quantifies it, for two purposes. First, the model separates two margins that fund-level flow–performance estimates combine: capital a manager draws from rivals by outperforming them, which amplifies that manager's growth, and capital private credit attracts as an asset class, which alone expands aggregate lending. Second, it measures the impaired exposure that rescue keeps out of reported returns and carries into the next fund generation. Appendix B collects the derivations.

7.1 Fund Generations and Timing

Time is discrete, and a period is one stage of the fund cycle. Each fund lives for two periods: it raises and deploys capital while young, then reports performance and exits while seasoned. A manager operates one fund of each generation at every date, so the funds, not the manager, are the overlapping generations. At date t, the seasoned fund O_t holds a book with deployed principal B_t, blended all-in loan return q_t, and true distressed exposure E_t, the stock of impaired claims whether or not the report recognizes them. The young fund Y_t, owned by a different cohort, has committed capital K_t, from which the manager originates ordinary and rescue loans. At the end of the date, O_t reports its return and exits, and Y_t becomes the seasoned fund at $t + 1$.



The timing assigns the rescue loan to Y_t and the claim it keeps current to O_t: the manager internalizes the seasoned fund's performance and its fundraising franchise, while the younger cohort bears the new credit exposure. The manager enters date t with state $s_t = (E_t, q_t, B_t, R_t, z_t)$, where R_t is the seasoned fund's most recent investor-net report and $z_t = (z_t^f, z_t^d)$ collects funding-supply and loan-demand conditions, which follow a Markov process. The report and the funding state determine the young fund's capital:

$$K_t = F(R_t; z_t^f), \qquad F_R > 0, \qquad F_{z^f} > 0. \tag{5}$$

Investors observe the report but not true distressed exposure. F_R is the flow–performance relation of Section 6; F_{z^f} captures investor-side variation unrelated to any borrower's condition, such as the mandates of Section 5.

7.2 Lending, Distress, and Reported Performance

We suppress time subscripts and mark next-period objects with a superscript +. The young fund makes rescue loans $Q_E \in [0, E]$ and ordinary loans Q_N. A rescue loan is new credit to a borrower whose existing debt the seasoned fund holds; it neither replaces nor modifies the old claim. One dollar of rescue keeps one dollar of existing exposure current, so the seasoned fund recognizes $E - Q_E$ as impaired while Q_E remains distressed but unrecognized. Because a rescued borrower may recover, this normalization does not assume that rescue is inefficient. Ordinary loans face downward-sloping demand, $Q_N = D_0(z^d) - \eta r$, where r is the spread the manager sets; rescue borrowers pay a reservation spread $\bar{r}_E > r$, which the manager charges (Appendix B.2). With a common base charge b, the young fund's originations become next period's seasoned book, $B^+ = Q_E + Q_N$, with blended return $q^+ = b + (\bar{r}_E Q_E + r Q_N)/B^+$.

Rescue defers distress rather than removing it. A share $1 - \delta$ of rescued borrowers recovers before the next date; the rest remain distressed, with the rescue loan's return capitalized into the

unresolved claim:

$$E^+ = \underbrace{\delta(1 + b + \bar{r}_E)Q_E}_{\text{unresolved rescue exposure}} + \underbrace{(\varepsilon_0 + \varepsilon_1 Q_N)Q_N}_{\text{distress from ordinary lending}}, \tag{6}$$

where ε_0 is the distress rate on ordinary loans and $\varepsilon_1 \geq 0$ lets marginal borrower quality decline as the ordinary book expands. Equation (6) is the model's risk-accumulation margin: a fraction δ of every rescued dollar, plus accrued return, reappears in the next generation's book.

The seasoned fund reports its contractual return net of recognized losses, levered through fund debt, and net of fees. With ℓ the reduction in gross return per dollar of unrescued distress, L fund debt relative to equity, c its cost, m the management fee, h the hurdle, and ϕ the carry share,

$$\widetilde{R} = q - \ell \frac{E - Q_E}{B}, \qquad y = (1 + L)\widetilde{R} - Lc, \qquad R^+ = y - m - \phi \max\{y - m - h, 0\}. \tag{7}$$

The investor-net report R^+ enters next period's fundraising through Equation (5). Rescue reaches the manager's compensation through two channels. Let $\iota = \mathbf{1}\{y - m - h > 0\}$ indicate that the seasoned fund is above its hurdle and $\lambda_R \equiv \partial R^+ / \partial \widetilde{R} = (1 + L)(1 - \phi\iota)$ the pass-through from gross to investor-net return. One more rescue dollar removes ℓ / B of recognized loss, so

$$\text{Carry effect:} \quad \frac{\partial \,\text{Carry}}{\partial Q_E} = \iota\phi\ell, \qquad\qquad \text{Report effect:} \quad \frac{\partial R^+}{\partial Q_E} = \lambda_R \frac{\ell}{B}. \tag{8}$$

The carry effect is the immediate performance-pay incentive and operates only above the hurdle; the report effect operates on both sides of the hurdle and reaches next-generation capital through Equation (5). These are the two incentives that Section 5 separates empirically.

7.3 The Manager's Problem

The manager chooses rescue lending and the ordinary spread to maximize current compensation plus the value of future funds. Current compensation is management fees mK plus carried interest on the seasoned fund. Two costs discipline lending. Rescue carries an organizational cost $\frac{\gamma}{2K} Q_E^2$, because supporting impaired borrowers takes monitoring and concentrates the portfolio.

Deployment carries a cost $\frac{\psi}{2K}(\text{dev}^E)^2$, where dev^E is the gap between the equity supporting current lending, $(Q_E + Q_N)/(1+L)$, and the fund's lifecycle benchmark $\bar{u}K$; this is the opportunity cost of capital, whose shadow value d^A rises with total deployment. The $1/K$ normalization makes both costs homogeneous of degree one in fund scale. With β the manager's continuation factor and $V(s)$ its franchise value,

$$V(s) = \max_{Q_E \in [0,E],\, r} \Big\{ \underbrace{mK + \phi \frac{B}{1+L} \max\{y - m - h, 0\}}_{\text{fees and carried interest}} - \underbrace{\frac{\gamma}{2K} Q_E^2}_{\text{rescue cost}} - \underbrace{\frac{\psi}{2K} \big(\text{dev}^E\big)^2}_{\text{deployment cost}} + \underbrace{\beta\, \mathbb{E}_{z^+|z}\big[V(E^+, q^+, B^+, R^+, z^+)\big]}_{\text{value of future fund generations}} \Big\}. \tag{9}$$

Write V_x^+ for the derivative of next period's value with respect to state x, and $\lambda_E \equiv -V_E$ for the shadow cost of true distress. At an interior solution,

$$0 = \mathcal{R}_{Q_E} = \underbrace{\iota\phi\ell}_{\text{current carry relief}} - \underbrace{\frac{\gamma}{K} Q_E}_{\text{marginal rescue cost}} - \underbrace{d^A}_{\text{capital opportunity cost}} + \beta\, \mathbb{E}\left[\begin{array}{l} \underbrace{V_R^+ \lambda_R \frac{\ell}{B}}_{\text{report-to-fundraising value}} + \underbrace{V_B^+}_{\text{larger future book}} \\ + \underbrace{V_q^+ \frac{b + \bar{r}_E - q^+}{B^+}}_{\text{return-composition benefit}} - \underbrace{\lambda_E^+ \delta (1 + b + \bar{r}_E)}_{\text{future distressed exposure}} \end{array} \right]. \tag{10}$$

The first line is the current tradeoff: rescue protects carry only above the hurdle, and its cost rises with rescue intensity and with the value of capital in ordinary lending. The second line collects the consequences for the next generation: rescue improves the report used in fundraising, enlarges the new book, tilts it toward the higher rescue-loan return, and carries unresolved distress forward. The condition delivers the complementarity tested in Table 4: a larger K lowers the marginal rescue cost and relaxes the deployment wedge, while $\iota = 1$ switches on the carry benefit, so capacity raises rescue by more when the seasoned fund is above its hurdle. Appendix B.2 gives the boundary

conditions and the spread condition.

7.4 Competition and the Feedback Loop

Managers compete separately for borrowers and for investor capital. On the lending side, N_B managers set spreads under differentiated Bertrand competition (Appendix B.2). On the funding side, N_F managers compete for a common pool of investor capital $a(z^f)$:

$$K_i = \frac{a(z^f)}{N_F}\left[\underbrace{1 + \kappa_f R_i}_{\text{class-level flow–performance relation}} + \underbrace{\kappa_g(R_i - \bar{R}_{-i})}_{\text{relative performance}}\right]. \tag{11}$$

The class-level loading κ_f lets stronger reported performance expand the capital supplied to private credit as an asset class; the relative-performance loading κ_g reallocates a given pool toward managers that outperform their rivals. In an individual manager's first-order condition, rival reports are fixed, so one more point of reported return brings $F_R^{\text{dev}} = (a/N_F)(\kappa_f + \kappa_g)$ of capital. In a symmetric equilibrium, all reports move together, the zero-sum relative term cancels, and the same point brings only $F_R^{\text{sym}} = (a/N_F)\kappa_f$. A fund-level flow–performance estimate such as Table 7 therefore loads on $\kappa_f + \kappa_g$, but only κ_f expands aggregate capital.

At a symmetric stationary equilibrium (Appendix B.3), the strength of the feedback follows from two local responses. Let $H_K \equiv dQ_E^*/dK$ be the response of optimal rescue to committed capital and $a_R \equiv F_R^{\text{sym}}\lambda_R\ell/B$ the next-generation capital that one rescue dollar raises through the report. Their product is the loop gain, $\mathcal{G} \equiv H_K a_R$. For an exogenous shift dC^0 in per-manager capital supply with direct loading F_C,

$$\begin{aligned}
\text{Capital response:} \quad \frac{dK^*}{dC^0} &= \underbrace{F_C}_{\text{direct response}} \times \underbrace{\frac{1}{1-\mathcal{G}}}_{\text{feedback multiplier}}, \\
\text{Rescue response:} \quad \frac{dQ_E^*}{dC^0} &= \underbrace{H_K F_C}_{\text{direct capacity channel}} \times \underbrace{\frac{1}{1-\mathcal{G}}}_{\text{feedback multiplier}}.
\end{aligned} \tag{12}$$

When $0 \le \mathcal{G} < 1$, each dollar of external inflow ultimately raises committed capital by $1/(1-\mathcal{G})$ dollars, as rescue supports the report and the report attracts further capital. Replacing F_R^{sym} with F_R^{dev} gives the manager-level gain, which holds rivals fixed and is larger by the factor $(\kappa_f + \kappa_g)/\kappa_f$. Because Equation (6) carries the distress that rescue conceals into the next generation, the same loop that amplifies growth accumulates deferred impairment. The calibration asks how large the two loop gains are and how much distress the mechanism keeps out of reported returns.

7.5 Calibration

We calibrate the model to the 2011–2013 private-credit market and estimate the three parameters that govern the loop. A period is three years, and the annual discount factor is 0.95. Contract terms are the standard terms of Section 2, a 1.5 percent management fee and 20 percent carried interest above an 8 percent hurdle, and fund debt is 65 percent of investor equity at 250 basis points over the Treasury bill rate. From our data, $\delta = 0.39$ is the three-year non-accrual rate among rescued borrowers, $\varepsilon_0 = 0.076$ is the same rate among borrowers that start healthy (we set $\varepsilon_1 = 0$), and $\bar{r}_E - r$ is the 45 basis point spread premium on rescue loans. The funding loadings come from the two margins of Section 6: the response of pension commitments to the asset class's lagged return (Table 6, Panel B) gives κ_f, and the response of successor-fund size to the prior fund's IRR (Table 7) gives $\kappa_f + \kappa_g$; the former is one-fifth of the latter. The reference allocation is anchored to 2011–2013 committed capital per active manager, loan spread, and deployed book, with the base return set so that the simulated mean reported IRR equals the 10 percent 2011–2013 mean.

We estimate the rescue-cost curvature γ, the deployment wedge ψ, and the concealment loading ℓ by simulated method of moments: for each candidate we solve the manager's problem, simulate 2,500 managers, run the paper's specifications on the simulated panel, and match five moments. The capacity $\times$ above-hurdle coefficients in the rescue regression (Table 4, with capital in log dollars) and in the three-year IRR change (Table 5, column 3) identify γ and ℓ: how strongly rescue responds to capital when carry is in the money, and how much reported return a rescue dollar protects. Recognized distress, $(E - Q_E)/B$, is matched to the 2011–2013 non-accrual share of

BDC book cost, and the mean and cross-fund dispersion of the deployment ratio, capital called by fund age three, discipline ψ. The ordinary spread is held at its 2011–2013 value, and simulated rescue dollars map to observed rescue loans through an event rate normalized to the 6.8 percent mean rescue rate.

7.6 Results

Table 8 reports the estimates and the fit: $\gamma = 0.70$, $\psi = 0.87$, and $\ell = 0.75$. The model reproduces the two interactions that constitute the mechanism with the right sign and order of magnitude and overshoots both: the capacity × above-hurdle effect on rescue is 0.016 against 0.010 in the data, and the effect on the seasoned fund's three-year IRR change is 2.3 percentage points against 1.4, each about two standard errors above the data. The two overshoots have a common source. The model's rescue response to capacity above the hurdle is somewhat too strong, and because rescue protects the report, the excess carries into the IRR moment. The model misses on two dimensions. The deployment ratio is more dispersed than in the data, 0.38 against 0.22, because the coarse two-point loan-demand shock makes the ordinary book alternate between two states; this miss is identical across candidates and uninformative about the parameters. Recognized distress is 5.9 percent of the book per period against 2.7 percent. This miss reflects a tension between measurements: δ and ε_0 are flows into non-accrual over three years, whereas the target is a point-in-time stock converted to a flow by assuming a one-year stay on the book. With these flows, total distress is 13.9 percent of the book, and the model would have to rescue four-fifths of it, rather than 57 percent, to recognize only 2.7 percent—more rescue than the interaction moments allow.

The estimates describe an economy in which rescue is cheap and effective. A rescue-cost curvature of 0.70 means that the marginal cost of rescue rises slowly with intensity, so managers keep 57 percent of newly distressed exposure current each period. A concealment loading of 0.75 means that a dollar of unrescued distress lowers the gross return by three-quarters of its share of the book, so a rescue dollar protects the report by somewhat less than one-for-one. Concealed distress, the exposure kept current by rescue loans and absent from reported returns, is 8.1 percent of the

seasoned book, against 5.9 percent recognized: reported returns understate distress by more than half. Total distress of 13.9 percent is of the same order as the 9 percent of borrower-years marked below 90 cents in Section 3, which counts borrowers rather than dollars.

Figure 6 shows the policy behind these numbers, the model's counterpart of Table 4. Rescue intensity rises with the young fund's capital at every level of the seasoned fund's IRR, because a larger fund lowers the marginal cost of rescue and relaxes the deployment wedge. Over the solved range, a 60 percent increase in capital raises rescue intensity from 51 to 64 percent when the seasoned fund reports 5 or 8 percent, and from 51 to 67 and 70 percent when it reports 12 and 15 percent. Below the hurdle the lines coincide, because only the fundraising channel operates; above it, the carry channel adds to the return on rescue. Capacity therefore raises rescue everywhere, but by one-quarter to one-half more when carry is in the money.

Figure 7 quantifies manager-level against industry-level amplification by following committed capital after a one-period improvement in reported return, in percent of the direct response. When one manager's report improves and rivals' are fixed (Panel A), the direct response is 4.9 times that of the same improvement across all managers (Panel B), because the fund-level relation loads on $\kappa_f + \kappa_g$ while only κ_f expands the asset class. The feedback then compounds the two differently. The manager-level loop gain is 0.22, so one manager's capital ultimately rises by 1.29 times its direct response, most of it within two periods; the symmetric gain is 0.045, so capital per manager rises by 1.05 times. The loop is thus an economically meaningful amplifier of an individual manager's growth and a modest amplifier of aggregate growth: a manager that preserves its track record gains capital largely at rivals' expense, and the industry-wide loop adds 5 percent to an inflow, not a multiple of it. The aggregate importance of the mechanism lies instead in what the capital sustains. Rescue keeps 8.1 percent of the seasoned book out of reported returns, and Equation (6) carries about half of that exposure into the next generation's book, where keeping it current requires the next young fund's capital. Reported returns in the calibrated economy are therefore a lagging measure of credit quality whose accuracy depends on continued inflows.

8 Conclusion and Policy Implications

Private credit's organization into successive funds makes portfolio management and fundraising interdependent. A manager can finance a troubled company already held by an older fund through a recently raised fund. Such follow-on financing can delay deterioration in the older fund's reported return, strengthen the record presented to prospective investors, and expand the resources under management. A decision intended to manage legacy exposure can thus alter both the volume and destination of later credit.

Our results support this interpretation at the borrower, manager, and investor levels. After accounting for a borrower's financing demand and each lender's credit supply, impaired borrowers are particularly likely to obtain their next loan from an incumbent. Subsequent valuations and non-accrual status provide no evidence that incumbents select troubled borrowers with better recovery prospects. At the manager level, rescue activity rises when deployable capital coexists with a mature fund above its incentive threshold, rather than with greater origination capacity alone. Those same circumstances predict improvements in the mature fund's later IRR, and the performance of an earlier fund predicts both its successor's formation and scale and whether prior investors commit again. These findings connect relationship-specific lending to compensation and capital raising across funds. The analysis concerns the incentive behind the pattern, not the efficiency of every transaction; additional financing can sometimes resolve temporary distress or preserve going-concern value.

The calibrated model organizes these findings. It separates two margins that the fund-level flow–performance response combines: capital that private credit attracts as an asset class when reported returns improve, and capital that a manager draws away from rivals by outperforming them. Only the first expands aggregate lending. In the calibration, four-fifths of the capital a manager attracts by outperforming comes from rivals, so the feedback amplifies an individual manager's capital response by 29 percent but the industry's by 5 percent. The model also makes explicit that rescue defers distress rather than removing it: at the estimates, rescue keeps more distress out of reported returns than those returns recognize, and about half of that exposure passes to the next

fund generation, where keeping it current depends on the arrival of new capital.

This mechanism shifts where fragility resides. Fund lockups limit withdrawals but do not prevent delayed recognition of credit deterioration. During periods of abundant inflows, continued financing can sustain weak borrowers and transfer additional exposure from earlier to later investor cohorts. If commitments retrench, managers may lose the means to continue that support just as weaker valuations make new capital more difficult to secure. Private credit can therefore be insulated from conventional run risk yet remain vulnerable to a fundraising-dependent reversal in loss recognition and credit supply.

The debate over private assets in 401(k) plans makes the incidence of these risks especially consequential. In March 2026, the U.S. Department of Labor proposed a process-based safe harbor for fiduciaries selecting diversified investment options with alternative-asset exposure (U.S. Department of Labor, 2026). Our findings do not determine whether private-credit products should be offered in defined-contribution plans. They identify information needed to evaluate them. Due diligence should supplement headline IRR with evidence on loan valuations, non-accrual, repeat financing of troubled companies, and the allocation of exposures across funds managed by the same organization. Fiduciaries and plan participants should be able to determine whether a later fund assumes risks that sustain an earlier fund's returns and the terms on which it does so.

The same principle applies to private-credit oversight more broadly. Fund-by-fund reporting can obscure transfers among investors in different vehicles. Manager-wide disclosure linking borrower exposures, new financing, valuations, and non-accrual status across funds would make such support visible. Stress analysis should also ask how asset values and credit supply would respond if fundraising slowed and refinancing became less available. These measures address the feedback documented in this paper while preserving the potential benefits of long-horizon, illiquid lending.

REFERENCES

Berk, Jonathan B. and Richard C. Green, "Mutual Fund Flows and Performance in Rational Markets," *Journal of Political Economy*, 2004, *112* (6), 1269–1295.

Block, Joern, Young Soo Jang, Steven N. Kaplan, and Anna Schulze, "A Survey of Private Debt Funds," *Review of Corporate Finance Studies*, 2024, *13* (2), 335–383.

Board of Governors of the Federal Reserve System, "Financial Stability Report," Technical Report, Board of Governors of the Federal Reserve System May 2026.

Caballero, Ricardo J., Takeo Hoshi, and Anil K. Kashyap, "Zombie Lending and Depressed Restructuring in Japan," *American Economic Review*, 2008, *98* (5), 1943–1977.

Cai, Fang and Sharjil Haque, "Private Credit: Characteristics and Risks," FEDS Notes, Board of Governors of the Federal Reserve System February 2024.

Chernenko, Sergey and David Scharfstein, "Private Credit and Financial Stability," 2025. Working paper, SSRN 5930356.

_ , Isil Erel, and Robert Prilmeier, "Why Do Firms Borrow Directly from Nonbanks?," *Review of Financial Studies*, 2022, *35* (11), 4902–4947.

_ , Robert Ialenti, and David Scharfstein, "Bank Capital and the Growth of Private Credit," 2026. Working paper, SSRN 5097437.

Chung, Ji-Woong, Berk A. Sensoy, Léa H. Stern, and Michael S. Weisbach, "Pay for Performance from Future Fund Flows: The Case of Private Equity," *Review of Financial Studies*, 2012, *25* (11), 3259–3304.

Davydiuk, Tetiana, Isil Erel, Wei Jiang, and Tatyana Marchuk, "Common Investors Across the Capital Structure: Private Debt Funds as Dual Holders," 2025. Working paper, SSRN 4992219.

_ , Tatyana Marchuk, and Samuel Rosen, "Direct Lenders in the U.S. Middle Market," *Journal of Financial Economics*, 2024, *162*, 103946.

_ , _ , and _ , "Market Discipline in the Direct Lending Space," *Review of Financial Studies*, 2024, *37* (4), 1190–1264.

Erel, Isil, Thomas Flanagan, and Michael S. Weisbach, "Risk-Adjusting the Returns to Private Debt Funds," NBER Working Paper 32278, National Bureau of Economic Research 2024.

Gompers, Paul and Josh Lerner, "Money Chasing Deals? The Impact of Fund Inflows on Private Equity Valuations," *Journal of Financial Economics*, 2000, *55* (2), 281–325.

Hochberg, Yael V., Alexander Ljungqvist, and Annette Vissing-Jørgensen, "Informational Holdup and Performance Persistence in Venture Capital," *Review of Financial Studies*, 2014, *27* (1), 102–152.

Institutional Limited Partners Association, "ILPA Principles 3.0: Fostering Transparency, Governance and Alignment of Interests for General and Limited Partners," Technical Report, Institutional Limited Partners Association 2019.

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

Ivashina, Victoria, "Private Credit: What Do We Know?," *Annual Review of Financial Economics*, 2026, *18*. Review in Advance.

Jang, Young Soo, Dasol Kim, Amir Sufi, and Xiangyu Chen, "The Lending Technology of Direct Lenders in Private Credit," NBER Working Paper 34500, National Bureau of Economic Research 2025. Revised May 2026.

_ , _ , and _ , "Private Credit at Scale: Evidence from U.S. Regulatory Data," 2025. Working paper.

Kaplan, Steven N. and Antoinette Schoar, "Private Equity Performance: Returns, Persistence, and Capital Flows," *Journal of Finance*, 2005, *60* (4), 1791–1823.

Matvos, Gregor, Tomasz Piskorski, and Amit Seru, "Private Credit, Balance Sheets and Financial Stability," NBER Working Paper 34991, National Bureau of Economic Research 2026.

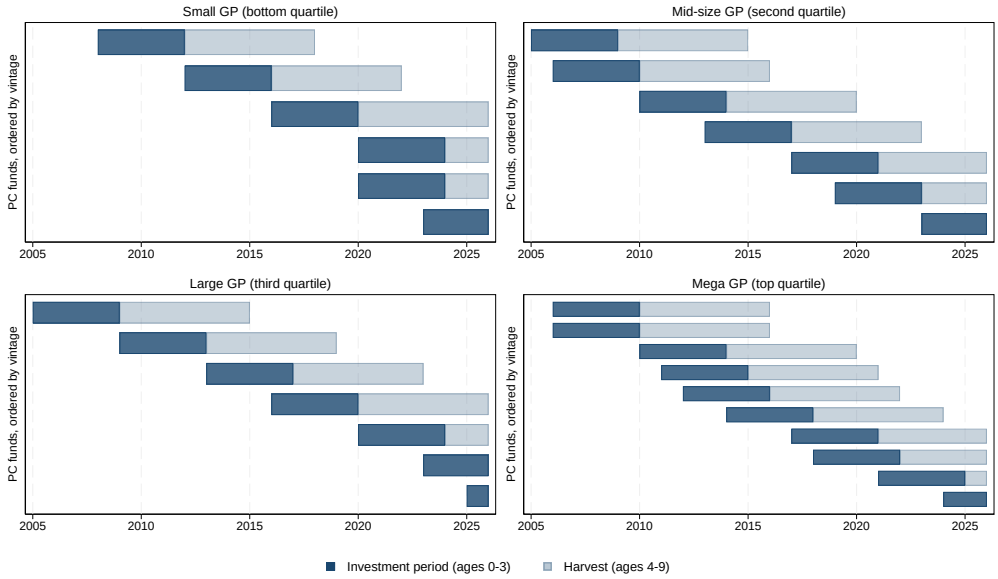
Pástor, Ľuboš and Robert F. Stambaugh, "On the Size of the Active Management Industry," *Journal of Political Economy*, 2012, *120* (4), 740–781.

Peek, Joe and Eric S. Rosengren, "Unnatural Selection: Perverse Incentives and the Misallocation of Credit in Japan," *American Economic Review*, 2005, *95* (4), 1144–1166.

Tantri, Prasanna, "Identifying Ever-greening: Evidence Using Loan-Level Data," *Journal of Banking and Finance*, 2021, *122*, 105997.

U.S. Department of Labor, "Fiduciary Duties in Selecting Designated Investment Alternatives," Proposed Rule, Employee Benefits Security Administration March 2026.

Panel A. Fund Lifecycles of Representative Managers



Panel B. Concurrently Active Funds, Share of Private Credit Capital

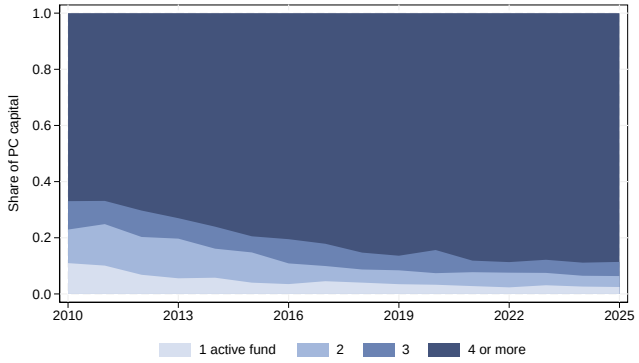


Figure 1. Overlapping Fund Generations. This figure describes how private credit managers operate several fund generations at once. The sample is PitchBook debt funds (mezzanine, direct lending, distressed debt, general debt, credit special situations, venture debt, bridge financing) with a reported vintage. Panel A plots the fund lifecycles of four representative managers, one from each quartile of managers ranked by total private credit capital raised across 2005–2025 vintages (Small: below $0.5 billion; Mid-size: $0.5–1.5 billion; Large: $1.5–4.9 billion; Mega: above $4.9 billion). Each bar is one fund; the dark segment is the investment period (fund ages 0–3) and the light segment the harvest period (ages 4–9), truncated at 2026. Within each quartile, the representative manager is the one whose fund count is closest to the median among managers with at least three funds, a first fund raised in 2005–2010, no gap above four years between consecutive vintages, and a fund raised in 2021 or later. Panel B plots, for each year 2010–2025, the share of private credit capital held by managers with one, two, three, or four or more active funds (aged 0–9). The sample is manager-years in which the manager has at least one active fund and lends in the deal data at least once in 2010–2025; each manager-year is weighted by the manager's cumulative private credit capital raised through the prior year, so manager-years with no recorded fund size are excluded.

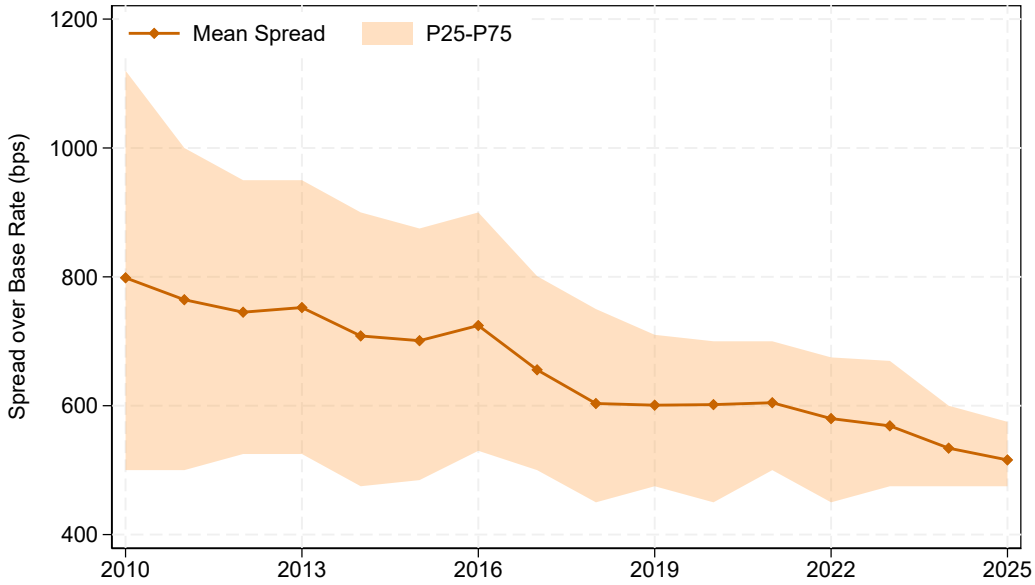


Figure 2. Private Credit Deal Spreads. This figure plots the pricing of private credit loans over time. The sample is all loan tranches in the deal universe with issue dates 2010–2025, from PitchBook. The line shows the mean spread over the base rate in basis points, where a floating-rate tranche contributes its stated margin and a fixed-rate tranche its coupon minus the 3-month rate in the issuance year; the shaded band is the interquartile range, and years with fewer than 50 priced tranches are masked.

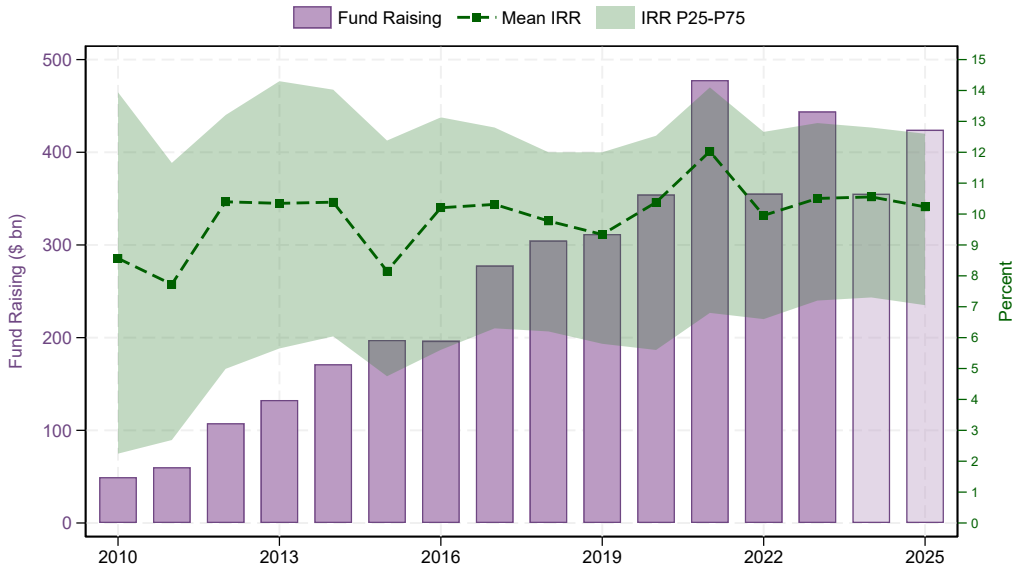
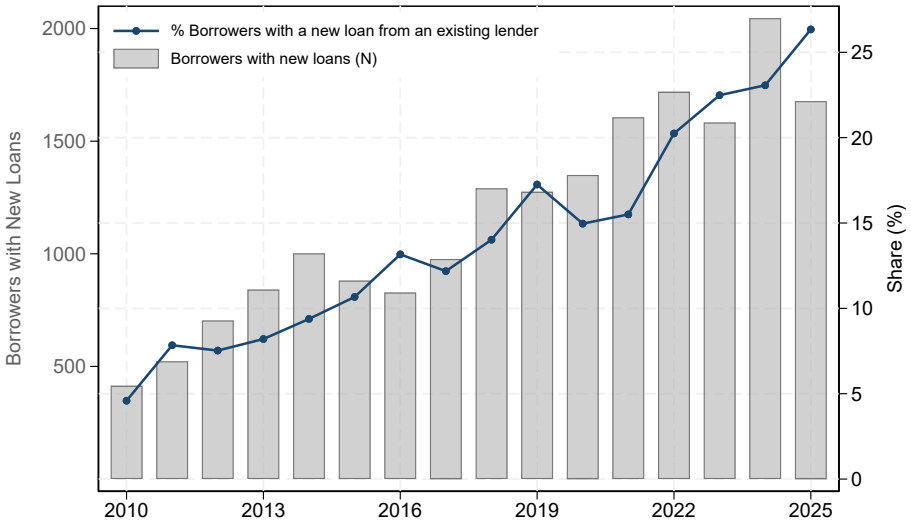


Figure 3. Private Credit Fundraising and Fund IRR. This figure plots private credit fundraising against the returns funds report. The sample is PitchBook debt funds (mezzanine, direct lending, distressed debt, general debt, credit special situations, venture debt, bridge financing). Bars show capital raised by vintage year in $ billions (left axis); the 2024–25 bars are spliced to PitchBook's published industry totals and shaded as provisional. The line shows the mean net IRR across funds reporting in each year, with the interquartile band, in percent (right axis).

Panel A. Over Time



Panel B. By Borrower Credit Quality

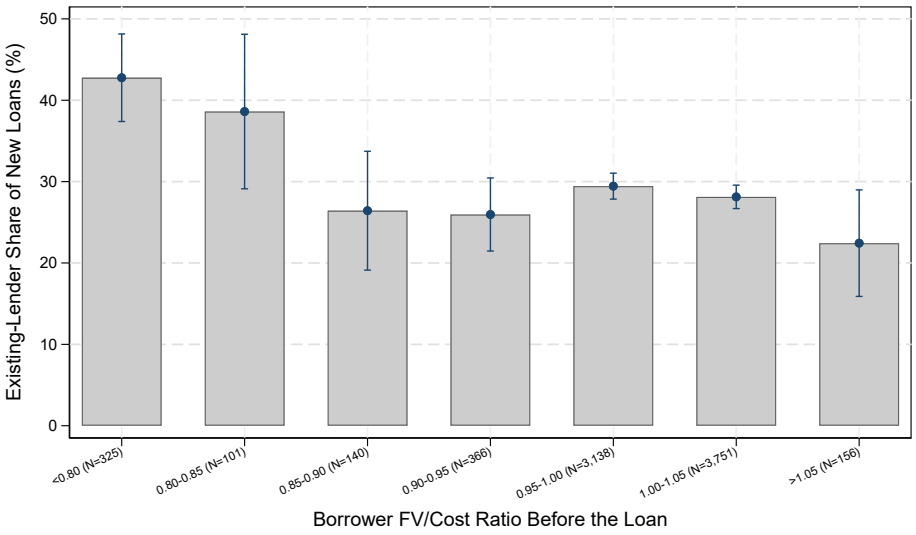


Figure 4. Existing-Lender Share of New Lending. This figure plots the share of borrowers with new loans that borrow from an existing lender — the borrower's most recent prior lender. Panel A plots the share over time (line, left axis), alongside the number of borrowers taking new loans (bars, right axis); the sample is borrower-years with at least one new loan, 2010–2025. Panel B plots the share by the borrower's pre-deal FV/Cost ratio — the latest BDC mark preceding the year's first loan, within the prior year — for the subsample with such a mark. Error bars are 95% binomial confidence intervals; bin sizes appear in the axis labels.

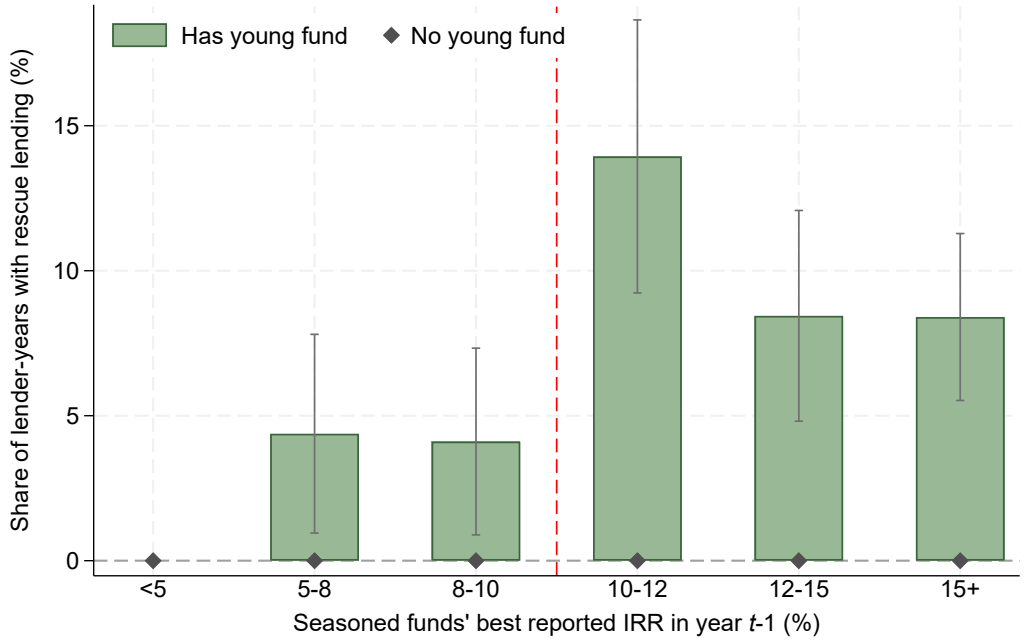


Figure 5. Rescue Lending and the Seasoned Funds' Reported IRR. This figure plots the share of lender-years with rescue lending — new money to an impaired existing borrower — by the best net IRR reported at $t-1$ across the lender's seasoned funds, splitting lender-years by whether the lender has a young fund (aged one to four). The sample is an annual panel of PitchBook private-debt lenders that lend in the deal data, 2010–2025, restricted to lender-years with a seasoned reporting fund; the same panel underlies Table 4. The dashed line marks the ten-percent hurdle. Diamonds mark lender-years without a young fund, which record no rescue lending in any bin. Error bars are 95% binomial confidence intervals for the bin shares.

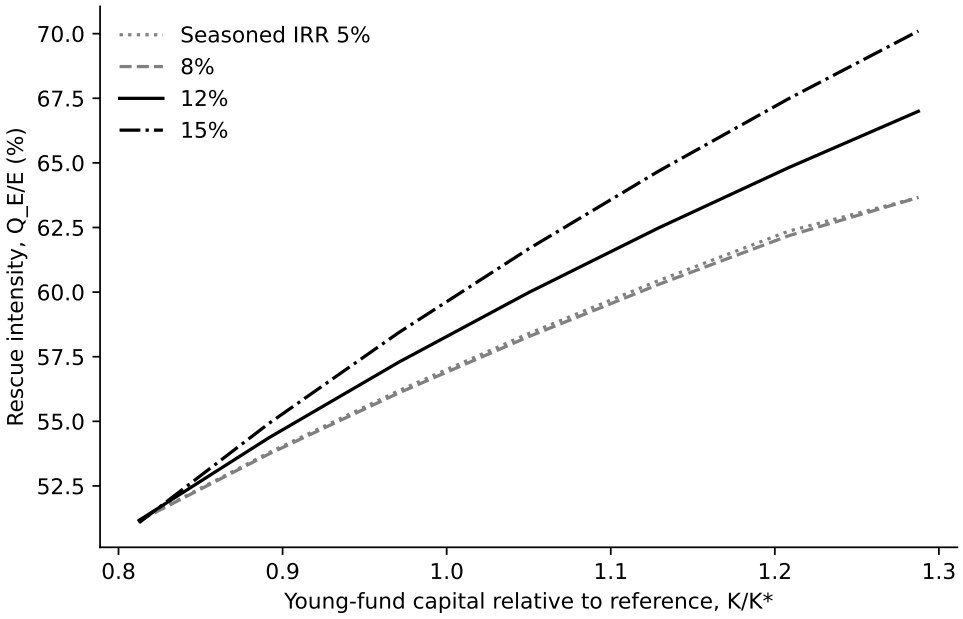


Figure 6. Optimal Rescue in the Calibrated Model. This figure plots the manager's optimal rescue intensity, the share of the seasoned fund's distressed exposure kept current by rescue loans, against the young fund's committed capital relative to its reference value, at the calibrated parameters of Table 8. Each line holds the seasoned fund's reported annual IRR at the indicated level; the 10 percent hurdle separates the two gray lines (below hurdle) from the two black lines (above). All other states are held at the reference allocation. The horizontal range is the range of capital covered by the solved state grid.

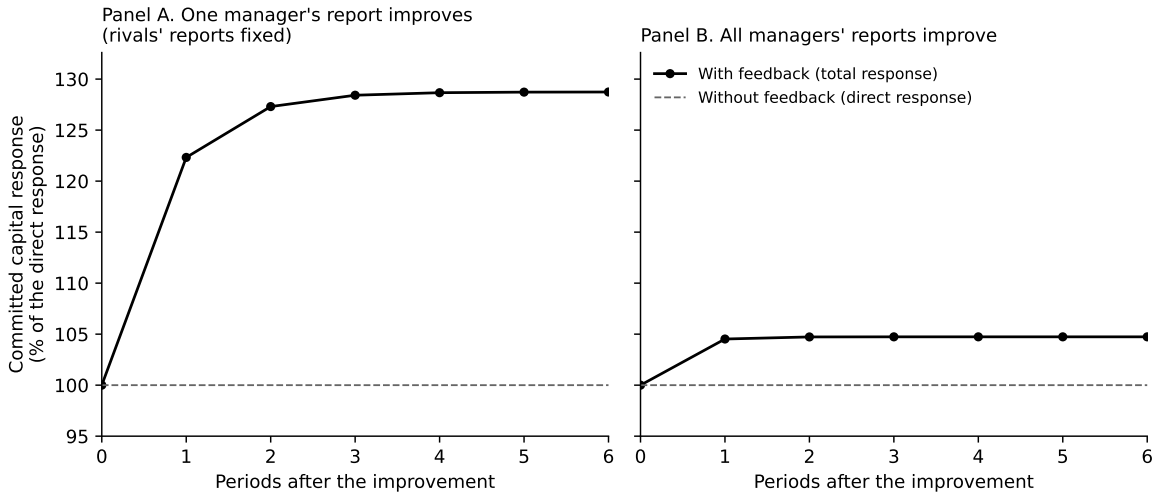


Figure 7. Manager-Level and Industry-Level Amplification. This figure traces the response of committed capital to a one-period improvement in reported return in the calibrated model, in percent of the direct response through the funding rule, Equation (11). Panel A is the response of one manager's capital when only its own report improves and rivals' reports are held fixed; Panel B is the response of capital per manager when every manager's report improves by the same amount. The dashed line is the response without feedback, the direct response through the funding rule, normalized to 100; the solid line is the total response, which adds the feedback through rescue, reported return, and fundraising period by period. Each period adds one round of the loop at the loop gain $\mathcal{G} = H_K a_R$ evaluated at the calibrated parameters of Table 8 with the unilateral funding response in Panel A and the symmetric response in Panel B. The cumulative response converges to $1/(1-\mathcal{G})$ times the direct response, Equation (12). The direct response itself is $(\kappa_f + \kappa_g)/\kappa_f$ times larger in Panel A than in Panel B, because a single manager's report also draws capital from rivals.

Table 1. Summary Statistics

Notes: This table reports summary statistics for the estimation samples. Panel A is the borrower-year panel behind Table 3: private credit borrowers whose existing loans — those on BDC books before year t — are observed in BDC holdings with a year-start mark, 2010–2024. FV/Cost is the cost-weighted fair value of those loans at their last report in year $t-1$; Impaired indicates FV/Cost below 0.9; Borrow from Existing and Borrow from New indicate a new loan in year t from the most recent prior lender or from any other lender; Maturity is the cost-weighted mean remaining contractual maturity. Panel B is the lender-year panel behind Table 4: PitchBook private-debt lenders with an active fund that lend in the deal data, 2010–2025. Rescue Lending indicates new money to an impaired existing borrower; New-Borrower Lending indicates a new loan to a borrower the lender has not lent to before; Above-Hurdle Seasoned Fund is defined for the lender-years with a seasoned reporting fund. Panel C is the fund-year panel behind Table 5: seasoned debt funds (aged five or more) with a reported net IRR. Estimation samples in the corresponding tables are slightly smaller where singleton fixed-effect groups drop.

Panel A. Borrower-Year Sample

	N	Mean	P10	P25	P50	P75	P90
FV/Cost	27,161	0.97	0.91	0.98	1.00	1.01	1.02
Impaired	27,161	0.09	0.00	0.00	0.00	0.00	0.00
Borrow from Existing	27,161	0.06	0.00	0.00	0.00	0.00	0.00
Borrow from New	27,161	0.15	0.00	0.00	0.00	0.00	1.00
Maturity	27,161	4.23	1.71	2.92	4.20	5.46	6.59

Panel B. Lender-Year Sample

	N	Mean	P10	P25	P50	P75	P90
Rescue Lending	4,534	0.03	0.00	0.00	0.00	0.00	0.00
New-Borrower Lending	4,534	0.43	0.00	0.00	0.00	1.00	1.00
Young Fund	4,534	0.76	0.00	1.00	1.00	1.00	1.00
Above-Hurdle Seasoned Fund	1,621	0.55	0.00	0.00	1.00	1.00	1.00
Impaired Borrowers (count)	4,534	0.62	0.00	0.00	0.00	0.00	1.00
Total Borrowers (count)	4,534	16.11	1.00	1.00	3.00	9.00	43.00
Lender Size ($bn raised)	4,534	8.76	0.05	0.27	1.07	5.19	20.43

Panel C. Seasoned-Fund IRR Sample

	N	Mean	P10	P25	P50	P75	P90
IRR (%)	6,397	9.65	1.90	6.10	8.87	12.20	17.30
Lender Has Young Fund	6,397	0.76	0.00	1.00	1.00	1.00	1.00
Lender Raises New Fund	6,397	0.47	0.00	0.00	0.00	1.00	1.00
Above-Hurdle	6,397	0.40	0.00	0.00	0.00	1.00	1.00
Fund Size ($bn)	6,378	1.11	0.12	0.26	0.58	1.30	2.50
Fund Age (years)	6,397	10.12	5.00	6.00	9.00	13.00	17.00

Table 2. Borrower–Lender Matching

Notes: This table reports regression estimates of lender–borrower matching in private credit. The sample is the Cartesian product of active borrowers and active lenders in each year over 2004–2025 (deals with at most five lenders), restricted to borrower-years with a pre-deal FV/Cost ratio — the borrower's latest mark preceding its first loan of year t. The dependent variable in all columns is an indicator for whether borrower i and lender j have a deal in year t. *Existing Lender* equals one if lender j was the most recent lender to borrower i prior to year t. *Impaired* is an indicator equal to one if the borrower's pre-deal FV/Cost ratio is below 0.9. *Impairment* is the continuous impairment depth, defined as 1−FV/Cost and equal to zero at par, so that larger values indicate deeper markdowns. Columns 1 and 3 include borrower, lender, and year fixed effects separately; columns 2 and 4 include borrower-year and lender-year fixed effects. Standard errors are two-way clustered at the borrower and lender levels, and *, **, *** denote 10%, 5%, and 1% statistical significance.

Dependent Variable: Borrower–Lender Match

	(1)	(2)	(3)	(4)
Existing Lender	0.271***	0.269***	0.271***	0.269***
	(0.024)	(0.024)	(0.024)	(0.024)
Existing Lender × Impaired	0.086***	0.088***		
	(0.026)	(0.025)		
Existing Lender × Impairment			0.269***	0.273***
			(0.075)	(0.075)
Borrower FEs	Yes	No	Yes	No
Lender FEs	Yes	No	Yes	No
Year FEs	Yes	No	Yes	No
Borrower-Year FEs	No	Yes	No	Yes
Lender-Year FEs	No	Yes	No	Yes
N	3,397,403	3,397,403	3,397,403	3,397,403
Adj. R^2	0.076	0.093	0.076	0.093

Table 3. Rescued Impaired Borrowers Do Not Outperform

Notes: This table reports regression estimates of the outcomes of borrowers' existing loans on lender type and borrower credit quality. The sample is a borrower-year panel of private credit borrowers whose existing loans — those on BDC books before year t — are observed in BDC holdings, 2010–2024. *Impaired* is an indicator equal to one if the cost-weighted FV/Cost of the borrower's existing loans, at their last report in year $t-1$, is below 0.9. *Borrow from Existing* equals one if the borrower receives a new loan in year t from its existing lender as defined in the matching table; *Borrow from New* equals one if from any other lender. The dependent variables track the same existing loans into years $t+1$, $t+2$, and $t+3$: in columns 1–3, the change in their FV/Cost from $t-1$; in columns 4–6, an indicator for the loans being on non-accrual at the horizon. *Maturity*, the cost-weighted mean remaining contractual maturity in years of these loans at the end of $t-1$, controls for scheduled roll-off. All specifications include industry-year fixed effects. Standard errors are clustered at the borrower level, and *, **, *** denote 10%, 5%, and 1% statistical significance.

	Δ FV/Cost			Non-Accrual		
	(1) $t+1$	(2) $t+2$	(3) $t+3$	(4) $t+1$	(5) $t+2$	(6) $t+3$
Impaired	0.063***	0.076***	0.086***	0.313***	0.265***	0.252***
	(0.009)	(0.013)	(0.021)	(0.018)	(0.024)	(0.035)
Impaired × Borrow from Existing	-0.023	-0.050	-0.133*	0.170***	0.132	0.036
	(0.036)	(0.054)	(0.078)	(0.060)	(0.091)	(0.136)
Impaired × Borrow from New	-0.012	0.012	0.080	-0.088**	-0.066	0.037
	(0.025)	(0.038)	(0.056)	(0.042)	(0.058)	(0.088)
Borrow from Existing	0.003	0.003	0.007	0.004	-0.001	-0.001
	(0.004)	(0.006)	(0.009)	(0.007)	(0.012)	(0.018)
Borrow from New	0.008***	0.005	-0.002	-0.008*	-0.008	-0.014
	(0.002)	(0.004)	(0.006)	(0.004)	(0.007)	(0.011)
Maturity	0.004***	0.006***	0.007***	-0.015***	-0.017***	-0.019***
	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.003)
Industry-Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	13,839	7,970	4,254	13,839	7,970	4,254
Adj. R^2	0.052	0.070	0.052	0.154	0.108	0.085

Table 4. Lender Incentives and Rescue Lending

Notes: This table relates lender-level rescue lending — new money to an impaired existing borrower — to the lender's fund-generation state. The sample is an annual panel of PitchBook private-debt lenders that lend in the deal data and have an IRR-reporting fund, 2010–2025. *Young Fund* indicates a fund aged one to four; *Above-Hurdle (Below-Hurdle) Seasoned Fund*, a seasoned fund reporting a net IRR above (below) ten percent. Panel A reports rescue rates by cell. In Panel B, *Lender Capacity* is the treatment named in the column header: *Has Young Fund* in columns 1–2, and *LP Mandates* — private-credit mandate announcements over the window shown, by LPs first committed to the lender at least five years earlier, standardized — in columns 3–4. Column 1 has no controls or fixed effects; columns 2–4 add lender and year fixed effects and controls for the lender's impaired and total incumbent borrowers and cumulative capital raised. Standard errors are clustered by lender, and *, **, *** denote 10%, 5%, and 1% statistical significance.

Panel A. Frequency of Rescue Lending by Lender's Fund Condition

	No Young Fund		Has Young Fund	
	Rescue %	N	Rescue %	N
Above-Hurdle Seasoned Fund	0.0	101	9.9	790
Below-Hurdle Seasoned Fund	0.0	115	4.8	477
No Seasoned Fund	–	0	5.8	138

Panel B. Regressions, Dependent Variable: Rescue Lending

Lender Capacity from:	Has Young Fund		LP Mandates	
Mandate window:			$[t-1]$	$[t-3, t-1]$
	(1)	(2)	(3)	(4)
Lender Capacity × Above-Hurdle Seasoned Fund	0.051**	0.026**	0.026***	0.025***
	(0.020)	(0.012)	(0.007)	(0.009)
Lender Capacity	0.048***	-0.008	-0.030***	-0.029**
	(0.014)	(0.014)	(0.008)	(0.011)
Above-Hurdle Seasoned Fund	-0.010	-0.089***	-0.072***	-0.073***
	(0.027)	(0.025)	(0.022)	(0.022)
Below-Hurdle Seasoned Fund	-0.010	-0.081***	-0.081***	-0.081***
	(0.027)	(0.022)	(0.021)	(0.021)
Controls	No	Yes	Yes	Yes
Lender FEs	No	Yes	Yes	Yes
Year FEs	No	Yes	Yes	Yes
Mean dep. var.	0.067	0.068	0.068	0.068
N	1,621	1,584	1,584	1,584
Lenders	226	189	189	189
R-squared	0.019	0.506	0.510	0.509

Table 5. Change in Seasoned Fund IRR

Notes: This table asks whether the lender's fundraising state is followed by better reported returns of its above-hurdle seasoned funds. The sample is a fund-year panel of seasoned debt funds — aged five or more — with a reported net IRR. The dependent variable is the change in the fund's reported IRR from year t to the horizon year given in the column header ($t+1$, $t+2$, or $t+3$). *Lender Capacity* is the treatment named in the top header: *Has Young Fund* — a fund aged one to four, as in Table 4 — in columns 1–3, and *LP Mandates* over the prior year — as defined in Table 4 — in columns 4–6. *Above-Hurdle* indicates the fund's own reported net IRR is at least ten percent in year t, the region where carry is plausibly in the money; true hurdles are unobserved. The fund's IRR level is controlled by decile fixed effects, so the fundraising state cannot be credited for the level itself; *Fund Size* (log) and *Fund Age* are controls. All columns compare funds within lender and year. Standard errors are clustered by lender, and *, **, *** denote 10%, 5%, and 1% statistical significance.

Dependent Variable: Change in Seasoned Fund IRR

Lender Capacity from:	Has Young Fund			LP Mandates		
Mandate window:					$[t-1]$	
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
Lender Capacity × Above-Hurdle	0.887***	1.140***	1.363***	0.172**	0.251**	0.339**
	(0.298)	(0.366)	(0.433)	(0.077)	(0.127)	(0.148)
Lender Capacity	-0.777***	-0.590	-0.442	-0.042	-0.090	-0.069
	(0.260)	(0.381)	(0.560)	(0.066)	(0.085)	(0.108)
Above-Hurdle	-0.405	-1.424**	-1.561***	0.266	-0.617	-0.653
	(0.480)	(0.623)	(0.597)	(0.378)	(0.497)	(0.514)
Fund Size	-0.002	0.014	0.065	-0.004	0.011	0.061
	(0.058)	(0.099)	(0.123)	(0.059)	(0.099)	(0.121)
Fund Age	0.042***	0.070***	0.088***	0.041***	0.067***	0.081***
	(0.011)	(0.018)	(0.023)	(0.011)	(0.018)	(0.023)
IRR-Decile FEs	Yes	Yes	Yes	Yes	Yes	Yes
Lender FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	5,008	4,168	3,454	5,008	4,168	3,454
Lenders	260	226	196	260	226	196
R-squared	0.114	0.186	0.178	0.111	0.184	0.176

Table 6. Asset-Class Performance and LP Allocations

Notes: This table relates the private-market allocations of limited partners — U.S. public pensions from the Public Plans Database, matched to their PitchBook commitment histories — to lagged asset-class performance. In Panel A, the dependent variable is the change in the pension's private-markets policy target weight (the PE, private credit, and real estate component targets, t to $t + 1$), regressed on the lagged private-minus-public composite return over the window ending at the pension's fiscal year end. Columns 2 and 4 add net cash flow, the funded ratio, and log plan assets; none enters significantly. The baseline sample fills a missing component target with zero only when the class's actual weight is a confirmed zero; the strict sample requires every component target reported at both endpoints. Standard errors are clustered by pension and return window in columns 1–2 and by pension in columns 3–4. In Panel B, 161 pensions are crossed with four private asset classes over 16 commitment years, 2009–2024; the dependent variables are an indicator for any commitment to the class in the year (columns 1 and 4), the class's share of the pension's commitment count (columns 2 and 5), and the commitment count, estimated by PPML (columns 3 and 6). The regressor is the class's trailing four-quarter benchmark return through December 31 of $t - 1$, a market index rather than the pension's own return. Columns 4–6 interact the return with class indicators in a single regression per column, so each coefficient is that class's own slope; the joint slope tests are reported beneath each of these columns. Standard errors are two-way clustered by pension and commitment year. *, **, *** denote 10%, 5%, and 1% statistical significance.

Panel A. Private-Markets Target Revisions, Pension–Year Level

	Baseline sample		Strict sample	
	(1)	(2)	(3)	(4)
Lagged private − public return	0.0165***	0.0173***	0.0388***	0.0407***
	(0.0054)	(0.0061)	(0.0140)	(0.0145)
Controls	N	Y	N	Y
Pension FEs	Y	Y	Y	Y
N	1,759	1,746	303	302
R-squared	0.069	0.070	0.129	0.136

Panel B. New Commitments, Pension–Asset Class–Year Level

	Pooled			By asset class		
	Any	Share	Count	Any	Share	Count
	(1)	(2)	(3)	(4)	(5)	(6)
Lagged asset-class return	0.1467**	0.0926**	0.5213*			
	(0.0520)	(0.0396)	(0.2873)			
Return × Private credit				0.3038**	0.1668*	0.8279***
				(0.1412)	(0.0797)	(0.2815)
Return × Private equity				-0.0083	-0.1539	0.1554
				(0.0825)	(0.0953)	(0.3857)
Return × Venture capital				-0.0062	0.0214	0.4101
				(0.0689)	(0.0361)	(0.3923)
Return × Real estate				0.3609***	0.2469**	0.9068**
				(0.0981)	(0.0956)	(0.4398)
Pension × Asset Class FEs	Y	Y	Y	Y	Y	Y
Pension × Year FEs	Y	Y	Y	Y	Y	Y
N	10,008	7,656	7,200	10,008	7,656	7,200
p: all four slopes equal				0.001	0.022	0.032
p: private credit = avg. of others				0.096	0.056	0.149

Table 7. Private Credit Fund Performance and Follow-On Fundraising

Notes: This table relates fundraising to the reported returns of the lender's funds. The sample is private credit funds of at least $50 million with vintages 2000–2017, following the specification of Chung et al. (2012); a follow-on counts if the same lender raises a successor fund of the same strategy within eight years. *Prior Fund Final IRR* is the fund's latest reported net IRR. Columns 1–2 are probit models for whether a follow-on is raised; columns 3–4 are estimated on raisers, with the log ratio of follow-on to prior fund size as the dependent variable. Standard errors are clustered at the lender level, and *, **, *** represent 10%, 5%, and 1% levels of statistical significance.

Follow-On Fundraising, Fund Level

Dependent Variable:	P(raise)		$\log \frac{\text{Fund}_{N+1}}{\text{Fund}_N}$	
Model:	Probit	Probit	OLS	OLS
	(1)	(2)	(3)	(4)
Prior Fund Final IRR	0.010**	0.011*	0.019***	0.016***
	(0.005)	(0.006)	(0.004)	(0.004)
Strategy FEs	Y	Y	Y	Y
Vintage FEs	N	Y	N	Y
N	1,024	1,024	607	607
Adj. R^2 / Pseudo R^2	0.023	0.236	0.012	0.052

Table 8. Model Calibration

Notes: This table reports the calibration of the model in Section 7. Panel A gives the three structural parameters estimated by simulated method of moments: for each candidate, the manager's problem is solved globally, a panel of 2,500 managers is simulated over 18 three-year periods after burn-in, and the paper's specifications are run on the simulated panel. Panel B lists the five targeted moments, the data value, and the model value at the estimate; each deviation is standardized by the data moment's standard error in the estimation objective (the estimated standard error for the two regression moments, the scale given in the text for the other three). The two interaction moments are the capacity × above-hurdle coefficients in the rescue regression (Table 4, with young-fund capital measured in log dollars) and in the three-year change in seasoned-fund IRR (Table 5, column 3); simulated rescue dollars are mapped to observed rescue loans through a one-parameter event rate. Recognized distress is the non-accrual share of BDC book cost in 2011–2013 (Section 3), scaled to the model's three-year period; its scale is the dispersion across the three years. The deployment ratio is committed capital called by fund age three across PitchBook debt funds (mean 0.76, cross-fund standard deviation 0.22); the model counterpart is total lending relative to committed capital. Panel C reports quantities implied by the estimate in the simulated stationary economy. The ordinary-loan spread is held at its 2011–2013 reference value in this calibration; the upfront-charge component of the base return is set so that the simulated mean reported return equals the 2011–2013 mean net IRR of 10 percent; and distress rates are the measured proxies without correction for exposure that rescue keeps out of the data. All remaining parameters are set externally or at the reference allocation as described in the text.

Panel A. Estimated parameters		
Rescue-cost curvature, γ	0.70	
Deployment wedge, ψ	0.87	
Concealment loading, ℓ	0.75	
Panel B. Targeted moments		
	Data	Model
Rescue: capacity × above-hurdle (per log $mn)	0.0102	0.0156
ΔIRR, 3 years: capacity × above-hurdle (pp)	1.36	2.26
Recognized distress, share of book per period (%)	2.72	5.89
Deployment ratio, mean	0.76	0.76
Deployment ratio, cross-fund SD	0.22	0.38
Panel C. Implied objects		
Rescue intensity, Q_E/E	0.57	
Concealed distress, Q_E/B (%)	8.1	
Recognized distress, $(E - Q_E)/B$ (%)	5.9	
Total distressed exposure, E/B (%)	13.9	
Reported net return, mean (% per year)	10.3	
Share of manager-years above the hurdle	0.57	

Internet Appendix

A Data Appendix and Additional Evidence

This appendix provides details on sample construction and measurement, followed by the coverage analysis and robustness tests summarized in the empirical sections.

A.1 Constructing the Linked Samples

The private-credit transaction universe consists of U.S. dollar-denominated, voluntary debt-purpose loans to U.S. borrowers from PitchBook. We exclude financing associated with buy-outs, mergers and acquisitions, growth and venture investments, recapitalizations, and debtor-in-possession proceedings. These restrictions isolate lending decisions in which the incumbent relationship is economically meaningful rather than mechanically determined by a change in control or a formal bankruptcy process. The complete transaction window runs from 2004 through 2025. Valuation-conditioned analyses begin in 2010, when BDC coverage becomes sufficiently broad.

We link transactions to BDC portfolio filings through a common PitchBook borrower identifier. For each transaction, we aggregate cost and fair value across the borrower's debt positions and reporting BDCs and retain the most recent valuation reported strictly before origination and within the preceding four quarters. The valuation is a borrower-level signal: a BDC must hold some debt of the borrower, but the new loan and its lender need not appear in a BDC portfolio. For borrower-year outcome tests, we instead use the year-start valuation of debt already held by BDCs at the end of year $t-1$ and follow the same debt forward. This construction prevents the new loan from mechanically changing the outcome.

PitchBook manager identifiers link each lender's transaction history to its fund roster, reported returns, and fundraising history. The fund records cover direct lending, mezzanine, distressed debt, general debt, credit special situations, venture debt, and bridge financing. LP identifiers connect managers to investors that previously committed to their funds and route mandate announcements through those established relationships.

Table 1 describes the three principal linked panels. Panel A contains 27,161 borrower-years

from 2010 through 2024 in which existing BDC-held debt has a year-start valuation. Panel B contains 4,534 manager-years that form the at-risk sample for the rescue-lending analysis. Panel C contains 6,397 seasoned fund-years used to study changes in reported performance. Estimation samples can be smaller when an analysis requires reported IRR, a future outcome, or non-singleton fixed-effect groups.

The remaining samples follow the requirements of each test. The borrower–lender matching analysis forms potential pairs of active borrowers and lenders in each year, restricts attention to deals with at most five lenders, and requires a pre-origination borrower valuation. The flow–performance sample contains private-credit funds of at least $50 million with vintages from 2000 through 2017 and defines a successor as a fund of the same strategy raised by the same manager within eight years. The LP reinvestment sample forms LP–manager–year observations and tests whether an existing investor commits again as a function of the manager's seasoned-fund performance.

A.2 Variable Construction

We aggregate BDC-reported fair value and cost across positions using cost weights. *FV/Cost* is the resulting borrower-level fair-value-to-cost ratio. The main impairment indicator equals one when FV/Cost is below 0.90. In specifications that use impairment depth, *Impairment* equals $1 - \text{FV}/\text{Cost}$, set to zero at par. BDC valuations are reported estimates rather than transaction prices, and the BDC providing the valuation need not be the lender making the new loan. We therefore use the valuation to measure borrower condition, not to attribute a valuation choice to the transaction lender.

An existing lender is a lender that participated in the borrower's most recent active lending year before year t; all other participating lenders are new. Rescue lending is new financing from an existing lender to a borrower whose pre-origination FV/Cost is below 0.90. This definition does not distinguish an efficient workout from evergreening. The borrower-outcome tests address advantageous selection into rescue, but they do not estimate the treatment effect of rescue on borrower value.

The fund-state measures are constructed annually. A young fund is one to four years old, and a seasoned fund is at least five years old. A seasoned fund is above hurdle when its reported net IRR is at least 10 percent. Contractual hurdles are not observed, so the cutoff is an empirical proxy rather than a fund-specific contractual threshold. The measures identify the manager's cross-fund state; the data do not assign each rescue loan to a particular young or seasoned fund. Reported net IRR is observed only for reporting funds and can change because of cash flows, valuations, or both. We therefore interpret IRR as reported performance, the object visible to investors during fundraising.

The mandate exposure counts private-credit mandates announced over the preceding one or three years by LPs whose first commitment to the manager occurred at least five years earlier, and standardizes the count within the estimation sample. The fixed relationship window prevents current fundraising from adding investors to the exposure measure. Appendix Table B.1 reports the associated fund-closing first stage and its timing and asset-class falsification tests.

A.3 Public-Pension Allocations and Commitments

This section documents the two pension-level samples in Table 6: the policy-target panel of Panel A and the commitment panel of Panel B. Both start from the Public Plans Database (PPD), an annual panel of U.S. state and local defined-benefit plans that reports each plan's total assets, cash flows, funded status, and asset-class allocation targets and actual weights. The PPD defines the set of investors; PitchBook supplies their commitments and the asset-class return benchmarks.

Linking pensions to PitchBook. The PPD covers 195 plans, which consolidate into 193 investment pools because some plans share an investment office. We match each pool to its PitchBook LP identity in three passes: curated aliases for treasury and investment-office names (for example, the Oregon State Treasury invests for Oregon PERS); exact and token-level string matches; and a guarded fuzzy match that requires agreement on distinctive tokens and headquarters state. We then run unmatched pools through PitchBook's entity lookup and verify each match by hand against the raw commitment records. The result is 161 pools with a PitchBook LP record

that carries commitments, holding 92 percent of PPD assets. Six further pools have an LP record with no commitments and are excluded as true zeros; the remainder have no PitchBook record. Where PitchBook carries two records for the same money (for example, a retirement system and its treasury department), we union the records and deduplicate at the fund level so that no commitment is counted twice.

Commitment panel (Panel B). The unit is a pension × asset class × commitment year. We date commitments by the commitment decision recorded in PitchBook, never by the fund's vintage year: vintage is a fund characteristic that can precede or follow the pension's decision by more than a year, so dating by vintage does not produce a valid flow test. Where PitchBook records only the commitment year, we take the pension's information set to be the end of the prior calendar year and apply the same convention to every commitment, so no observation uses same-year information. The four asset classes are PitchBook's classifications of the fund committed to: private equity, venture capital, private credit, and real estate. Other private classes have no benchmark return and are excluded, and PPD's commodities and real-assets bucket is never mapped to PitchBook real assets. The window is commitment years 2009–2024: PitchBook's asset-class index returns begin in 2008, so the lagged signal first exists in 2009, and we drop 2025 for reporting lag. The regressor is PitchBook's trailing four-quarter index return for the asset class through December 31 of $t-1$, a market benchmark common to every pension rather than any pension's own realized return. Because the signal varies only at the asset class-by-year level, it takes 64 distinct values, and inference rests on the 16 commitment-year clusters rather than on the number of cells; standard errors are two-way clustered by pension and commitment year. The sample waterfall runs from 23,622 pool-class-years for the 161 pools, to 15,018 in the window, 12,515 with a mapped benchmark, and 10,012 in the four analysis classes. The estimation samples are 10,008 for the commitment indicator; 7,656 for the class's share of the pension's commitment count, which is undefined in pension-years with no commitment; and 7,200 for the commitment count by Poisson pseudo-maximum likelihood, which drops pension-class cells with no commitments in any year. Pension × asset class fixed effects absorb each pension's standing allocation to a class, and pension × commitment-year fixed

effects absorb the pension's total activity in the year, so the coefficient is identified from whether the pension tilts that year's commitments toward the class whose benchmark just outperformed. Because the panel is balanced across the four classes, replacing the return with the class's return minus the four-class mean leaves the pooled coefficient unchanged.

Robustness of the commitment response. Two checks are not tabulated. First, adding the supply of funds in market and the number of active managers in the class-year, together with differential class trends, leaves the pooled coefficient at 0.113 (standard error 0.037). Second, the class-specific slopes in columns 4–6 are less stable than the pooled response. The private-credit slope is identified largely from 2009, the only year in which the lagged private-credit benchmark is negative: dropping that year reduces the slope by three-fifths, and replacing the return with its deviation from the four-class mean, which the fixed effects do not absorb in the interacted specification, reduces it to 0.063 (0.116). The real-estate slope survives both checks. We therefore rest any private-credit-specific statement on the pooled response and on the policy-target evidence in Panel A rather than on the class-specific slope alone. Appendix Table B.6 reports the moderator analysis. Net cash flow, defined as contributions minus benefit payments over plan assets at $t-1$, and log plan assets are each standardized and entered interacted with the return, because the pension-year fixed effects absorb their levels.

Policy-target panel (Panel A). The unit is a pension-year. Base years are fiscal 2009–2023, and target changes are measured from t to $t+1$, so revisions run through fiscal 2024. The outcome is the change in the pension's private-markets policy target: the sum of the PPD component targets for private equity, private credit, and real estate as a fraction of plan assets. We harmonize these classes from PPD's reported fields. Private equity is PPD's private-equity total net of its private-debt and private-placement components; private credit is the sum of PPD's private-debt and credit-opportunities components; and real estate is PPD's real-estate total. PPD reports no separate venture-capital target, so venture exposure remains inside the private-equity component and cannot receive its own coefficient.

The outcome is the policy target rather than the realized weight because realized weights move

with valuation drift and mechanical rebalancing. After an asset class outperforms, its realized weight rises and is then rebalanced back toward target, so realized weights mean-revert against past class returns and do not isolate the allocation decision. The regressor is the lagged private-minus-public composite return over the four quarters ending at the pension's fiscal year end. The private composite weights the PitchBook benchmark returns of private equity, private credit, and real estate by the pension's own actual allocation weights to those classes at $t - 1$, renormalized to sum to one; the public composite weights the MSCI ACWI net return for public equity and the Bloomberg U.S. Aggregate total return for fixed income by the pension's lagged actual weights to those two classes, renormalized in the same way. Weights are therefore strictly predetermined, and within a given window the spread varies across pensions only through their lagged portfolio mix; a variance decomposition attributes about 87 percent of its variation to the component common to all pensions in the window.

Two reporting regimes are nested. The baseline sample fills a missing component target with zero only when PPD reports that class's actual weight as exactly zero; the strict sample keeps only pension-years in which every component target is explicitly reported at both endpoints. Both regimes drop pension-years in which PPD records a classification change in any component between the two endpoints. Because most pensions never write down a separate private-credit target, the strict sample is effectively the set of plans running an itemized private-credit program: 303 pension-years across 52 pensions, against 1,759 pension-years across 154 pensions in the baseline. The larger strict coefficient therefore partly reflects composition. However, the strict sample is also the less contaminated one, because the baseline imputes a zero target for some plans that demonstrably hold private credit. Pension fixed effects give each pension its own trend in the target level. Standard errors are two-way clustered by pension and return window in the baseline sample; the strict sample is too small for the two-way estimator, and those columns cluster by pension. The controls (net cash flow, funded ratio, and log plan assets, all at t) are individually insignificant in both samples and are suppressed from the displayed table.

Why commitments rather than allocation weights. PPD's asset-class weights cannot measure

private-credit exposure. Among pension-years in which a pension committed to at least one private-credit fund in PitchBook, 74 percent report a PPD private-credit weight of exactly zero with complete class detail, and the zero rate is nearly the same for pensions that did not commit. Private-credit holdings sit inside PPD's private-equity, hedge-fund, fixed-income, or opportunistic buckets and are itemized by only a minority of plans. Commitments dated by decision are therefore the appropriate measure of the investor-side flow to private credit, and the policy target is the appropriate measure of the allocation decision; we use realized PPD weights for neither.

A.4 Coverage of BDC Valuations

Appendix Figure B.2 describes the reach of the BDC valuations used to measure borrower condition. Panel A reports both same-year coverage and the stricter pre-origination coverage used in the transaction analysis. By count, pre-origination coverage rises from 21% of debt instruments in 2010 to 61% in 2024, while same-year coverage rises from 41% to 79%. Pre-origination coverage by debt volume rises from 42% to 57% over the same years. The thin coverage before 2010 motivates beginning the valuation-conditioned analysis in that year.

The covered sample is broader than the set of loans held directly by BDCs. Panel B shows that 42% of debt instruments issued to borrowers with a same-year valuation never appear in BDC holdings; the corresponding share is 47% by dollar volume. The valuation therefore provides a borrower-level signal that extends to loans made outside the reporting BDC portfolios.

Coverage nevertheless selects borrowers connected to the BDC reporting system. Covered debt instruments are smaller on average than uncovered instruments ($39 million versus $68 million), and covered borrowers have more debt issues (3.6 versus 1.6) and distinct lenders (3.0 versus 1.9). Among floating-rate instruments with pricing information, however, average margins are nearly identical at 585 and 584 basis points. The linked sample consequently overrepresents repeat borrowers, an important limitation for interpreting the relationship-lending results, even though its observable pricing closely resembles that of the uncovered sample.

A.5 Borrower Outcomes and Loan Attrition

Appendix Table B.3 examines whether existing loans disappear from BDC holdings after a borrower obtains new financing. The dependent variable is the share of the loans' initial cost no longer reported at each horizon. For impaired borrowers financed by their incumbent, incremental attrition is 10.3 percentage points after one year, 7.7 percentage points after two years, and 6.3 percentage points after three years.

These estimates should not be read as repayment rates. A position can leave the BDC data because of repayment or refinancing, but also because of a sale, write-off, or loss of reporting coverage. The table therefore shows why assigning a favorable outcome to every disappearing position would be inappropriate. In the main outcome analysis, we instead condition changes in FV/Cost and non-accrual on observing the preexisting loans at the relevant horizon.

Appendix Table B.2 restricts the sample to borrower-years financed by exactly one lender type. The omitted category is financing by a new lender, so the interaction directly compares impaired borrowers financed by their incumbent with impaired borrowers refinanced by outsiders. Incumbent-financed impaired borrowers have a 19.4-percentage-point lower relative change in FV/Cost after three years and a 28.8-percentage-point higher relative probability of non-accrual after one year. The smaller sample makes several longer-horizon estimates imprecise, but the qualitative conclusion remains: incumbent rescue is not concentrated among impaired borrowers that subsequently outperform on the observable credit measures.

A.6 Private-Equity Flow–Performance Benchmark

Appendix Table B.5 applies the flow–performance specifications to buyout funds. Prior-fund IRR positively predicts both successor-fund formation and successor-fund growth, including within GP. The benchmark confirms that the relation documented for private credit is a recognizable feature of sequential private-capital fundraising. We do not interpret differences in coefficient magnitudes across the two asset classes because their samples and return distributions differ.

A.7 New-Borrower Lending Placebo

Appendix Table B.4 repeats the four columns of Table 4, Panel B with an indicator for lending to at least one new borrower as the dependent variable. Capacity moves ordinary deployment: the young-fund level effects are 34.0 percentage points unconditionally and 10.5 percentage points with controls and fixed effects. The capacity-times-above-hurdle interactions, by contrast, are 0.4, −8.6, 0.4, and 1.3 percentage points—all statistically indistinguishable from zero against a 54.7% base rate.

The placebo does not imply that fundraising is unrelated to ordinary deployment; it shows the opposite margin responds to capacity alone. The incentive fingerprint—the interaction of capacity with an above-hurdle track record—appears only for rescue, where it amounts to roughly 40% of the outcome mean. The concentration of rescue in incentive-relevant states is therefore difficult to attribute to a general increase in lending capacity.

B Model Derivations

This appendix derives the solution to the model in Section 7. Section B.1 gives the envelope system, Section B.2 the rescue and pricing conditions with their boundary cases, and Section B.3 the stationary equilibrium and the feedback multiplier.

B.1 Envelope System

We suppress time subscripts and mark next-period objects with +. Two conventions from Section 7 are used throughout. The management-fee rate m both reduces the seasoned fund's report and scales the manager's fee income mK, which avoids carrying the seasoned fund's committed capital as a separate state. The deployment cost in Equation (9) is written in terms of the equity-equivalent deviation of deployment from the fund's lifecycle benchmark,

$$\mathrm{dev}^E \equiv \underbrace{\frac{Q_E + Q_N}{1+L}}_{\text{equity supporting current deployment}} - \underbrace{\bar{u}K}_{\text{lifecycle deployment benchmark}}. \tag{13}$$

Define the effective carry threshold per dollar of gross assets as

$$\tau_L \equiv \frac{m + h + Lc}{1+L}. \tag{14}$$

On the above-hurdle branch, carried interest is

$$\phi \frac{B}{1+L}(y - m - h) = \phi(q - \tau_L)B - \phi\ell(E - Q_E), \tag{15}$$

which isolates the three seasoned-fund states that affect carry: the blended return q, the book B, and unrescued distress $E - Q_E$.

Define the equity-scale deployment wedge and the marginal value of committed capital through that wedge:

$$d^{Q} \equiv \frac{\psi}{K}\,\mathrm{dev}^{E}, \qquad d^{A} \equiv \frac{d^{Q}}{1+L}, \qquad d^{K} \equiv \bar{u}\,d^{Q}. \tag{16}$$

The wedge d^{A} is the shadow value of committed capital in ordinary lending that appears as the capital opportunity cost in the rescue condition (10); d^{K} is the value of relaxing the deployment constraint with one more dollar of committed capital.

On a smooth carry branch, the envelope with respect to the preceding report is

$$V_R = \left(m + d^{K} + \frac{\gamma Q_E^2}{2K^2} + \frac{\psi(\mathrm{dev}^{E})^2}{2K^2} \right) F_R. \tag{17}$$

A stronger preceding report raises current committed capital; the marginal capital dollar earns management fees m, moves deployment toward its benchmark by d^{K}, and lowers both cost intensities through their $1/K$ terms.

The remaining envelopes are

$$V_E = -\iota\phi\ell - \beta\,\mathbb{E}\left[V_R^{+}\lambda_R \frac{\ell}{B} \right], \tag{18}$$

$$V_B = \iota\phi(q - \tau_L) + \beta\,\mathbb{E}\left[V_R^{+}\lambda_R \ell \frac{E - Q_E}{B^2} \right], \tag{19}$$

$$V_q = \iota\phi B + \beta\,\mathbb{E}\left[V_R^{+}\lambda_R \right]. \tag{20}$$

True distress lowers current carry and, through the report, future fundraising. A larger book raises carry but dilutes the report effect of a dollar of recognized distress. A higher contractual return raises both current carry and the report. With $\lambda_E \equiv -V_E$, Equations (18) and (20) imply

$$\lambda_E = \ell \frac{V_q}{B}, \tag{21}$$

because a marginal dollar of distress and a marginal change in the blended return reach value through the same reporting channel. The identity is a numerical check on the solved value

derivatives.

B.2 Lending Conditions

Rescue lending. Differentiating the Bellman maximand with respect to Q_E yields Equation (10), whose right-hand side we denote $\mathcal{R}_{Q_E}$. At the boundaries of $[0, E]$, the same residual gives the Kuhn–Tucker conditions

$$\mathcal{R}_{Q_E} \begin{cases} = 0, & 0 < Q_E < E, \\ \leq 0, & Q_E = 0, \\ \geq 0, & Q_E = E. \end{cases} \tag{22}$$

To display the resource-cost structure of the interior condition, define

$$\Gamma_L \equiv \gamma + \frac{\psi}{(1+L)^2}, \qquad \mathcal{K}_L \equiv \frac{\psi}{1+L}\left(\bar{u}K - \frac{Q_N}{1+L}\right). \tag{23}$$

Let $\Delta^{\mathrm{GP}}(Q_E)$ collect current carry relief and all continuation terms in Equation (10). The first-order condition rearranges to

$$Q_E^* = \left[\frac{K\Delta^{\mathrm{GP}}(Q_E^*) + \mathcal{K}_L}{\Gamma_L}\right]_{[0,E]}. \tag{24}$$

Equation (24) is an implicit projection, not a closed-form policy: the marginal benefit depends on rescue through the next state and the continuation derivatives. On a branch where the rescue residual is decreasing in Q_E, the projection is equivalent to the Kuhn–Tucker system in Equation (22). Holding Δ^{GP} and Q_N fixed, the mechanical response to committed capital is

$$Q_K^0 \equiv \left.\frac{\partial Q_E^*}{\partial K}\right|_{\Delta^{\mathrm{GP}}, Q_N} = \frac{\Delta^{\mathrm{GP}} + \psi\bar{u}/(1+L)}{\Gamma_L}. \tag{25}$$

The equilibrium response H_K used in Section 7 and defined in Equation (41) below also includes changes in future reports, states, and continuation values.

Ordinary-loan pricing. The ordinary spread governs deployment, future returns, and new distress, but not the current report, because newly originated loans sit in the young fund and report only once that fund is seasoned. Because

$$\frac{\partial q^+}{\partial r} = \frac{D(r)}{B^+} + \frac{D_r(b + r - q^+)}{B^+}, \qquad \frac{\partial E^+}{\partial r} = (\varepsilon_0 + 2\varepsilon_1 Q_N) D_r, \tag{26}$$

the interior spread condition in the single-manager problem is

$$0 = \underbrace{-d^A D_r}_{\text{deployment value}} + \beta\, \mathbb{E}\left[\underbrace{V_B^+ D_r}_{\text{book-size effect}} + \underbrace{V_q^+ \left\{ \frac{D(r)}{B^+} + \frac{D_r(b + r - q^+)}{B^+} \right\}}_{\text{return-composition effect}} - \underbrace{\lambda_E^+ (\varepsilon_0 + 2\varepsilon_1 Q_N) D_r}_{\text{future-distress effect}} \right]. \tag{27}$$

Since $D_r < 0$, a higher spread reduces originations; the manager weighs the value of deployment against the effect of pricing on the size, return, and risk of the next book. With a bounded rate choice $r \in [\underline{r}, \bar{r}]$, let $\mathcal{R}_r$ denote the right-hand side of Equation (27). The rate policy satisfies

$$\mathcal{R}_r \begin{cases} = 0, & \underline{r} < r < \bar{r}, \\ \leq 0, & r = \underline{r}, \\ \geq 0, & r = \bar{r}. \end{cases} \tag{28}$$

Competition for borrowers. Managers compete for borrowers under differentiated Bertrand competition. Manager i's ordinary lending is

$$D_i(r_i, \mathbf{r}_{-i}; z^d) = \underbrace{\frac{D_0(z^d) - \eta \bar{r}}{N_B}}_{\text{share of aggregate loan demand}} + \underbrace{\frac{\theta}{N_B} \sum_{j \neq i} (r_j - r_i)}_{\text{business stealing}}, \qquad \bar{r} = \frac{1}{N_B} \sum_j r_j. \tag{29}$$

At a symmetric allocation, total demand across the N_B lenders equals $D_0 - \eta r$. The slope relevant to a unilateral pricing decision is

$$e_B = \underbrace{\frac{\eta}{N_B^2}}_{\text{aggregate-demand response}} + \underbrace{\theta\frac{N_B-1}{N_B}}_{\text{business-stealing response}}, \tag{30}$$

so $-e_B$ replaces D_r in Equation (27), and the exact interior residual at date t is

$$\begin{aligned}\mathcal{R}_{r,t} = e_B d_t^A + \beta\,\mathbb{E}_t\Bigg[-e_B V_{B,t+1} + V_{q,t+1}\frac{D_{B,t} - e_B(b + r_t - q_{t+1})}{B_{t+1}} \\ + e_B\lambda_{E,t+1}(\varepsilon_0 + 2\varepsilon_1 D_{B,t})\Bigg],\end{aligned} \tag{31}$$

where $D_{B,t} = (D_{0,t} - \eta r_t)/N_B$ is per-manager ordinary lending at a symmetric allocation.

Rescue-loan pricing. Before imposing the reservation rate, let the rescue spread $r_E \le \bar{r}_E$ be a control. Using Equation (21),

$$\frac{\partial \mathcal{O}}{\partial r_E} = \beta Q_E\,\mathbb{E}\left[\frac{V_q^+}{B^+} - \delta\lambda_E^+\right] = \beta Q_E(1-\delta\ell)\mathbb{E}\left[\frac{V_q^+}{B^+}\right]. \tag{32}$$

Hence $r_E^* = \bar{r}_E$ whenever $Q_E > 0$, $\delta\ell < 1$, and $V_q^+ \ge 0$, with a strict corner when the expectation is positive. The ordinary-rate choice set is restricted below $\bar{r}_E$, preserving the rescue premium over the feasible state space.

Regularity and branches. The solution is computed on a compact state and control space. We impose $F_R > 0$, $D_r < 0$, $\delta(1 + b + \bar{r}_E) < 1$, and $0 \le \varepsilon_0 + \varepsilon_1 Q_N \le 1$ on the feasible range, and bound ordinary lending away from zero. On each smooth carry branch, the Bellman maximand is locally concave in (Q_E, r). A proposed above-hurdle solution is retained only when $y - m - h > 0$ and a below-hurdle solution only when $y - m - h < 0$; equality uses a common subgradient in $[0, 1]$. Under standard Feller and compactness conditions, the Bellman operator is a contraction with a unique bounded value function. The policy need not be unique at kinks or switching boundaries.

B.3 Stationary Equilibrium and Feedback

At a deterministic stationary allocation, the state reproduces itself under the policies Q_E^* and r^*. In the single-manager formulation,

$$K^* = F(R^*), \tag{33}$$

$$Q_N^* = D(r^*), \qquad B^* = Q_E^* + Q_N^*, \tag{34}$$

$$q^* = b + \frac{\bar{r}_E Q_E^* + r^* Q_N^*}{B^*}, \tag{35}$$

$$E^* = \delta(1 + b + \bar{r}_E)Q_E^* + \varepsilon_0 Q_N^* + \varepsilon_1 (Q_N^*)^2, \tag{36}$$

$$y^* = (1 + L)\left(q^* - \ell \frac{E^* - Q_E^*}{B^*}\right) - Lc, \tag{37}$$

$$R^* = y^* - m - \phi \max\{y^* - m - h, 0\}. \tag{38}$$

Under symmetric competition, $Q_N^* = (D_0 - \eta r^*)/N_B$ and $K_i^* = a(1 + \kappa_f R^*)/N_F$. The report derivative of the funding rule (11) then takes two values. In an individual manager's first-order conditions, rival reports are held fixed,

$$F_R^{\text{dev}} = \frac{a(z^f)}{N_F}(\kappa_f + \kappa_g), \tag{39}$$

whereas in a symmetric equilibrium response all reports move together and the zero-sum relative term cancels:

$$F_R^{\text{sym}} = \frac{a(z^f)}{N_F}\kappa_f. \tag{40}$$

Equations (33)–(38) are closed by the stationary envelope equations, the rescue and rate Kuhn–Tucker conditions, and the carry-branch inequality.

The feedback multiplier follows from two local mappings. The exact rescue residual defines the policy response

$$dQ_E = H_K\, dK, \qquad H_K = -\left.\frac{\partial \mathcal{R}_{Q_E}/\partial K}{\partial \mathcal{R}_{Q_E}/\partial Q_E}\right|_{*}, \tag{41}$$

and rescue affects capital through the report:

$$dK = F_C\, dC^0 + F_R^{\text{sym}} \lambda_R \frac{\ell}{B}\, dQ_E = F_C\, dC^0 + a_R\, dQ_E. \tag{42}$$

Substituting Equation (41) into Equation (42),

$$(1 - a_R H_K) dK = F_C\, dC^0, \tag{43}$$

which is Equation (12) with loop gain $\mathcal{G} = a_R H_K$. The derivation shows why the symmetric fundraising derivative excludes κ_g: when all reports move together, relative performance does not change aggregate capital.

The scalar calculation conditions on the remaining stationary objects. For the full equilibrium, collect the stationary equations in $\mathcal{H}(X, C^0) = 0$ with $X = (K, Q_E, r, E, q, B, R)$. The exact local response is

$$\frac{dX^*}{dC^0} = -\left[\frac{\partial \mathcal{H}}{\partial X}\right]^{-1} \frac{\partial \mathcal{H}}{\partial C^0}. \tag{44}$$

The quantitative implementation reports both the scalar loop diagnostic and the response from re-solving the full system after a small change in the funding intercept. At a binding rescue cap $Q_E = E$, the local response is one-sided and remains zero until the cap ceases to bind.

C Additional Figures and Tables

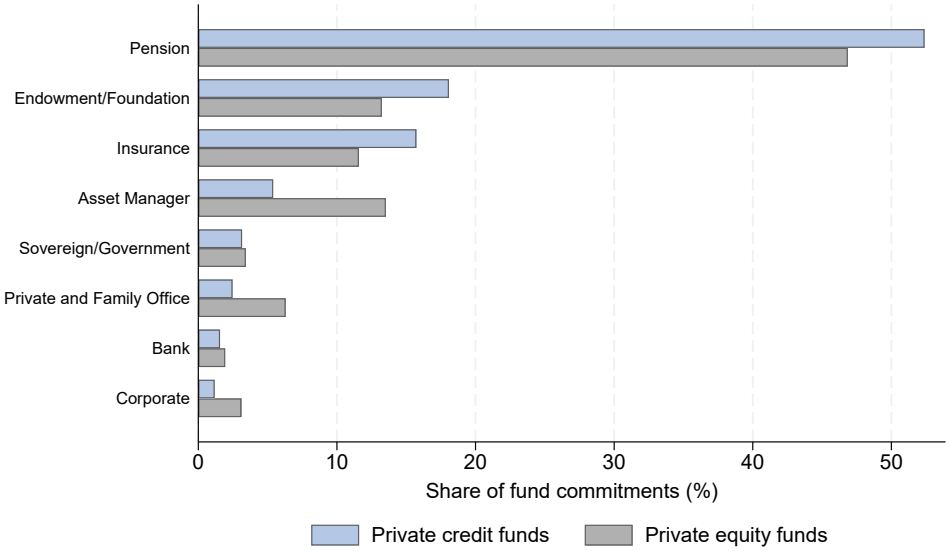
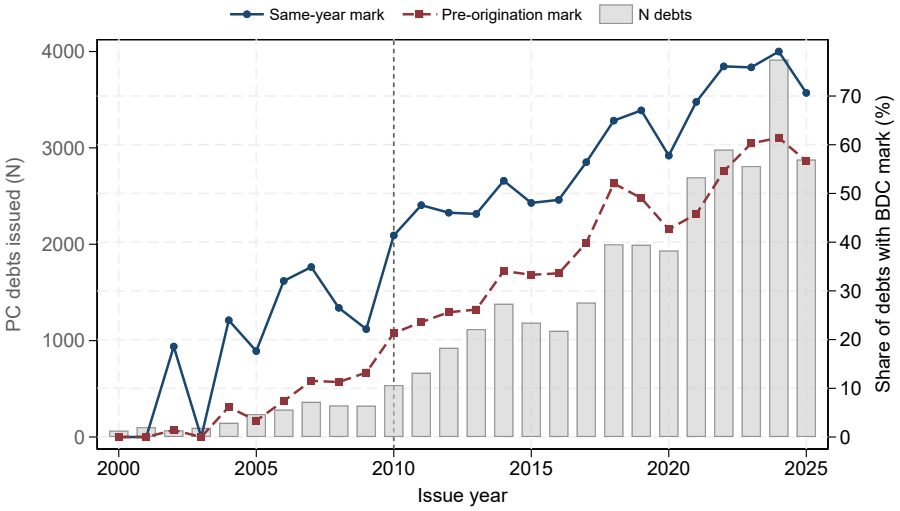


Figure B.1. LP Types of Fund Commitments, Private Credit vs Private Equity. This figure plots the LP-type composition of fund commitments in private credit and in PE. Bars show each broad LP segment's share of LP–fund commitment records, in percent, separately for commitments to private credit funds and to private equity funds (buyout, growth/expansion, diversified private equity, restructuring).

Panel A. Coverage of universe loans, by count



Panel B. Covered borrowers' loans never held by any BDC

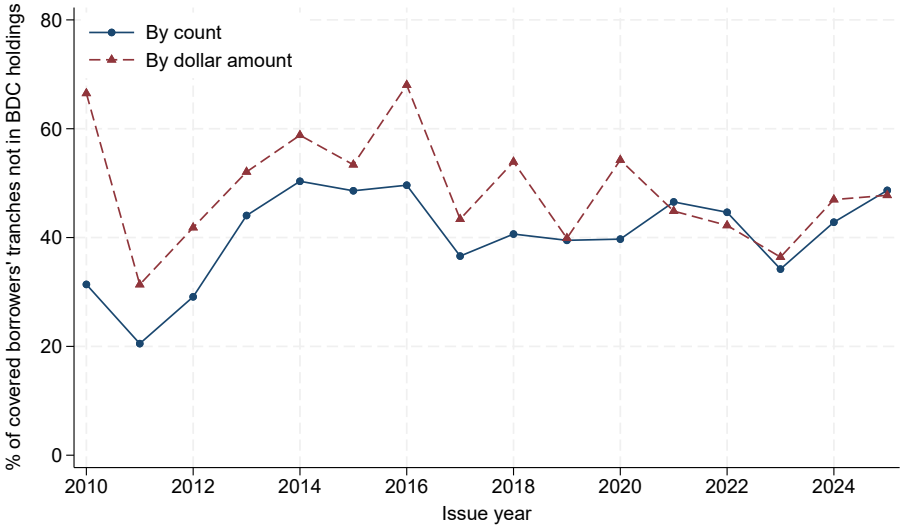


Figure B.2. BDC Valuation Coverage of Private-Credit Transactions. This figure reports the coverage of the BDC valuations used to measure borrower condition. Panel A plots, by issue year, the share of private-credit debt instruments whose borrower has either a same-year valuation or a valuation reported before origination and within the preceding four quarters. Pre-origination coverage rises from 21% in 2010 to 61% in 2024, motivating the 2010 starting date for valuation-conditioned analyses. Panel B shows, among debt instruments issued to borrowers with a same-year valuation, the share that never appears in BDC holdings. Approximately two-fifths of covered instruments fall into this category, showing that the borrower-level valuations extend beyond loans held directly by BDCs.

Table B.1. LP Mandates and Lender Fundraising

Notes: This table reports the relation between related-LP private-credit mandates and lender fundraising — the first stage behind the mandate columns of Tables 4 and 5. The sample is the lender-year panel of Table 4. The dependent variable indicates that the lender closes a debt fund in year t. *Related-LP PC Mandates* counts private-credit mandate announcements over $[t-3, t-1]$ (column 4: $[t-1]$) by LPs first committed to the lender at least five years earlier, standardized. *Future PC Mandates* counts the same LPs' announcements over $[t+1, t+3]$; *Related-LP Non-PC Mandates* counts their non-private-credit announcements over $[t-3, t-1]$. Standard errors are clustered by lender, and *, **, *** represent 10%, 5%, and 1% levels of statistical significance.

Dependent Variable: Closed a Fund

Mandate window:	$[t-1]$			$[t-3, t-1]$		
	(1)	(2)	(3)	(4)	(5)	(6)
Related-LP PC Mandates	0.021***	0.024***	0.021***	0.023***	0.022***	0.022***
	(0.007)	(0.007)	(0.007)	(0.008)	(0.008)	(0.008)
Future PC Mandates		-0.009			-0.011	
		(0.008)			(0.009)	
Related-LP Non-PC Mandates			-0.004			0.002
			(0.007)			(0.007)
Lender FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	0.354	0.354	0.354	0.354	0.354	0.354
N	4,497	4,497	4,497	4,497	4,497	4,497
Lenders	509	509	509	509	509	509
R-squared	0.381	0.381	0.381	0.381	0.381	0.381

Table B.2. Rescued Borrowers, Borrowing Years Only

Notes: This table repeats the Table 3 specification on the subsample of borrower-years that take a new loan from exactly one lender type — either the existing lender or new lenders; the small group borrowing from both in the same year is excluded. The omitted category is borrowing from new lenders, so *Borrow from Existing* compares incumbent-financed borrowers directly against outside-refinanced ones, and *Impaired* × *Borrow from Existing* gives the differential outcome of impaired borrowers financed by their incumbent relative to impaired borrowers refinanced by new lenders. Other variables, fixed effects, and clustering are as in Table 3, and *, **, *** denote 10%, 5%, and 1% statistical significance.

	Δ FV/Cost			Non-Accrual		
	(1)	(2)	(3)	(4)	(5)	(6)
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
Impaired	0.047*	0.098**	0.177***	0.267***	0.228***	0.325***
	(0.026)	(0.040)	(0.057)	(0.045)	(0.065)	(0.099)
Impaired × Borrow from Existing	-0.008	-0.050	-0.194*	0.288***	0.238*	0.035
	(0.051)	(0.082)	(0.112)	(0.082)	(0.128)	(0.196)
Borrow from Existing	-0.007	-0.002	0.016	0.014	-0.001	-0.009
	(0.005)	(0.008)	(0.012)	(0.010)	(0.015)	(0.021)
Maturity	0.003*	0.004	0.009***	-0.010***	-0.012**	-0.018***
	(0.002)	(0.003)	(0.003)	(0.003)	(0.005)	(0.005)
Industry-Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	2,654	1,549	860	2,654	1,549	860
Adj. R^2	0.074	0.169	0.128	0.167	0.145	0.137

Table B.3. Attrition of Existing Loans

Notes: This table repeats the Table 3 specification with attrition as the dependent variable: the share of the existing loans' year $t-1$ cost no longer reported in BDC holdings by year $t+1$, $t+2$, or $t+3$. Exit from BDC reporting mixes full repayment and refinancing away with position sales, write-offs, and the lender leaving BDC coverage, so attrition bounds repayment from above rather than measuring it. The sample, regressors, fixed effects, and clustering are as in Table 3, and *, **, *** denote 10%, 5%, and 1% statistical significance.

	Attrition		
	(1) $t+1$	(2) $t+2$	(3) $t+3$
Impaired	-0.053***	-0.052***	-0.060***
	(0.011)	(0.011)	(0.010)
Impaired × Borrow from Existing	0.103***	0.077**	0.063**
	(0.035)	(0.034)	(0.032)
Impaired × Borrow from New	0.032	0.051**	0.031
	(0.028)	(0.025)	(0.024)
Borrow from Existing	-0.081***	-0.080***	-0.050***
	(0.011)	(0.012)	(0.012)
Borrow from New	0.019**	0.023***	0.014**
	(0.008)	(0.007)	(0.007)
Maturity	-0.076***	-0.075***	-0.064***
	(0.002)	(0.002)	(0.002)
Industry-Year FEs	Yes	Yes	Yes
N	27,161	23,410	19,874
Adj. R^2	0.138	0.171	0.182

Table B.4. Lender Incentives Placebo — Lending to New Borrowers

Notes: This table repeats the four columns of Table 4, Panel B with an indicator for lending to at least one new borrower as the dependent variable: capital deployment should move this margin, while reporting incentives should not. Because ordinary lending is far more common than rescue lending, the comparison across tables is proportional, not in levels. Standard errors are clustered by lender, and *, **, *** denote 10%, 5%, and 1% statistical significance.

Dependent Variable: Lending to New Borrower

Lender Capacity from:	Has Young Fund		LP Mandates	
Mandate window:			$[t-1]$	$[t-3, t-1]$
	(1)	(2)	(3)	(4)
Lender Capacity × Above-Hurdle Seasoned Fund	0.004	-0.086	0.004	0.013
	(0.083)	(0.077)	(0.018)	(0.021)
Lender Capacity	0.340***	0.105*	-0.005	-0.012
	(0.063)	(0.061)	(0.018)	(0.022)
Above-Hurdle Seasoned Fund	0.086	0.081	-0.019	-0.021
	(0.100)	(0.094)	(0.064)	(0.064)
Below-Hurdle Seasoned Fund	0.044	0.026	0.003	0.003
	(0.066)	(0.059)	(0.060)	(0.060)
Controls	No	Yes	Yes	Yes
Lender FEs	No	Yes	Yes	Yes
Year FEs	No	Yes	Yes	Yes
Mean dep. var.	0.542	0.547	0.547	0.547
N	1,621	1,584	1,584	1,584
Lenders	226	189	189	189
R-squared	0.058	0.526	0.525	0.525

Table B.5. Flow–Performance, Private Equity Benchmark

Notes: Same specifications as Table 7, for the universe of private equity funds — buyout, growth/expansion, diversified private equity, and restructuring, the benchmark asset class where the flow–performance relation is well established. Panel A repeats the fund-level follow-on regressions (strategy fixed effects are over the four PE strategies). Panel B reports LP-level reinvestment regressions following Hochberg et al. (2014): the sample is LP–lender–years in which the LP commits to some fund and the lender has previously raised from the LP's segment; the dependent variable indicates that the LP commits to a fund of the lender that year; *Seasoned-Fund (Young-Fund) IRR* is the lender's mean reported net IRR in the prior year across its private equity funds aged five or more (under five), and *Existing LP* indicates a prior commitment to the lender; standard errors are two-way clustered at the LP and lender levels.

Panel A. Follow-On Fundraising, Fund Level

Dependent Variable:	P(raise)		$\log \frac{\text{Fund}_{N+1}}{\text{Fund}_N}$	
Model:	Probit	Probit	OLS	OLS
	(1)	(2)	(3)	(4)
Prior Fund Final IRR	-0.001	0.011***	0.009***	0.009***
	(0.002)	(0.002)	(0.002)	(0.002)
Strategy FEs	Y	Y	Y	Y
Vintage FEs	N	Y	N	Y
N	3,531	3,531	2,484	2,484
Adj. R^2 / Pseudo R^2	0.024	0.208	0.011	0.041

Panel B. LP Reinvestment, LP–Lender–Year Level

	(1)	(2)	(3)	(4)
Seasoned-Fund IRR × Existing LP	0.004***	0.004***		0.005***
	(0.001)	(0.001)		(0.001)
Young-Fund IRR × Existing LP			0.001***	0.001*
			(0.000)	(0.000)
Seasoned-Fund IRR	0.000***	0.000*		0.000**
	(0.000)	(0.000)		(0.000)
Young-Fund IRR			0.000***	0.000***
			(0.000)	(0.000)
Existing LP	0.113***	0.111***	0.186***	0.098***
	(0.013)	(0.013)	(0.010)	(0.017)
LP-Year FEs	Yes	Yes	Yes	Yes
Lender FEs	No	Yes	Yes	Yes
N	2,625,921	2,625,916	2,473,070	1,856,177
Adj. R^2	0.131	0.134	0.140	0.144

Table B.6. What Moderates the Pension Commitment Response

Notes: This table asks what moderates the commitment response of Table 6, Panel B. Columns 1–4 interact the lagged asset-class return with the pension's net cash flow; columns 5–8 with log plan assets, each with the same four specifications so the two moderators can be compared directly. Net cash flow is contributions minus benefit payments over total plan assets at $t-1$, from the Public Plans Database; it is operating cash flow, excludes investment income, and 88.9 percent of pension-years are net outflow, so higher values mean less liquidity pressure. Both moderators are standardized, so an interaction is the change in the response per standard deviation. Because pension × year fixed effects absorb any pension-year level, a characteristic can enter only interacted, and the moderator × asset class controls are included because the fixed effects do not absorb them. Columns 3 and 7 add asset class × year fixed effects, which absorb the lagged return itself and identify the interaction alone. Net cash flow moderates the share margin and not the extensive margin; pension size moderates nothing, which forecloses the reading that the response is confined to large plans. Standard errors are two-way clustered by pension and commitment year, and *, **, *** denote 10%, 5%, and 1% statistical significance.

	Net cash flow				Log plan assets			
	Share (1)	Share (2)	Share (3)	Any (4)	Share (5)	Share (6)	Share (7)	Any (8)
Lagged asset-class return	0.0924**	0.0946**		0.1466**	0.1084**	0.1080*		0.1359**
	(0.0409)	(0.0397)		(0.0514)	(0.0468)	(0.0516)		(0.0497)
Return × Net cash flow (sd)	0.1205**	0.0981**	0.1051**	0.0558				
	(0.0489)	(0.0432)	(0.0451)	(0.0465)				
Return × Log plan assets (sd)					-0.0487	-0.0556	-0.0504	0.0437
					(0.0428)	(0.0417)	(0.0423)	(0.0521)
Pension × Asset Class FEs	Y	Y	Y	Y	Y	Y	Y	Y
Pension × Year FEs	Y	Y	Y	Y	Y	Y	Y	Y
Asset Class × Year FEs	N	N	Y	N	N	N	Y	N
Z × Asset Class controls	N	Y	Y	Y	N	Y	Y	Y
N	7,648	7,648	7,648	10,000	7,652	7,652	7,652	10,004
R-squared	0.506	0.507	0.525	0.677	0.505	0.506	0.525	0.678