# Threshold Invariance: Why Covenant Tightness Does Not Protect LPs in Private Credit

April 13, 2026

**Abstract**

In private credit, covenant tightness does not protect limited partners. A fund manager (GP) who privately decides whether to enforce covenant violations has an enforcement cutoff that depends on the management-fee-to-co-investment ratio $\alpha/\delta$ and is invariant to the covenant threshold $\theta$. Tightening the covenant increases violations but not enforcement: every marginal violator is amended. Three extensions develop this result. First, carried interest acts as *conditional* co-investment, effective only above the hurdle rate; a fund that crosses below the hurdle loses carry discipline discretely, creating a cliff in enforcement quality. Second, when GPs compete for LP capital, naive LPs generate a race to maximum tightness with minimum co-investment: "covenant theater" that maximizes the appearance of protection while minimizing enforcement. Third, co-investment dominates both covenant tightness and carry as the robust instrument for LP protection.

*JEL Classification:* G23, G32, G34.

*Keywords:* Private credit, covenants, evergreening, delegated monitoring, fund management fees.

# 1 Introduction

Private credit has grown from \$158 billion in 2010 to roughly \$1.7 trillion in 2024, making direct lending funds a primary source of corporate finance for middle-market borrowers. A defining feature of this market is covenant intensity: virtually all direct lending transactions include maintenance financial covenants, quarterly tests on leverage or interest coverage that the borrower must satisfy. By contrast, over 90% of syndicated leveraged loans are now covenant-lite. Limited partners (LPs) evaluating private credit funds routinely scrutinize covenant tightness as a measure of credit discipline. Tighter covenants, the reasoning goes, give the fund manager (GP) earlier warning and stronger intervention rights, protecting LP capital.

Covenant tightness is irrelevant to LP welfare. The sufficient statistic for enforcement quality is the management-fee-to-co-investment ratio $\alpha/\delta$, not the covenant threshold $\theta$. Tightening the covenant increases violations but does not increase enforcement: every marginal violator falls above the GP's enforcement cutoff and is amended rather than liquidated. LP welfare is therefore invariant to covenant tightness for all thresholds above the enforcement cutoff.

The mechanism is direct. A GP who earns management fees proportional to assets under management (AUM) benefits from amending a covenant violation: the loan stays on the books at face value rather than being marked down to liquidation recovery. A GP who co-invests alongside LP capital bears a share of the continuation risk: if the amended loan subsequently defaults, the GP loses co-invested capital. The GP's enforcement cutoff, the quality level below which the GP liquidates a violating borrower, balances these two forces. Because this tradeoff depends on borrower quality $q$ and the compensation parameters $(\alpha, \delta)$ but not on the covenant threshold $\theta$, the enforcement cutoff is invariant to $\theta$. When the LP tightens the covenant, the new violations have quality near $\theta$, which exceeds the enforcement cutoff. The GP amends every one. Not a single additional loan is liquidated.

Threshold invariance is a separability result: because the GP's enforcement decision depends pointwise on borrower quality and fee parameters but not on the covenant threshold, the threshold drops out of the LP's welfare function. The welfare loss from under-enforcement equals $(1-D)\alpha^2/[8(\alpha+\delta)^2]$ per dollar of lending. At illustrative parameters (2% management fee, 3% co-investment, 20% default recovery), this loss is 1.6% of face value per loan over the loan's life. Applied to the closed-end, invested-capital-fee segment of the \$1.7 trillion private credit market (the scope condition discussed in Section 6), the embedded lifetime loss runs into the tens of billions; the annualized figure depends on portfolio turnover.

A natural question is why covenant thresholds exceed the GP's enforcement cutoff in practice, given that this excess creates no additional LP protection. Three institutional features generate this pattern. First, covenant thresholds in private credit are often set by market convention or credit agreement templates before fund-specific fee terms are finalized; the same leverage covenant appears across funds with different fee structures and hence different enforcement cutoffs. Second, rating agencies and LP advisory committees use covenant tightness as a due diligence metric, creating pressure to write tight covenants regardless of the fee structure that governs enforcement. Third, GPs may signal quality by offering tight covenants, even though enforcement depends on fee incentives rather than the covenant itself. Under any of these explanations, $\theta$ is determined by forces outside the model while $\alpha/\delta$ determines what happens when violations occur.

The paper contributes to three literatures. First, the delegated monitoring literature following Diamond [1984] and Rajan and Winton [1995] treats covenants as devices that incentivize monitoring. In Rajan and Winton [1995], tighter covenants give the lender more opportunities to intervene, strengthening monitoring incentives. This logic holds when the monitor's interests are aligned with investors'. When the monitor earns fees on AUM and privately decides whether to enforce, tighter covenants create more intervention opportunities that the monitor uses for forbearance, not discipline. The present model nests the Rajan and Winton [1995] aligned-interest benchmark as the limiting case $\delta/\alpha \to \infty$, where the GP's enforcement cutoff converges to the LP-optimal cutoff and covenants function as intended.

Second, the covenant design literature, including Gârleanu and Zwiebel [2009] and Donaldson et al. [2025], studies the optimal allocation of control rights through covenants. Gârleanu and Zwiebel [2009] show that tight-covenant-then-waive is the efficient equilibrium outcome under adverse selection: the lender sets a tight covenant to acquire a put option, then waives violations when the borrower reveals favorable information. In the present model, tight-covenant-then-amend is an inefficient outcome of fee-driven agency. The GP amends not because the borrower reveals favorable information, but because amendment preserves AUM. The two models generate the same qualitative prediction (tighter covenants, more waivers) but differ in cross-sectional determinants and welfare implications, providing partially distinguishable testable predictions.

Third, the evergreening literature, including Faria-e Castro et al. [2024] and Acharya et al. [2021], explains forbearance through balance-sheet pressure: banks evergreen to avoid recognizing losses on regulatory capital. Private credit funds mark to model and hold no regulatory capital, so this mechanism does not apply. The present paper provides a fee-driven evergreening mechanism specific to delegated asset management: the GP evergreens to preserve AUM and management fees, not to protect regulatory capital.

Table 1 summarizes the positioning relative to existing work.

Table 1: This paper relative to existing theory.

| Paper | Setting | Mechanism | Prediction |
|---|---|---|---|
| Rajan and Winton [1995] | Bank lending | Covenants incentivize monitoring | Tighter → more enforcement |
| Gârleanu and Zwiebel [2009] | Bank lending | Covenant-then-waive is efficient | Waivers are optimal |
| Faria-e Castro et al. [2024] | Banks | Regulatory capital avoidance | Evergreening ↑ capital scarcity |
| Donaldson et al. [2025] | Bank lending | Covenants vs. collateral trade-off | Excess tightness destroys borrower incentives |
| **This paper** | **Private credit funds** | **AUM fee preservation** | **Tighter → more amendment, not enforcement** |

Section 2 presents the model. Section 3 derives the main results. Section 4 extends the model to a portfolio with carried interest. Section 5 endogenizes covenant tightness through GP competition for LP capital. Section 6 compares the model with Gârleanu and Zwiebel [2009], develops testable predictions, and discusses policy implications. Section 7 concludes.

# 2 Model

## 2.1 Environment

Three dates index time: $t \in \{0, 1, 2\}$. A single loan has face value normalized to 1.

$t = 0$ **(Origination).** A fund manager (GP) originates the loan. The covenant threshold $\theta \in [0, 1]$ is determined by market convention (see Section 1). Higher $\theta$ corresponds to a tighter covenant. Limited partners (LPs) provide capital of 1.

$t = 1$ **(Monitoring).** Borrower quality $q$ follows $U[0, 1]$; the GP observes $q$ privately.

- If $q \geq \theta$: no covenant violation occurs. The loan continues to $t = 2$.
- If $q < \theta$: a covenant violation occurs. The GP chooses to *enforce* (liquidate, recovering $L$) or *amend* (waive the violation; the loan continues to $t = 2$).

LPs observe aggregate NAV but cannot distinguish "no violation" from "violation amended." This reflects the information environment in private credit: LP advisory committees typically meet quarterly and review portfolio-level metrics (aggregate NAV, default rates, sector exposures), not loan-level covenant compliance decisions. Individual amendment decisions are made between quarterly reports and disclosed, if at all, in aggregate form. Even when LPAs grant information rights, the GP controls the timing and framing of loan-level disclosures.

$t = 2$ **(Resolution).** Any loan reaching $t = 2$ either recovers with probability $q$, repaying 1, or defaults with probability $1 - q$, recovering $D$.

**Assumption 1** (Recovery ordering)**.** $0 < D < L < 1$ and $L = (1 + D)/2$.

The liquidation recovery $L$ lies at the midpoint of face value and default recovery. All results extend to $L > (1 + D)/2$; Appendix B verifies this.

Figure 1 summarizes the timing.

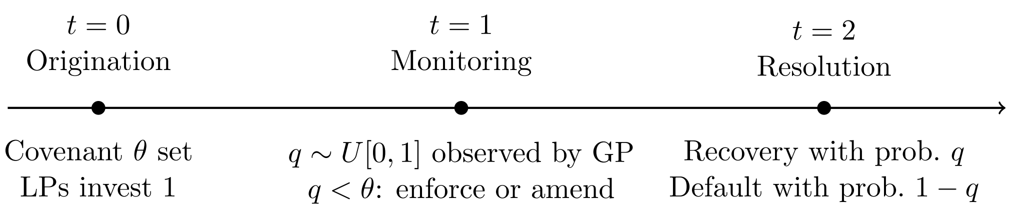


Figure 1: Model timing.

## 2.2 Payoffs

The *continuation value* of a loan with quality $q$ that reaches $t = 2$ is

$$V(q) = q \cdot 1 + (1 - q) \cdot D = q(1 - D) + D. \tag{1}$$

This value is increasing in $q$, with $V(0) = D$ and $V(1) = 1$.

The LP prefers enforcement when $L > V(q)$, that is, when $q < q^*$ where

$$q^* \equiv \frac{L - D}{1 - D} = \frac{1}{2}. \tag{2}$$

The LP-optimal policy enforces all violations with $q < 1/2$ and allows continuation for $q \geq 1/2$.

## 2.3 GP compensation

The GP receives two forms of compensation:

1. *Management fee*: $\alpha \cdot \text{AUM}_t$ per period, where $\alpha > 0$.

2. *Co-investment*: alongside the LP's unit of capital, the GP commits $\delta > 0$ of own capital to the loan, so the loan is financed by $1+\delta$ in total and the GP holds a $\delta/(1+\delta)$ share of the principal. The parameter $\delta$ is the GP's commitment per LP dollar (industry convention; 1–5% in private credit) and is also reported below as the GP's pro-rata exposure for compactness.

The co-investment provides the GP's "skin in the game" without requiring LPs to observe amendment decisions. When the GP amends a loan that subsequently defaults, the GP loses co-invested capital regardless of whether LPs detected the amendment.

**Enforcement payoff.** When the GP enforces at $t = 1$, liquidation proceeds $L$ remain as AUM through $t = 2$. The GP earns management fees on $L$ at both $t = 1$ and $t = 2$, plus recovers $\delta L$ from co-investment:

$$\Pi^E = (2\alpha + \delta)L. \tag{3}$$

**Amendment payoff.** When the GP amends, AUM stays at face value 1 at $t = 1$ and has expected value $V(q)$ at $t = 2$. The GP receives:

$$\Pi^A(q) = \alpha \cdot 1 + (\alpha + \delta)\, V(q) = \alpha + (\alpha + \delta)\big[q(1 - D) + D\big]. \tag{4}$$

**Assumption 2** (Scope conditions)**.** The model applies to closed-end private credit funds where: (i) management fees are charged on net asset value (invested capital), not committed capital; (ii) maintenance covenants give the GP discretion over enforcement; and (iii) liquidation proceeds remain as fee-earning AUM within the fund during the investment period.

Under committed-capital fee structures, the management fee does not depend on AUM, eliminating the fee channel that drives the results. The distinction between invested-capital and committed-capital fees is itself a testable implication of the model.

# 3 Results

## 3.1 The GP's enforcement cutoff

**Proposition 1** (Enforcement cutoff)**.** *The GP enforces a covenant violation if and only if* $q < q_e$, *where*

$$q_e = \frac{\delta}{2(\alpha + \delta)}. \tag{5}$$

*This cutoff satisfies:*

*(i)* $0 < q_e < q^* = 1/2$ *for all* $\alpha, \delta > 0$ *(the GP under-enforces relative to the LP optimum);*

*(ii)* $q_e$ *is independent of the covenant threshold* $\theta$*;*

*(iii)* $\partial q_e/\partial\alpha < 0$ *(higher management fees reduce enforcement);*

*(iv)* $\partial q_e/\partial\delta > 0$ *(higher co-investment increases enforcement).*

*Proof.* The GP amends a violated loan when $\Pi^A(q) \geq \Pi^E$. Using (4) and (3) with $L = (1+D)/2$:

$$\begin{aligned}
\Pi^A(q) - \Pi^E &= \alpha + (\alpha+\delta)\big[q(1-D) + D\big] - (2\alpha+\delta)\frac{1+D}{2} \\
&= \alpha + (\alpha+\delta)q(1-D) + \alpha D + \delta D - \alpha - \alpha D - \frac{\delta}{2} - \frac{\delta D}{2} \\
&= (\alpha+\delta)q(1-D) - \frac{\delta(1-D)}{2} \\
&= (1-D)\left[(\alpha+\delta)q - \frac{\delta}{2}\right]. \qquad (6)
\end{aligned}$$

Since $(1-D) > 0$, the GP amends if and only if $(\alpha+\delta)q \geq \delta/2$, yielding $q \geq q_e = \delta/[2(\alpha+\delta)]$.

Part (i): $q_e > 0$ because $\delta > 0$. The inequality $q_e < 1/2$ requires $\delta < \alpha + \delta$, which holds because $\alpha > 0$.

Part (ii): the expression for $q_e$ contains only $\alpha$ and $\delta$.

Part (iii): $\partial q_e/\partial\alpha = -\delta/[2(\alpha+\delta)^2] < 0$.

Part (iv): $\partial q_e/\partial\delta = \alpha/[2(\alpha+\delta)^2] > 0$. □

**Economic intuition.** When a covenant violation occurs, the GP weighs two forces. Amending preserves AUM at face value rather than marking the loan down to $L$, generating additional management fees. This fee channel scales with $\alpha$. Amending also exposes the GP's co-investment to continuation risk: with probability $1-q$, the borrower defaults and the co-investment recovers only $D < L$. This skin-in-the-game channel scales with $\delta$.

High-quality violators (high $q$) are likely to recover, making the fee benefit dominant. Low-quality violators (low $q$) are likely to default, making the co-investment loss dominant. The enforcement cutoff $q_e$ balances these forces.

The GP under-enforces ($q_e < q^*$) because the management fee benefit of preserving AUM is private to the GP, while the cost of delayed default falls on both the GP's co-investment and LP capital. The GP bears a share $\delta/(1+\delta)$ of any continuation loss (small for industry-standard $\delta \approx 0.03$), while capturing the full management fee on preserved AUM. The fee creates a wedge that pushes the GP toward excessive forbearance.

The enforcement cutoff can be rewritten as

$$q_e = \frac{1}{2 + 2\alpha/\delta}, \tag{7}$$

showing that the ratio $\alpha/\delta$ is the sufficient statistic for the GP's enforcement behavior.

**Numerical example.** With $\alpha = 0.02$ (2% management fee), $\delta = 0.03$ (3% GP co-investment), and $D = 0.2$ (so $L = 0.6$ and $q^* = 0.5$):

$$q_e = \frac{0.03}{2(0.02 + 0.03)} = \frac{0.03}{0.10} = 0.30.$$

The GP enforces only violations with $q < 0.30$. Borrowers with $q \in [0.30, 0.50]$ are "evergreened": their covenant violations are amended despite the LP preferring liquidation.

## 3.2 The covenant tightness trap

**Definition 1** (Amendment and evergreening zones)**.** *For $\theta > q_e$, the amendment zone is $\mathcal{A}(\theta) = [q_e, \theta)$, with probability measure $A(\theta) = \theta - q_e$. The evergreening zone (harmful amendments) is $\mathcal{E}(\theta) = [q_e, \min(\theta, q^*))$, with measure $E(\theta) = \min(\theta, q^*) - q_e$.*

**Proposition 2** (Covenant tightness trap)**.** *For $\theta > q_e$:*

*(i) $dA/d\theta = 1$. Every unit increase in covenant tightness produces one unit increase in the amendment zone.*

*(ii) The enforcement zone $[0, q_e)$ is invariant to $\theta$.*

*(iii) $dE/d\theta = \mathbf{1}\{\theta < q^*\}$. Tightening below $q^*$ increases harmful evergreening one-for-one; tightening above $q^*$ adds only value-neutral amendments.*

*Proof.* Since $q_e$ is independent of $\theta$ (Proposition 1, part ii):

Part (i): $A(\theta) = \theta - q_e$, so $dA/d\theta = 1$.

Part (ii): The enforcement zone is $[0, q_e)$, which does not involve $\theta$.

Part (iii): For $q_e < \theta < q^*$: $E(\theta) = \theta - q_e$, so $dE/d\theta = 1$. For $\theta \geq q^*$: $E(\theta) = q^* - q_e$, constant in $\theta$. □

**Economic intuition.** Consider tightening the covenant from $\theta$ to $\theta + \epsilon$. This step adds a mass $\epsilon$ of new violations with quality $q \in [\theta, \theta + \epsilon)$. These marginal violators have quality near $\theta > q_e$, placing them squarely in the amendment zone. The GP amends every one. The GP liquidates none of them.

The result reverses the conventional wisdom from Rajan and Winton [1995]. In their framework, tighter covenants give the lender more opportunities to intervene, strengthening monitoring incentives. This logic holds when the monitor's interests are aligned with investors'. When the monitor earns fees on AUM and privately decides whether to enforce, tighter covenants create more intervention opportunities that the monitor uses for forbearance, not discipline.

**The case $\theta \leq q_e$.** When the covenant is set below the GP's enforcement cutoff, every violation ($q < \theta \leq q_e$) falls in the enforcement zone. The GP enforces all violations, achieving a 100% enforcement rate conditional on a violation. This higher rate is purely mechanical: a loose covenant restricts the GP from intervening in the interval $[\theta, q_e]$ where the GP *would* have enforced. Total LP welfare under $\theta < q_e$ is strictly lower than under $\theta \geq q_e$, by $\int_\theta^{q_e} (L - V(q))\, dq > 0$. The LP-best covenant within the model is therefore $\theta = q_e$; any tighter $\theta$ adds amendment volume without changing welfare, and any looser $\theta$ forgoes efficient liquidations. Even the LP-best covenant produces under-enforcement relative to first-best because $q_e < q^*$.

**Numerical example.** Using $q_e = 0.30$ and $q^* = 0.50$:

Table 2: Covenant tightness and enforcement outcomes.

| $\theta$ | Violations | Enforced | Amended | Harmful amend. | Enforcement rate |
|---|---|---|---|---|---|
| 0.20 | 20.0% | 20.0% | 0.0% | 0.0% | 100.0% |
| 0.30 | 30.0% | 30.0% | 0.0% | 0.0% | 100.0% |
| 0.40 | 40.0% | 30.0% | 10.0% | 10.0% | 75.0% |
| 0.50 | 50.0% | 30.0% | 20.0% | 20.0% | 60.0% |
| 0.60 | 60.0% | 30.0% | 30.0% | 20.0% | 50.0% |
| 0.70 | 70.0% | 30.0% | 40.0% | 20.0% | 42.9% |
| 0.80 | 80.0% | 30.0% | 50.0% | 20.0% | 37.5% |

For $\theta \leq 0.30 = q_e$, the covenant works perfectly: all violations are enforced. For $\theta > 0.30$, each additional violation is amended. The enforcement rate collapses from 100% to 37.5% as $\theta$ increases, but the volume of enforcement stays fixed at 30%. Figure 2 illustrates the decomposition.

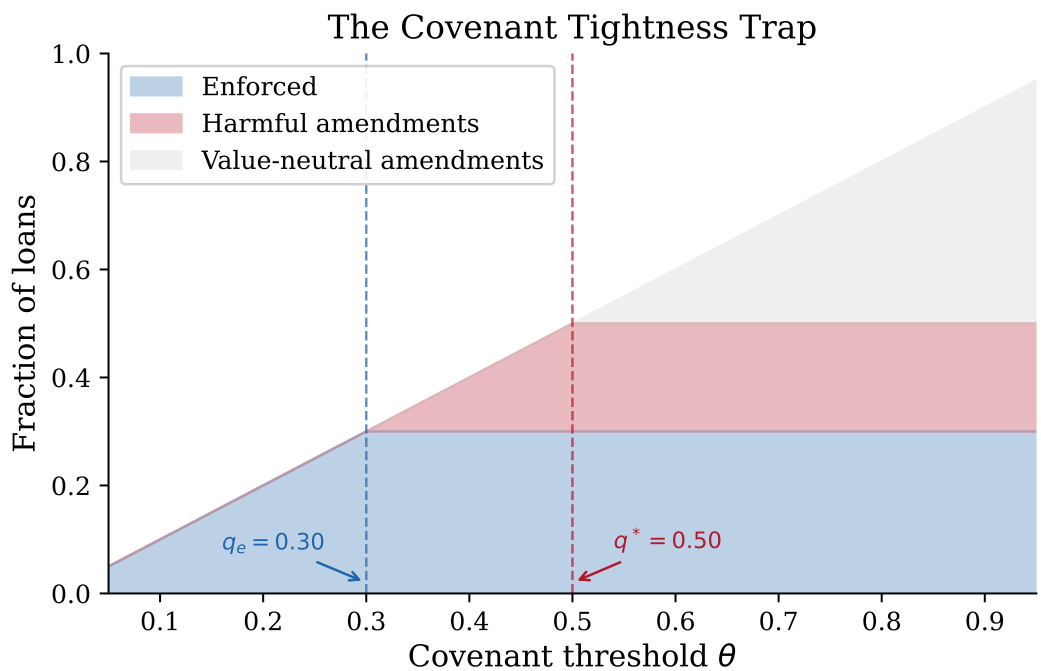


Figure 2: The covenant tightness trap. As the covenant threshold $\theta$ increases, the enforced fraction (blue) stays flat at $q_e = 0.30$ while harmful amendments (red) and value-neutral amendments (gray) grow. Tighter covenants produce more violations but no additional enforcement.

### 3.3 Threshold invariance of LP welfare

**Proposition 3** (Threshold invariance)**.** *The gross fund return per LP dollar (and, holding fees fixed, the LP's net payoff) is independent of* $\theta$ *for all* $\theta \geq q_e$*:*

$$W(\theta) = q_e L + \int_{q_e}^{1} V(q)\, dq \tag{8}$$

*for all* $\theta \geq q_e$*. This payoff is strictly less than the first-best:*

$$W^{FB} = q^* L + \int_{q^*}^{1} V(q)\, dq. \tag{9}$$

*Proof.* Given enforcement cutoff $q_e$ and covenant threshold $\theta \geq q_e$, the LP's payoff decomposes as:

$$
\begin{aligned}
W(\theta) &= \underbrace{\int_0^{q_e} L\,dq}_{\text{enforced}} + \underbrace{\int_{q_e}^{\theta} V(q)\,dq}_{\text{amended}} + \underbrace{\int_{\theta}^{1} V(q)\,dq}_{\text{no violation}} \\
&= q_e L + \int_{q_e}^{1} V(q)\,dq. \qquad (10)
\end{aligned}
$$

The $\theta$-dependent boundary between "amended" and "no violation" is irrelevant because both produce the same continuation payoff $V(q)$. The only boundary that matters is $q_e$, which separates enforcement (payoff $L$) from continuation (payoff $V(q)$).

For the welfare comparison:

$$
W^{FB} - W = \int_{q_e}^{q^*} \big[L - V(q)\big]\,dq > 0,
$$

since $L > V(q)$ for $q < q^*$ by definition of $q^*$, and $q_e < q^*$ by Proposition 1(i). □

**Welfare loss in closed form.** Under $U[0,1]$ with $V(q) = q(1-D)+D$ and $L = (1+D)/2$:

$$
\begin{aligned}
W^{FB} - W &= \int_{q_e}^{q^*} \left[\frac{1-D}{2} - q(1-D)\right] dq = (1-D)\int_{q_e}^{1/2} \left(\frac{1}{2} - q\right) dq \\
&= \frac{(1-D)}{2}\left(q_e - \frac{1}{2}\right)^2. \qquad (11)
\end{aligned}
$$

The welfare loss is proportional to the squared distance between the GP's enforcement cutoff and the LP-optimal cutoff. Substituting $q_e = \delta/[2(\alpha+\delta)]$:

$$
W^{FB} - W = \frac{(1-D)\,\alpha^2}{8(\alpha+\delta)^2}. \qquad (12)
$$

**Terminology.** The term "threshold invariance" distinguishes the present result from the covenant irrelevance results in Davydenko et al. [2020], which concern default timing under individual rationality. The present result concerns the irrelevance of covenant tightness to welfare under agency conflicts, a different mechanism.

**Economic intuition.** The LP pays a welfare cost not because the covenant is too tight, but because enforcement is too lax. Covenant tightness determines how many violations occur, but all violations above $q_e$ are amended, producing the same outcome as if no violation had occurred. The LP's loss comes entirely from the zone $[q_e, q^*]$ where the GP amends loans that the LP would prefer liquidated.

An LP who evaluates funds by comparing covenant tightness is looking at the wrong variable. Two funds with $\theta = 0.4$ and $\theta = 0.8$ but identical fee structures deliver identical LP welfare. The LP should instead scrutinize the fee and co-investment structure $(\alpha, \delta)$.

**Numerical example.** With $q_e = 0.30$, $q^* = 0.50$, $L = 0.6$, and $D = 0.2$:

$$W(\theta) = 0.30 \times 0.6 + \int_{0.30}^{1} \big[0.8q + 0.2\big]\, dq = 0.18 + 0.504 = 0.684.$$

First-best welfare: $W^{FB} = 0.50 \times 0.6 + \int_{0.50}^{1} [0.8q + 0.2]\, dq = 0.30 + 0.40 = 0.70$. The welfare loss is 0.016, or 1.6% of face value, regardless of $\theta$.

## 3.4 Optimal contracting

**Proposition 4** (Optimal fee structure)**.** *LP welfare is increasing in the GP's co-investment ratio $\delta/\alpha$:*

*(i) $dW/d\delta > 0$ for all finite $\delta$ (more co-investment always improves LP welfare).*

*(ii) The welfare loss $W^{FB} - W = (1 - D)\alpha^2/[8(\alpha + \delta)^2]$ vanishes as $\delta/\alpha \to \infty$.*

*(iii) Reducing the management fee to $\alpha' < \alpha$ while holding $\delta$ fixed raises $q_e$ toward $q^*$ and improves LP welfare.*

*(iv) The LP's preferred contract sets $\alpha$ as low as feasible and $\delta$ as high as feasible. First-best requires $\alpha = 0$ (pure co-investment).*

*Proof.* Part (i): From (8),

$$\frac{dW}{dq_e} = L - V(q_e).$$

Since $q_e < q^*$, we have $V(q_e) < V(q^*) = L$, so $dW/dq_e > 0$. By Proposition 1(iv), $\partial q_e/\partial\delta > 0$. The chain rule gives $dW/d\delta > 0$.

Part (ii): From (12), as $\delta/\alpha \to \infty$, $(\alpha + \delta)^2 \to \infty$, so the loss vanishes.

Part (iii): By Proposition 1(iii), $\partial q_e/\partial\alpha < 0$, so reducing $\alpha$ raises $q_e$ toward $q^*$. The chain rule and $dW/dq_e > 0$ give the welfare improvement.

Part (iv): From (12), the loss is minimized by minimizing $\alpha$ and maximizing $\delta$. Setting $q_e = q^* = 1/2$ requires $\delta/[2(\alpha + \delta)] = 1/2$, which implies $\alpha = 0$. □

**Why not contract on enforcement directly?** A natural question is why the LP does not simply specify a contractual enforcement rule ("enforce all violations with $q < q^*$") rather than relying on co-investment to align incentives. The answer is that $q$ is the GP's private information: borrower quality at the time of a covenant violation is a judgment call based on soft information (management quality, market conditions, restructuring feasibility) that cannot be verified by a court or an auditor. The covenant threshold $\theta$ is contractible because it is based on hard financial ratios, but the enforcement decision conditional on violation requires the GP's private assessment of $q$. Co-investment aligns the GP's incentive to use this private information correctly; it does not require $q$ to be contractible.

**Economic intuition.** The LP's problem is not that the covenant is wrong. The problem is that the fee structure creates a wedge between the GP's enforcement incentive and the LP's enforcement preference. The covenant determines the scope of the GP's discretion but cannot direct how that discretion is exercised. Co-investment is the sharp instrument: larger $\delta$ makes the GP bear more of the continuation risk, aligning the GP's enforcement cutoff with the LP's.

The result implies a complementarity between covenants and incentive design. A tight covenant with low co-investment (the current industry standard) produces the illusion of protection. High co-investment with a moderate covenant achieves actual protection.

**Connection to Rajan and Winton [1995].** The model nests the Rajan-Winton aligned-incentive benchmark as a limiting case. As $\delta/\alpha \to \infty$, the GP's co-investment cost dominates the fee incentive, $q_e \to q^* = 1/2$, and the GP enforces at the LP-optimal cutoff. In this limit, tighter covenants produce more enforcement (up to $\theta = q^*$), exactly as Rajan and Winton [1995] predict. The distortion arises specifically when $\alpha/\delta$ is large.

**Numerical example.** With $\alpha = 0.02$ and $D = 0.2$:

Table 3: Co-investment and enforcement quality.

| $\delta$ | $q_e$ | Evergreening zone | Welfare loss | Loss ($ per $1,000) |
|---|---|---|---|---|
| 0.01 (1%) | 0.167 | 0.333 | 0.0444 | $44.4 |
| 0.03 (3%) | 0.300 | 0.200 | 0.0160 | $16.0 |
| 0.05 (5%) | 0.357 | 0.143 | 0.0082 | $8.2 |
| 0.10 (10%) | 0.417 | 0.083 | 0.0028 | $2.8 |
| 0.20 (20%) | 0.455 | 0.045 | 0.0008 | $0.8 |

Industry-standard co-investment of 3% produces a 20-percentage-point evergreening zone and a welfare loss of 1.6% of face value. Raising co-investment to 10% would reduce the welfare loss by roughly 82% (from 1.6% to 0.28% of face value).

# 4 Portfolio Model with Performance Fees

The baseline model omits carried interest and analyzes a single loan. This section extends the model to a portfolio of $N$ loans with carry and shows that (i) threshold invariance survives, (ii) carry acts like *conditional* co-investment that is effective only when the fund exceeds its hurdle rate, and (iii) carry provides no enforcement discipline precisely when the portfolio is distressed, creating a cliff effect at the hurdle.

## 4.1 Setup

The GP manages a fund with $N$ loans, each with face value $1/N$, so total fund size is 1. Loan $j$ has quality $q_j$ drawn i.i.d. from $U[0,1]$. The GP observes each $q_j$ privately and makes loan-by-loan enforcement decisions. Covenant threshold $\theta$ applies to all loans.

The GP's compensation adds a carry component:

1. Management fee: $\alpha \times$ AUM (as before).
2. Co-investment: fraction $\delta$ of own capital, pro rata.
3. **Carry:** $\beta \times \max(R_{\text{fund}} - h, 0)$, where $R_{\text{fund}}$ is the fund's gross return per dollar lent and $h$ is the hurdle rate, both expressed in the model's normalization (recovery cash flows only, with face value normalized to 1). Industry carry is typically 20% over an 8% total-return hurdle; in the model's recovery-cash-flow units, the corresponding $h$ depends on the omitted coupon component. We treat $\beta$ and $h$ as primitives to be calibrated.

## 4.2 Fund return and the large-$N$ limit

The fund's gross return under enforcement cutoff $q_e$ is

$$R_{\text{fund}} = \frac{1}{N}\left[\sum_{\text{enforced}} L + \sum_{\text{continuing}} V(q_j)\right]. \tag{13}$$

By the law of large numbers, as $N \to \infty$:

$$R_{\text{fund}} \to q_e L + \int_{q_e}^{1} V(q)\,dq = W(q_e), \tag{14}$$

where $W(q_e)$ is exactly the LP welfare function from Proposition 3. The fund return and LP welfare coincide in the large-$N$ limit.

## 4.3 The enforcement cutoff with carry

In the large-$N$ limit, each loan has a negligible effect on the fund return. The GP's marginal benefit from amending loan $j$ at quality $q$ (rather than enforcing) is:

$$\Delta\Pi_j = \frac{1}{N}\left\{\alpha(1-L) + \left(\alpha + \delta + \tilde{\beta}\right)[V(q) - L]\right\}, \tag{15}$$

where $\tilde{\beta} \equiv \beta \cdot \mathbf{1}\{W(q_e) > h\}$ is the *effective carry sensitivity*: the carry rate when the fund is above the hurdle, zero otherwise. The $\alpha(1-L)$ term captures the management fee benefit of keeping AUM at face value rather than marking down to $L$. The $(V(q) - L)$ term captures the effect on the GP's residual claim (co-investment plus carry) from continuation rather than liquidation.

Setting $\Delta\Pi_j = 0$ and following the same algebra as Proposition 1, with $\delta$ replaced by $\delta + \tilde{\beta}$:

**Proposition 5** (Enforcement cutoff with carry)**.** *In the large-$N$ portfolio model with carry rate $\beta$ and hurdle $h$, the GP's enforcement cutoff is*

$$q_e^C = \frac{\delta + \tilde{\beta}}{2(\alpha + \delta + \tilde{\beta})}, \tag{16}$$

*where $\tilde{\beta} = \beta \cdot \mathbf{1}\{W(q_e^C) > h\}$.*

*This cutoff satisfies:*

*(i) $q_e^C \geq q_e$ (carry weakly increases enforcement).*

*(ii) $q_e^C$ is independent of $\theta$ (threshold invariance survives).*

*(iii) Below the hurdle ($W(q_e^C) \leq h$): $\tilde{\beta} = 0$ and $q_e^C = q_e$ (carry has no effect).*

*(iv) Above the hurdle ($W(q_e^C) > h$): carry acts as additional co-investment of magnitude $\beta$, and the sufficient statistic for enforcement becomes $\alpha/(\delta + \beta)$.*

*Proof.* The marginal amendment benefit (15) is linear in $q$ with the same structure as the baseline, replacing $\delta$ with $\delta + \tilde{\beta}$. The derivation of $q_e$ in Proposition 1 applies with this substitution, yielding (16).

(i) Since $\tilde{\beta} \geq 0$, the function $g(x) = x/[2(\alpha + x)]$ is increasing in $x$, so $q_e^C = g(\delta + \tilde{\beta}) \geq g(\delta) = q_e$.

(ii) The fund return $W(q_e^C) = q_e^C L + \int_{q_e^C}^1 V(q)\,dq$ depends on $q_e^C$ but not on $\theta$ (by the same argument as Proposition 3: the boundary between "amended" and "no violation" does not affect the integrand). Therefore $\tilde{\beta}$ does not depend on $\theta$, and neither does $q_e^C$.

(iii) When $W(q_e^C) \leq h$, $\tilde{\beta} = 0$ by definition. Then (16) reduces to the baseline.

(iv) When $W(q_e^C) > h$, $\tilde{\beta} = \beta$. Then $q_e^C = (\delta+\beta)/[2(\alpha+\delta+\beta)]$ and the GP's enforcement behavior depends on $\alpha/(\delta+\beta)$. □

**Why threshold invariance survives.** The carry payoff $\beta \max(W(q_e) - h, 0)$ depends on the fund return $W(q_e)$, which by Proposition 3 does not depend on $\theta$. Carry introduces a portfolio-level incentive, but because the portfolio return itself is invariant to covenant tightness, the carry incentive is also invariant. The GP's enforcement decision for loan $j$ depends on $q_j$, $\alpha$, $\delta$, and whether the fund is above the hurdle, but never on $\theta$.

## 4.4 The cliff effect

**Proposition 6** (Carry cliff)**.** *The enforcement cutoff $q_e^C$ is discontinuous at the hurdle: if a small deterioration in portfolio quality pushes the fund return from $W > h$ to $W \leq h$, the enforcement cutoff drops discretely from $(\delta + \beta)/[2(\alpha + \delta + \beta)]$ to $\delta/[2(\alpha + \delta)]$.*

*Proof.* Below the hurdle, $\tilde{\beta} = 0$ and $q_e^C = \delta/[2(\alpha + \delta)]$. Above the hurdle, $\tilde{\beta} = \beta$ and $q_e^C = (\delta + \beta)/[2(\alpha + \delta + \beta)]$. The gap is:

$$\Delta q_e \equiv \frac{\delta + \beta}{2(\alpha + \delta + \beta)} - \frac{\delta}{2(\alpha + \delta)} = \frac{\alpha\beta}{2(\alpha + \delta)(\alpha + \delta + \beta)} > 0. \tag{17}$$

This gap is positive and discrete: the cutoff jumps by $\Delta q_e$ when the fund crosses the hurdle. □

Figure 3 illustrates the discontinuity.

**Equilibrium selection.** Because $\tilde{\beta}$ depends on $\mathbf{1}\{W(q_e^C) > h\}$ and $W(q_e^C)$ in turn depends on $\tilde{\beta}$, equation (16) is a fixed point rather than a closed-form expression. For $h \in (W(q_e), W(q_e^C))$ both the low-cutoff equilibrium ($\tilde{\beta} = 0$, $q_e^C = q_e$, $W \leq h$, self-consistent) and the high-cutoff equilibrium ($\tilde{\beta} = \beta$, $q_e^C > q_e$, $W > h$, self-consistent) coexist. The cliff in Proposition 6 should be read as a comparison across these self-fulfilling regimes: a fund that begins life expecting to clear its hurdle enforces strictly and earns the carry, while an otherwise identical fund coordinated on the pessimistic equilibrium forgoes carry and reverts to baseline enforcement. Outside this range, the equilibrium is unique. The numerical examples below adopt the high-cutoff equilibrium where it exists.

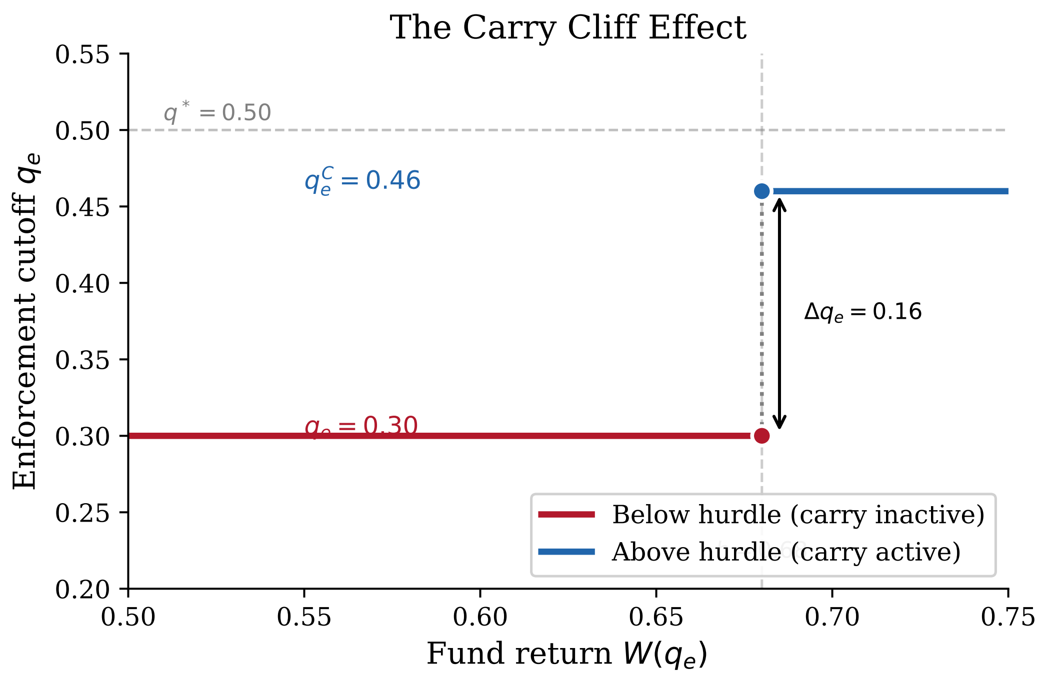


Figure 3: The carry cliff effect. The enforcement cutoff $q_e$ jumps discretely when the fund crosses the hurdle rate $h$. Above the hurdle, carry acts as additional co-investment ($q_e^C = 0.46$). Below the hurdle, carry is zero and enforcement reverts to the baseline ($q_e = 0.30$). The gap $\Delta q_e = 0.16$ represents the discrete loss of enforcement discipline.

**Economic interpretation.** Carry provides enforcement discipline only when the GP stands to earn it. A fund performing above the hurdle has a GP whose effective co-investment is $\delta + \beta$, producing strong alignment. A fund that slips below the hurdle loses this alignment discretely: the GP no longer benefits from marginal improvements in portfolio quality and reverts to the baseline enforcement cutoff driven by co-investment alone. Enforcement deteriorates precisely when the portfolio is struggling, creating a pro-cyclical pattern: good funds enforce well, distressed funds enforce poorly.

**When does carry matter?** At baseline parameters ($\alpha = 0.02$, $\delta = 0.03$, $D = 0.2$), the fund return is $W(q_e) = 0.684$ (from the numerical example in Proposition 3). Note that this measures recovery cash flows only; an 8% real-world total-return hurdle includes coupon income that the model abstracts from. To translate hurdles into the model's units one would set $h$ relative to the maximum achievable $W^{FB} = 0.70$. For any $h$ above $W(q_e^C) \approx 0.70$, carry never activates and the baseline applies.

For carry to matter, the hurdle must be low enough that the fund clears it despite having distressed loans. Since $W(0.30) = 0.684$, any hurdle below 0.684 places the baseline fund above the hurdle, activating carry discipline. With $h = 0.60$:

- With carry active: $q_e^C = 0.46$, welfare loss = 0.064%.
- Without carry (below hurdle): $q_e = 0.30$, welfare loss = 1.6%.

The difference is stark, but the carry-active region is narrow in the model's normalization. Real-world 6–8% total-return hurdles, which include coupon income, fall above the model's recovery-only ceiling $W^{FB} = 0.70$, so on a like-for-like basis distressed portfolios sit firmly below the hurdle and carry provides no enforcement benefit.

## 4.5 Carry does not help when needed most

**Corollary 1** (Asymmetry of carry and co-investment)**.** *Co-investment $\delta$ disciplines enforcement unconditionally: it raises $q_e$ for every portfolio, regardless of performance. Carry $\beta$ disciplines enforcement only conditionally: it raises $q_e$ only for above-hurdle portfolios. Therefore:*

*(i) Co-investment is a robust instrument for LP protection. Carry is fragile.*

*(ii) In the parameter region where under-enforcement causes the largest welfare losses (distressed portfolios with many low-quality loans), carry provides zero enforcement benefit.*

*(iii) The LP's preferred contract raises $\delta$, not $\beta$, to address enforcement distortions.*

**Implication for fund design.** Industry practice treats carry as the primary GP incentive mechanism and co-investment as a secondary "alignment" tool. The model reverses this ranking for enforcement: co-investment is the first-order instrument because it operates unconditionally, while carry is a fair-weather incentive that disappears in distress. LPs who rely on carry for enforcement discipline face a form of wrong-way risk: the incentive vanishes exactly when the portfolio deteriorates.

**Numerical comparison.** Table 4 compares enforcement outcomes across compensation structures.

Table 4: Carry versus co-investment: enforcement outcomes.

| Compensation | $\delta$ | $\beta$ | $q_e$ (above $h$) | $q_e$ (below $h$) | Robust? |
|---|---|---|---|---|---|
| Low co-invest, high carry | 0.01 | 0.20 | 0.457 | 0.167 | No |
| High co-invest, no carry | 0.10 | 0.00 | 0.417 | 0.417 | Yes |
| Balanced | 0.05 | 0.15 | 0.455 | 0.357 | Partial |
| Industry standard | 0.03 | 0.20 | 0.460 | 0.300 | No |

The "low co-invest, high carry" structure looks well-aligned above the hurdle ($q_e = 0.457$, close to $q^*$) but collapses below it ($q_e = 0.167$, massive evergreening). The "high co-invest, no carry" structure delivers $q_e = 0.417$ regardless of fund performance. For enforcement purposes, unconditional co-investment dominates conditional carry.

# 5 Endogenous Covenant Tightness

The baseline model takes the covenant threshold $\theta$ as exogenous. This section endogenizes $\theta$ through GP competition for LP capital. Two GPs compete by offering contracts that specify both covenant tightness and co-investment. The key finding: when LPs evaluate GPs partly on covenant tightness, competition drives covenants to maximum tightness and co-investment to its minimum, producing "covenant theater" that maximizes the appearance of protection while minimizing actual enforcement.

## 5.1 Setup

Two GPs ($i = 1, 2$) compete for capital from a unit mass of LPs. Each GP offers a contract $(\theta_i, \delta_i)$ where $\theta_i \in [0, 1]$ is the covenant threshold and $\delta_i \in [\underline{\delta}, \bar{\delta}]$ is the co-investment share. The management fee $\alpha > 0$ is fixed at the industry norm.

A fraction $\lambda \in [0, 1]$ of LPs are *naive*: they allocate capital to the GP with the tightest covenant. A fraction $1 - \lambda$ are *sophisticated*: they understand threshold invariance and allocate to the GP with the highest co-investment. With $\theta_1 = \theta_2$, naive LPs split equally; with $\delta_1 = \delta_2$, sophisticated LPs split equally.

Co-investment is costly: the GP has convex personal capital costs $c(\delta) = \frac{\psi}{2}\delta^2$ per dollar managed, capturing wealth constraints and risk aversion. GP $i$'s profit from managing capital $K_i$ with co-investment $\delta_i$ is

$$\Pi_i = K_i \left[ 2\alpha - \frac{\psi}{2}\delta_i^2 \right]. \tag{18}$$

Sophisticated LP capital allocation follows a logistic function $s_i = e^{\kappa\delta_i}/(e^{\kappa\delta_i} + e^{\kappa\delta_j})$ with sensitivity $\kappa > 0$, so GP $i$'s total capital is

$$K_i = \lambda \cdot \mathbf{1}\{\theta_i > \theta_j\} + \frac{\lambda}{2} \cdot \mathbf{1}\{\theta_i = \theta_j\} + (1 - \lambda) \cdot s_i(\delta_i, \delta_j), \tag{19}$$

A strict winner in $\theta$ captures all naive capital; ties split equally.

## 5.2 Results

**Proposition 7** (Race to tightness)**.** *In any equilibrium with $\lambda > 0$, both GPs set $\theta_1^* = \theta_2^* = 1$ (maximum tightness). The equilibrium co-investment $\delta^*(\lambda)$ satisfies the first-order condition*

$$\phi\left[2\alpha - \tfrac{\psi}{2}(\delta^*)^2\right] = \tfrac{\psi\delta^*}{2}, \tag{20}$$

*where $\phi \equiv (1-\lambda)\kappa/4$, and is decreasing in the naive LP share $\lambda$:*

$$\frac{d\delta^*}{d\lambda} < 0. \tag{21}$$

*Proof. Covenant tightness.* Fix any $(\delta_1, \delta_2)$. If GP $i$ sets $\theta_i < 1$, GP $j$ can capture all naive capital by setting $\theta_j = 1$. Since tightening is costless to the GP (the enforcement cutoff $q_e$ does not depend on $\theta$ by Proposition 1), deviating to $\theta_i = 1$ weakly increases capital without changing the GP's per-dollar profit. Any $\theta_i < 1$ is weakly dominated. Therefore $\theta_i^* = 1$.

*Co-investment.* With $\theta_1 = \theta_2 = 1$, naive capital splits equally. The GP chooses $\delta_i$ to maximize $\Pi_i = K_i[2\alpha - \psi\delta_i^2/2]$. At the symmetric equilibrium $\delta_1 = \delta_2 = \delta^*$, $K_i = 1/2$ and $\partial s_i/\partial\delta_i = \kappa/4$. The FOC is (20).

*Comparative static.* The FOC defines $\delta^*$ implicitly by $F(\delta^*, \phi) \equiv \phi(\delta^*)^2 + \delta^* - 4\alpha\phi/\psi = 0$. By the implicit function theorem:

$$\frac{d\delta^*}{d\phi} = -\frac{\partial F/\partial\phi}{\partial F/\partial\delta^*} = -\frac{(\delta^*)^2 - 4\alpha/\psi}{2\phi\delta^* + 1}.$$

From the FOC, $4\alpha/\psi = \delta^*(1 + \phi\delta^*)/\phi = \delta^*/\phi + (\delta^*)^2$, so $(\delta^*)^2 - 4\alpha/\psi = -\delta^*/\phi < 0$. The denominator $2\phi\delta^* + 1 > 0$. Therefore $d\delta^*/d\phi = \delta^*/[\phi(2\phi\delta^* + 1)] > 0$. Since $\phi = (1-\lambda)\kappa/4$ is decreasing in $\lambda$, the chain rule gives $d\delta^*/d\lambda < 0$. ☐

**Boundary cases.** At $\lambda = 1$ (all naive): $\phi = 0$, the FOC gives $\delta^* = 0$, so $\delta^* = \underline{\delta}$ (the bound binds). Enforcement is worst: $q_e = \underline{\delta}/[2(\alpha + \underline{\delta})]$. At $\lambda = 0$ (all sophisticated): $\phi = \kappa/4$, competition for sophisticated capital is maximal, and $\delta^*$ approaches $\bar{\delta}$ as $\kappa \to \infty$. Enforcement is best: $q_e \to \bar{\delta}/[2(\alpha + \bar{\delta})]$.

**Corollary 2** (LP sophistication and welfare)**.** *LP welfare $W(q_e(\delta^*(\lambda)))$ is decreasing in the naive share $\lambda$. The welfare loss from LP naivety is*

$$\Delta W(\lambda) = \frac{(1-D)}{2}\left[\left(\frac{\alpha}{2(\alpha + \delta^*(\lambda))}\right)^2 - \left(\frac{\alpha}{2(\alpha + \bar{\delta})}\right)^2\right].$$

*Proof.* $W$ is increasing in $q_e$ (Proposition 3), $q_e$ is increasing in $\delta$ (Proposition 1), and $\delta^*$ is decreasing in $\lambda$ (Proposition 7). The chain gives $dW/d\lambda < 0$. □

**Economic mechanism.** Covenant tightness is a costless signal: threshold invariance guarantees that tightening has no effect on enforcement or LP welfare. Co-investment is a costly signal: it commits GP wealth and shifts the enforcement cutoff. Naive LPs reward the cheap signal; sophisticated LPs reward the costly one. The LP composition determines which competitive force dominates.

When naive LPs dominate, GPs compete on a zero-cost, zero-value dimension (covenant tightness). This is a pure Bertrand game that produces maximum tightness. Since naive LPs do not reward co-investment, GPs minimize it. The result: impressive covenants, minimal enforcement, maximum evergreening. When sophisticated LPs dominate, competition extracts genuine skin-in-the-game, producing near-first-best enforcement.

**Numerical example.** Set $\alpha = 0.02$, $\psi = 1$, $\kappa = 20$, $\underline{\delta} = 0.01$, $\bar{\delta} = 0.20$, $D = 0.2$. Solving the quadratic (20) for each $\lambda$:

Table 5: LP sophistication, co-investment, and enforcement quality.

| $\lambda$ | $\phi$ | $\delta^*$ | $q_e$ | Evergreen zone | Welfare loss |
|---|---|---|---|---|---|
| 0.00 | 5.00 | 0.200 | 0.455 | 0.045 | 0.083% |
| 0.25 | 3.75 | 0.179 | 0.450 | 0.050 | 0.10% |
| 0.50 | 2.50 | 0.146 | 0.440 | 0.060 | 0.14% |
| 0.75 | 1.25 | 0.090 | 0.409 | 0.091 | 0.33% |
| 1.00 | 0.00 | 0.010 | 0.167 | 0.333 | 4.44% |

Figure 4 plots the relationship. Moving from all-sophisticated ($\lambda = 0$) to all-naive ($\lambda = 1$) LPs increases the welfare loss by a factor of roughly 54 (exact: $(0.22/0.03)^2 \approx 53.8$).

The drop is nonlinear: the first 75% of naive LPs raises the welfare loss from 0.083% to 0.33%, while the last 25% raises it from 0.33% to 4.44%.

**Connection to Berk and Green (2004).** In Berk and Green [2004], rational investors compete for capacity-constrained alpha and competition drives expected net alpha to zero in equilibrium. The mechanism here is distinct but analogous: competition for naive capital drives covenant tightness to its maximum and enforcement quality to its minimum. The Berk-Green result is a frictionless rational-expectations outcome; the present result requires a behavioral friction (a fraction of LPs evaluating on a payoff-irrelevant dimension). Both share

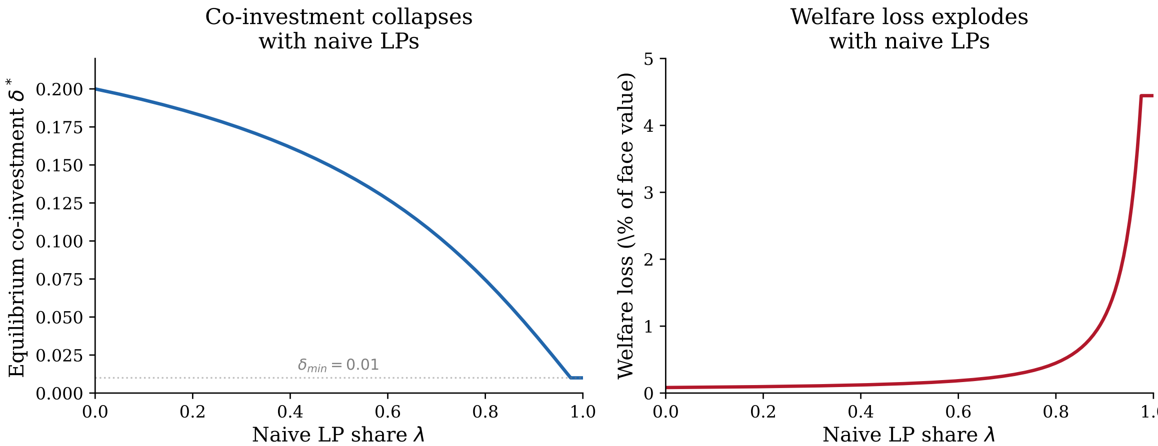


Figure 4: LP sophistication and equilibrium outcomes. Left: equilibrium co-investment $\delta^*$ collapses as the naive LP share $\lambda$ rises. Right: welfare loss (as % of face value) is convex in $\lambda$, exploding as naive LPs dominate.

the structural lesson that competitive pressure dissipates rents along whichever dimension investor capital responds to.

**Testable prediction.** Funds with a larger share of less-experienced LPs (family offices, smaller pensions entering private credit) should have lower co-investment. Funds with sophisticated LPs (large endowments, sovereign wealth funds with dedicated private credit teams) should have higher co-investment. Across sub-markets that differ in LP composition $\lambda$, equilibrium co-investment $\delta^*(\lambda)$ varies negatively with $\lambda$. (Within the basic model, $\theta^* = 1$ for any $\lambda > 0$, so the cross-sectional comparison concerns equilibrium $\delta^*$ rather than the joint pair $(\theta^*, \delta^*)$. Appendix A shows that convex portfolio reputation costs make $q_e$ increasing in $\theta$, which would interior-ize $\theta^*$ if extended to the competition stage; we leave a fully reputation-augmented competition model for future work.)

# 6 Discussion

## 6.1 Comparison with Gârleanu and Zwiebel [2009]

Gârleanu and Zwiebel [2009] (hereafter G-Z) show that tight-covenant-then-waive is the efficient equilibrium outcome of optimal contracting under adverse selection. In their model, the lender sets a tight covenant to acquire a put option on the borrower, then waives violations when the borrower reveals favorable private information. Waivers are efficient: they implement the first-best allocation of control.

In the present model, tight-covenant-then-amend is the inefficient outcome of fee-driven agency. The GP amends not because the borrower reveals favorable information, but because amendment preserves AUM and management fees. Amendments are inefficient: they transfer wealth from LPs to the GP.

**Proposition 8** (Partially distinguishable predictions)**.** *The two models generate the same qualitative prediction (tighter covenants, more waivers/amendments) but differ in cross-sectional determinants:*

*(i) G-Z predicts that waiver rates are higher for borrowers with greater information asymmetry (more opaque borrowers), controlling for fee structure.*

*(ii) The present model predicts that amendment rates are higher for funds with higher management-fee-to-co-investment ratios ($\alpha/\delta$), controlling for borrower opacity.*

*(iii) In G-Z, cross-sectional variation in waiver rates across borrower types should not predict LP returns (waivers are efficient). In the present model, cross-sectional variation in amendment rates across fund types (driven by $\alpha/\delta$) should predict LP returns negatively (amendments are inefficient).*

*Proof.* Part (i) follows from G-Z Proposition 3: the optimal covenant is tighter for borrowers whose private information is more valuable to the lender, and waivers are more frequent when the covenant is tighter.

Part (ii) follows from Proposition 1: $q_e$ depends only on $(\alpha, \delta)$. The amendment zone $A(\theta) = \theta - q_e$ for $\theta > q_e$; since $q_e$ decreases in $\alpha/\delta$, higher $\alpha/\delta$ increases amendment rates for any given $\theta$.

Part (iii): In G-Z, waivers move the allocation toward first-best, so higher waiver rates do not reduce LP value. In the present model, amendments move the allocation away from first-best (Proposition 3), so higher amendment rates imply larger welfare losses. □

**Empirical implementation.** To separate the two mechanisms, regress amendment rates on (a) borrower opacity measures (credit rating availability, audited financials, analyst coverage) and (b) fund fee structure measures ($\alpha$, $\delta$, $\alpha/\delta$). G-Z predicts that (a) is significant and (b) is not. The present model predicts that (b) is significant and (a) is not. If both are significant, the two mechanisms coexist and the magnitudes reveal their relative importance. Because borrower opacity and fund fee structure are likely correlated in practice (private credit by definition involves opaque borrowers), this regression requires careful attention to multicollinearity; instrumental variables based on fund vintage or GP organizational structure may help isolate the fee channel.

## 6.2 Connections to the evergreening literature

The evergreening literature documents forbearance by capital-constrained banks. Faria-e Castro et al. [2024] model bank evergreening driven by regulatory capital incentives: recognizing a loss depletes regulatory capital, so the bank extends distressed loans to avoid the capital hit. Acharya et al. [2021] model zombie lending and policy traps in which weakly capitalized banks under-recognize losses.

The present model provides a distinct mechanism for private credit. The GP is never capital-constrained; the distortion arises from the fee structure alone. A well-capitalized GP with ample co-investment resources still under-enforces as long as $\alpha > 0$, because the management fee creates a private benefit from AUM preservation that does not exist in the LP's payoff. The predictions differ accordingly: bank evergreening increases with capital scarcity, while fund evergreening increases with the $\alpha/\delta$ ratio.

Donaldson et al. [2025] study how covenant design affects control-right allocation, deriving conditions under which excessive covenant tightness destroys borrower incentives. In their model, the lender exercises acquired control rights. In the present model, the lender (GP) does not exercise the additional control rights that tighter covenants provide, rendering covenant tightness irrelevant through a different channel.

## 6.3 Testable predictions

The model generates five testable predictions:

1. **Tighter covenants raise violations but not enforcement volume.** Across private credit funds with similar $\alpha/\delta$ ratios, tighter covenant thresholds produce more violations but the same volume of enforcements. The enforcement *volume* (loans liquidated as a share of the portfolio) is invariant to covenant tightness; the enforcement *rate* (enforcements divided by violations) is mechanically decreasing in tightness.

2. **The $\alpha/\delta$ ratio predicts enforcement.** Funds with higher management-fee-to-co-investment ratios have lower enforcement cutoffs and more evergreening, controlling for covenant tightness.

3. **Amendment rates predict LP returns.** Cross-sectional variation in amendment rates driven by fee structure ($\alpha/\delta$) negatively predicts LP returns. Cross-sectional variation in waiver rates driven by borrower opacity does not predict returns (the G-Z channel).

4. **Co-investment disciplines enforcement.** Funds with higher GP co-investment enforce more aggressively, controlling for management fees and covenant tightness.

5. **The trap is most severe in low-recovery segments.** The welfare cost of evergreening is $(1-D)(q_e-1/2)^2/2$, which increases in $(1-D)$. Subordinated or unsecured private credit should exhibit larger welfare losses than senior secured lending.

**What should not hold if the model is correct.** Tight covenants should not be associated with better LP outcomes controlling for $\alpha/\delta$. Amendment rates should not be uncorrelated with fee structure. GP co-investment should not be uncorrelated with enforcement behavior.

## 6.4 Scope conditions and limitations

Several scope conditions bound the result.

**Closed-end fund structure.** The model applies to closed-end private credit funds where liquidation proceeds remain as fee-earning AUM during the investment period. Open-end funds or funds that return liquidation proceeds to LPs have a weaker AUM-preservation motive, reducing the fee channel. Under committed-capital fee structures (where management fees are charged on committed rather than invested capital), the management fee does not depend on enforcement decisions, and the GP's enforcement cutoff depends only on co-investment. In the limit $\alpha_{\text{AUM}} = 0$ (pure committed-capital fees), threshold invariance still holds but the welfare loss shrinks because $q_e$ increases toward $q^*$. The committed-capital versus invested-capital fee distinction is itself a testable prediction.

**Endogenous covenant-setting.** Section 5 endogenizes $\theta$ through GP competition for LP capital. The key finding: when naive LPs reward covenant tightness, competition drives $\theta$ to 1 (maximum tightness) and $\delta$ to its minimum, producing covenant theater. The baseline model's exogenous $\theta$ is the correct reduced form for a market where institutional forces (LP naivety, market convention, regulatory metrics) push $\theta$ above $q_e$.

**Single loan.** The baseline model analyzes a single loan, abstracting from portfolio-level interactions. Appendix A shows that convex portfolio reputation costs break threshold invariance, making the enforcement cutoff increasing in $\theta$. For reasonable parameterizations, this effect is quantitatively small relative to the baseline invariance, but portfolio-level forces represent a meaningful boundary condition for the result.

**Uniform distribution.** The $U[0,1]$ distribution generates the clean comparative static $dA/d\theta = 1$ in Proposition 2. Under a general distribution $F$, the result becomes $dA/d\theta = f(\theta)$, which depends on the density at the covenant threshold. Threshold invariance of LP welfare (Proposition 3) holds for any continuous distribution, because the irrelevance of $\theta$ follows from the pointwise equivalence of continuation payoffs for amended and non-violated loans, not from the distributional assumption.

**Performance fees (carry).** Section 4 extends the model to a portfolio with carry. The key finding: carry acts as conditional co-investment that disciplines enforcement only above the hurdle rate. Below the hurdle, carry provides zero enforcement benefit, and the baseline model applies. This creates a cliff effect (Proposition 6) where enforcement deteriorates discretely when the fund crosses below the hurdle.

## 6.5 Policy implications

The model suggests that regulatory focus on covenant tightness in private credit is misguided. Three policy instruments can reduce the welfare loss from fee-driven under-enforcement:

**Co-investment mandates.** Raising $\delta$ aligns GP incentives with LP preferences. At $\delta = 2\alpha$, the enforcement cutoff reaches $q_e = 1/3$, shrinking the evergreening zone from 0.20 to 0.167 (a one-sixth reduction relative to the industry standard). At $\delta = 5\alpha$, $q_e = 5/12 \approx 0.42$, and the welfare loss falls below 0.3% of face value.

**Fee structure caps.** Reducing $\alpha$ directly shrinks the distortion wedge. A cap on invested-capital management fees (or a mandatory shift to committed-capital fees) would eliminate the AUM-preservation motive.

**Transparency requirements.** Requiring loan-level disclosure of amendment activity would collapse the information asymmetry between the GP and LPs. If LPs can observe which covenant violations were amended, they can condition capital allocation on actual enforcement behavior, disciplining the GP through the capital market rather than through contract design. Form PF filings are confidential and visible only to regulators, so they cannot serve this function directly; the relevant policy lever would be mandated LP-facing disclosure (for example, in Form ADV brochures, in audited fund financials, or through standardized loan-level reporting in LP advisory committee materials).

**Illustrative magnitude.** At baseline parameters ($\alpha = 0.02$, $\delta = 0.03$, $D = 0.2$), the per-loan welfare loss is 1.6% of face value. This figure is illustrative rather than calibrated: it applies to closed-end funds charging fees on invested capital (a subset of the market), assumes uniform borrower quality, and abstracts from carry, portfolio diversification, and GP competition. Across the invested-capital segment of private credit, the aggregate loss would be a fraction of $1.7 trillion times 1.6%, suggesting the distortion has economic significance even under conservative assumptions about market scope.

# 7 Conclusion

The GP's enforcement cutoff in private credit depends on the management-fee-to-co-investment ratio $\alpha/\delta$ and is invariant to the covenant threshold $\theta$. Tightening the covenant increases violations without increasing enforcement. Extending the baseline to portfolios with carried interest, carry acts as conditional co-investment: it disciplines enforcement above the hurdle rate but provides no benefit below it, creating a cliff effect where enforcement deteriorates precisely when the fund is struggling. Endogenizing covenant tightness through GP competition, naive LPs generate a race to maximum tightness with minimum co-investment, while sophisticated LPs ignore covenants and compete on fee structure. These results redirect attention from covenant design to fee structure design as the binding instrument for credit discipline in delegated lending. Co-investment, not covenant tightness or carry, is the robust lever for LP protection.

# References


Viral V. Acharya, Simone Lenzu, and Olivier Wang. Zombie lending and policy traps. Technical Report 29606, NBER, 2021.

Jonathan B. Berk and Richard C. Green. Mutual fund flows and performance in rational markets. *Journal of Political Economy*, 112(6):1269–1295, 2004.

Sergei A. Davydenko, Redouane Elkamhi, and Marco Salerno. When are financial covenants relevant? SSRN Working Paper, abstract_id=3554454, 2020.

Douglas W. Diamond. Financial intermediation and delegated monitoring. *Review of Economic Studies*, 51(3):393–414, 1984.

Jason R. Donaldson, Denis Gromb, and Giorgia Piacentino. Conflicting priorities: A theory of covenants and collateral. *Journal of Finance*, 80(3):1739–1768, 2025.

Miguel Faria-e Castro, Pascal Paul, and Juan M. Sanchez. Evergreening. *Journal of Financial Economics*, 153:103778, 2024.

Nicolae Gârleanu and Jeffrey Zwiebel. Design and renegotiation of debt covenants. *Review of Financial Studies*, 22(2):749–781, 2009.

Raghuram Rajan and Andrew Winton. Covenants and collateral as incentives to monitor. *Journal of Finance*, 50(4):1113–1146, 1995.

# A Portfolio Reputation Extension

The baseline model evaluates each loan in isolation. This section introduces a portfolio-level reputation cost and shows how threshold invariance breaks, providing a boundary condition for the main result.

Suppose the GP faces a convex reputation cost $R(n)$ that depends on the total mass of amendments $n$, with $R' > 0$ and $R'' > 0$. Assume a continuum of i.i.d. loans, so by the law of large numbers the realized mass equals the ex-ante probability: under $U[0, 1]$, $n(\theta, q_e) = \theta - q_e$ for $\theta > q_e$.

The GP's net benefit from amending a marginal loan at quality $q = q_e$ must now account for the reputation cost:

$$\Pi^A(q) - \Pi^E(q) = (1 - D)\left[(\alpha + \delta)q - \frac{\delta}{2}\right] - R'\big(n(\theta, q_e)\big). \tag{22}$$

Setting the expression to zero and solving for the enforcement cutoff:

$$q_e(\theta) = \frac{\delta/2 + R'(\theta - q_e(\theta))/(1 - D)}{\alpha + \delta}. \tag{23}$$

Now $q_e$ depends on $\theta$ through $n = \theta - q_e$. Taking the total derivative:

$$\frac{dq_e}{d\theta} = \frac{R''(\theta - q_e)}{(\alpha + \delta)(1 - D) + R''(\theta - q_e)} > 0 \quad \text{when } R'' > 0. \tag{24}$$

With convex portfolio reputation costs, tighter covenants raise the enforcement cutoff. The GP becomes stricter when facing more violations because the marginal reputation cost of each additional amendment increases.

**Quantitative assessment.** The sensitivity of the enforcement cutoff to covenant tightness depends on $R''$ relative to $(\alpha + \delta)(1 - D)$. For small $R''$:

$$\frac{dq_e}{d\theta} \approx \frac{R''}{(\alpha + \delta)(1 - D)}.$$

With $\alpha = 0.02$, $\delta = 0.03$, and $D = 0.2$, the denominator is $0.05 \times 0.8 = 0.04$. For $dq_e/d\theta$ to exceed 0.1 (a 10% passthrough from covenant tightness to enforcement), $R''$ must exceed 0.004 per amendment squared. The baseline model ($R = 0$) represents the limit where portfolio-level reputation forces are absent. The result that $q_e$ does not depend on $\theta$ is a meaningful economic claim: it holds when the GP evaluates each loan in isolation, which is the natural starting point for private credit where enforcement decisions are made loan-by-loan based on borrower fundamentals.

# B General Liquidation Recovery

The baseline model sets $L = (1 + D)/2$. This section verifies that threshold invariance holds for all $L$ satisfying $D < L < 1$.

Define $\gamma \equiv 1 + D - 2L$. The baseline sets $\gamma = 0$. For $L > (1 + D)/2$, we have $\gamma < 0$.

**Proposition 9** (General enforcement cutoff). *For general $L \in (D, 1)$, the GP's enforcement cutoff is*

$$q_e^{gen} = \frac{\alpha(2L - 1 - D) + \delta(L - D)}{(\alpha + \delta)(1 - D)}, \tag{25}$$

*which does not depend on $\theta$. The LP-optimal cutoff is $q^* = (L - D)/(1 - D)$, and $q_e^{gen} < q^*$ for all $L < 1$.*

*Proof.* The GP amends when

$$\alpha + (\alpha + \delta)\big[q(1 - D) + D\big] \geq (2\alpha + \delta)L.$$

Rearranging:

$$(\alpha + \delta)q(1 - D) \geq (2\alpha + \delta)L - \alpha - (\alpha + \delta)D = \alpha(2L - 1 - D) + \delta(L - D) = -\alpha\gamma + \delta(L - D).$$

Solving:

$$q \geq \frac{-\alpha\gamma + \delta(L - D)}{(\alpha + \delta)(1 - D)} = q_e^{\text{gen}}.$$

This expression does not depend on $\theta$. For $L \geq (1 + D)/2$, $\gamma \leq 0$ and the numerator is non-negative; for $L < (1 + D)/2$, $\gamma > 0$ and the fee term reduces (rather than augments) the cutoff, but the expression remains valid (and non-negative because $\delta(L - D) \geq \alpha\gamma$ whenever $\delta(L - D)/\alpha \geq 1 + D - 2L$, which holds for $\delta/\alpha$ large enough; otherwise $q_e^{\text{gen}} = 0$ and the GP enforces no violations).

For $q_e^{\text{gen}} < q^*$: the condition reduces to $-\alpha\gamma < \alpha(L - D)$, i.e., $2L - 1 - D < L - D$, i.e., $L < 1$. This always holds.

At the baseline $\gamma = 0$: $q_e^{\text{gen}} = \delta(L - D)/[(\alpha + \delta)(1 - D)] = \delta(1 - D)/[2(\alpha + \delta)(1 - D)] = \delta/[2(\alpha + \delta)] = q_e$. The baseline is nested. □

Because $q_e^{\text{gen}}$ does not depend on $\theta$, the threshold invariance result (Proposition 3) holds for all $L \in (D, 1)$.

# C Comparative Statics

The key endogenous quantities are the enforcement cutoff $q_e$, the evergreening zone $q^* - q_e$, and the welfare loss. Table 6 summarizes.

Table 6: Comparative statics of the enforcement cutoff and welfare loss.

| Parameter | Effect on $q_e$ | Effect on welfare loss | Intuition |
|---|---|---|---|
| $\alpha$ (mgmt fee) | $\downarrow$ | $\uparrow$ | Higher fees increase AUM-preservation benefi |
| $\delta$ (co-investment) | $\uparrow$ | $\downarrow$ | More skin-in-the-game disciplines GP |
| $D$ (default recovery) | 0 | $\downarrow$ | Higher recovery reduces cost per bad amendn |
| $\theta$ (covenant) | 0 | 0 | Threshold invariance |

**The $\alpha/\delta$ ratio.** The enforcement cutoff is $q_e = 1/(2 + 2\alpha/\delta)$, making $\alpha/\delta$ the sufficient statistic. Table 7 displays the relationship.

Table 7: Fee-to-co-investment ratio and enforcement quality.

| $\alpha/\delta$ | $q_e$ | Evergreening zone | Welfare loss | Interpretation |
|---|---|---|---|---|
| 0.1 | 0.455 | 0.045 | 0.0008 | High co-investment, minimal distortion |
| 0.5 | 0.333 | 0.167 | 0.0111 | Moderate distortion |
| 1.0 | 0.250 | 0.250 | 0.0250 | Substantial distortion |
| 2.0 | 0.167 | 0.333 | 0.0444 | Severe distortion |
| 5.0 | 0.083 | 0.417 | 0.0694 | Near-total forbearance |

**Note on $D$.** Under $L = (1 + D)/2$, the enforcement cutoff $q_e = \delta/[2(\alpha + \delta)]$ does not depend on $D$, because both the fee benefit and the co-investment cost of amendment scale with $(1 - D)$, which cancels. The welfare loss does depend on $D$: $(1 - D)\alpha^2/[8(\alpha + \delta)^2]$ decreases in $D$. Markets with higher default recovery have smaller welfare losses because the cost of each inappropriate amendment $(L - D = (1 - D)/2)$ is smaller.