EVLI

Co–Investments in Private Equity: Lemons or Gems?

Evidence from 8200 Deals

Kristian Metsämäki | Evli Plc | 2026

Key Takeaways

No adverse selection. Across four propensity score matching specifications with progressively richer controls, no specification produces evidence of adverse selection. Co–investments are not where GPs offload their weakest deals. If anything, the data points in the opposite direction. The findings are consistent for both MOIC and IRR.

Robust downside protection. Co–investments have a loss rate of 18.1%, versus 25.7% for non–co–investments (NCI) and 25.3% for the broader deal universe ($p < 0.001$). This finding persists across geographies, time periods, sectors, and analytical methods.

Entry discipline matters. A ranking based on entry characteristics (company size, entry valuation, and profitability) reveals that deal selection matters at least as much as co–investment status. The spread from Q1 to Q5 is nearly 1.0x for co–investments (3.04x down to 2.06x) but only about 0.5x for non–co–investments. In a co–investment, the LP takes concentrated single–deal exposure; overpaying has a more direct impact on returns.

The quality advantage has persisted. Across every vintage period, co–investments exhibit no adverse selection and lower loss rates. The gross return differential was larger in pre–2010 vintages and narrower in more recent years, but the underlying quality signal has been consistent throughout. Combined with fee savings of 200–400 basis points annually, the case for co–investment rests on durable fundamentals.

Unrealized portfolio confirms the pattern. Among unrealized deals held more than four years with current valuations, co–investments show lower loss rates (15.8% vs 22.5% NCI) and higher median multiples (1.95x vs 1.70x). While sample sizes preclude statistical significance, the directional consistency with realized findings reinforces the main conclusions.

EVLI

Why This Matters

Co–investment has moved from a niche arrangement for the largest institutional investors to a mainstream allocation strategy. An estimated 82% of PE firms now offer co–investment rights (Torys LLP, 2025), and $33.2 billion was raised across co–investment vehicles in 2024 alone (Chronograph, 2024). For allocators (from pension funds to family offices), co–investment offers direct deal exposure at reduced or zero management fees and carried interest, with potential savings of 200–400 basis points annually (Fang, Ivashina, and Lerner, 2015).

The appeal comes with a persistent question: **are GPs offering their best deals, or their worst?** The academic literature has produced conflicting answers. Fang, Ivashina, and Lerner (2015) found that co–investments underperform, attributing the result to adverse selection. Braun, Jenkinson, and Schemmerl (2020) challenged that conclusion, finding no underperformance on a gross basis. More recently, Adams Street Partners (2024) found no correlation between co–investment quality and GP performance across 270+ funds.

We bring new evidence to this debate: a proprietary dataset of over 8200 deals spanning 20+ years and $1.25 trillion in deployed capital. Critically, our three–way comparison design (which distinguishes co–investments, confirmed non–co–investments, and the broader unclassified deal universe) addresses a methodological gap that has left room for doubt in prior studies.

How We Built the Dataset

The raw data comes from Evli's proprietary fund–of–funds database. Co–investment status is recorded as a free–text field with **837 unique values**, encompassing explicit positive and negative indicators, co–investor names, and numeric amounts representing co–investment capital. We classify each deal into one of three groups:

- **Co–Investment** (**CI**)**:** Deals with clear positive co–investment indicators, including named co–investors and recorded co–investment amounts. After classification, we identified 252 deals where the co–investment field recorded GP or fund vehicle names rather than LP co–investor identities. These represent inter–fund transactions or parallel fund investments, not genuine LP co–investments, and are excluded.
- **Non–Co–Investment** (**NCI**)**:** Deals where co–investment status was actively tracked and explicitly confirmed as absent.
- **Unclassified:** Deals with missing or ambiguous co–investment information. These 6528 deals are not "missing data" in the usual sense; they predominantly come from funds where co–investment was not systematically reported. They represent the normal deal population of a fund–of–funds database.

The three–way design is the methodological core of this study. Most prior work compares co–investments to non–co–investments within a classified subset. By including the unclassified majority as a third group, we can test whether the NCI comparison group is representative. **The result validates the comparison: unclassified deals perform almost identically to NCI deals** (**median 2.25x vs 2.33x,**

p = 0.47). This means our co–investment findings are not an artifact of how the comparison group was selected; they hold against the full deal universe. Importantly, the large unclassified share (80% of the sample) does not distort the analysis; it strengthens it, because the unclassified group serves as an independent benchmark that confirms the NCI group is representative of ordinary deal flow.

Previous studies have relied on binary classification (co–investment vs. non–co–investment) with substantially smaller datasets: Fang, Ivashina, and Lerner (2015) analyzed 179 co–investments, while Braun, Jenkinson, and Schemmerl (2020) used approximately 10000 deals but without a third comparison group. Our three–way design with 8229 deals adds the unclassified majority as a validation benchmark.

EVLI

Our Data

Our analysis draws on Evli's proprietary fund-of-funds database. After classifying co-investment status and excluding inter-fund transactions, the sample breaks down as follows:

	Co-Investment (CI)	Non-Co-Investment (NCI)	Unclassified	Total
All deals	1067	634	6528	8229
Realized deals (analysis sample)	613	382	4143	5138

Table 1: Sample by co-investment classification. Analysis focuses on realized deals with valid gross MOIC.

Each deal carries gross MOIC, geography, sector, vintage year, deal size, and co-investment status. We focus on realized deals with valid gross MOIC, as their performance is final and not subject to valuation uncertainty. All return figures in this paper are gross of fees and carried interest, which isolates deal quality from fee structure. This is the right lens for evaluating whether GPs offer strong deals to co-investors.

	Co-Investment	Non-Co-Investment	Unclassified
N (realized)	613	382	4143
(% of sample)	*(11.9%)*	*(7.4%)*	*(80.6%)*
Median MOIC	2.58x	2.33x	2.25x
Mean MOIC	3.08x	2.77x	2.73x
25th percentile	1.41x	0.95x	0.99x
75th percentile	3.78x	3.60x	3.53x
Loss rate (< 1.0x)	18.1%	25.7%	25.3%
Severe loss (< 0.5x)	12.7%	20.4%	19.7%
Home run rate (≥ 3.0x)	40.6%	35.9%	33.8%

Table 2: Descriptive statistics for realized deals. All MOIC figures are gross, capped at 20x. The three-way difference is statistically significant (Kruskal-Wallis $p < 0.001$).

The data tells a clear story at first glance: co-investments returned more and lost less across every metric. The two comparison groups perform nearly identically, validating the comparison framework. The findings are consistent for both MOIC and IRR: co-investments show a median gross IRR of 26.9% versus 25.3% for non-co-investments ($p = 0.032$).

EVLI

Sample Composition

The three groups differ in composition, which matters for interpreting raw performance comparisons (and motivates our propensity score matching analysis).

Variable	CI	NCI	Unclassified
Geography			
Europe	53.3%	45.0%	43.4%
North America	45.0%	53.2%	46.2%
Other	1.7%	1.9%	10.3%
Vintage Period			
Pre–2005	13.7%	12.8%	21.0%
2005–2009	12.6%	15.9%	15.3%
2010–2014	21.5%	19.9%	15.9%
2015–2019	28.4%	31.5%	23.5%
2020+	23.9%	19.9%	24.2%

Table 3: Group composition by geography and vintage period. Compositional differences are controlled for in the propensity score matching analysis.

CI deals tilt toward Europe (53% vs 45% NCI) and recent vintages. NCI deals are more concentrated in North America. The unclassified group is broadly similar in geographic composition to CI and NCI, though with a moderately higher share of "Other" geographies (10.3% vs under 2% for CI/NCI) and earlier vintages (21.0% pre–2005 vs approximately 13% for CI/NCI), reflecting less systematic co–investment reporting in older and non–European/NA funds. These compositional differences are controlled in the propensity score matching analysis.

No Evidence of Adverse Selection

Co–investments outperform both comparison groups across every metric (see Table 1). The three–way difference is statistically significant (Kruskal–Wallis $p < 0.001$), and the two comparison groups perform nearly identically.

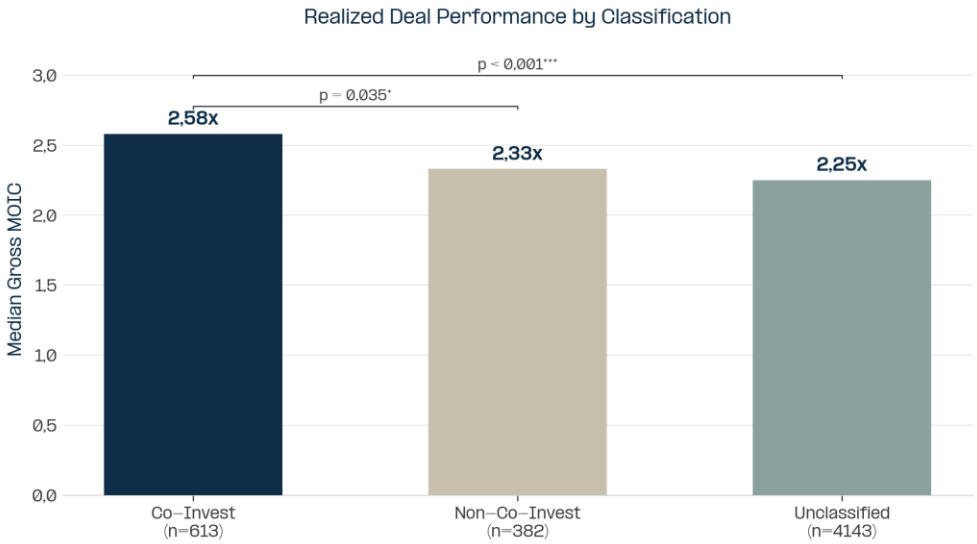

Figure 1: Return distributions for the three groups. Co–investments (dark blue) show a rightward shift relative to NCI (tan) and Unclassified (sage), with a visibly thinner left tail.

Propensity Score Matching: The Rigorous Test

The headline comparison may reflect differences in the types of deals offered as co–investments rather than co–investment status itself. Propensity score matching (PSM) controls for this by comparing co–investments to similar non–co–investments matched on observable characteristics. We run four progressively richer PSM specifications:

Specification	Controls	ATT (MOIC)	p–value
v1: Baseline	Vintage, size, geography	+0.08x	0.551
v2: + Valuation	+ entry EV/EBITDA, EBITDA margin	+0.59x	< 0.001
v3: + Deal structure	+ transaction process, seller type	+0.18x	0.256
v4: + Operations	+ add–on acquisitions	+0.59x	0.078

Table 4: Progressive propensity score matching results. ATT = average treatment effect on the treated.

The baseline specification shows the raw co–investment premium is largely explained by observable deal characteristics: co–investments tend to be concentrated in certain vintages and geographies that performed well. The enhanced specifications (v2, v4) point to a larger treatment effect, though with wider confidence intervals due to smaller matched samples and covariate balance limitations.

The key finding is not in any single specification, but across all four: no specification produces a negative treatment effect. All point estimates are positive. If GPs were systematically offloading their weakest deals, at least one specification should show underperformance. None does. This robustly rejects the adverse selection hypothesis of Fang, Ivashina, and Lerner (2015).

Within-Manager Evidence

A natural concern is that co-investments may be concentrated among higher-quality managers. Among the 9 GPs with sufficient deals in both CI and NCI categories, the majority produce higher median MOIC for their co-investment deals. The within-manager CI premium is +0.11x (p = 0.30). This is not statistically significant given the limited number of managers, but the consistently positive direction eliminates manager selection as an explanation.

Better Downside Protection

This finding has direct implications for portfolio construction, and it is robust across every analytical cut.

	Co-Investment	Non-Co-Investment	Unclassified
Loss rate (MOIC < 1.0x)	18.1%	25.7%	25.3%
Severe loss (MOIC < 0.5x)	12.7%	20.4%	19.7%
Home run rate (≥ 3.0x)	40.6%	35.9%	33.8%

Table 5: Loss and performance rates by deal classification. The loss rate difference across all three groups is highly significant (chi-squared p = 0.0005).

The loss rate difference is highly significant (chi-squared p = 0.0005 across all three groups). The pattern holds in every subgroup: Europe, North America, pre-2010, post-2010.

The loss rate advantage is not limited to older vintages. In the post-2010 era, co-investments still exhibit lower loss rates: 14.3% versus 20.8% for non-co-investments. While the smaller sample of recent realized deals limits statistical power (chi-squared p = 0.066), the directional pattern is consistent with the full-sample finding. The co-investment downside advantage persists across every sub-period we tested.

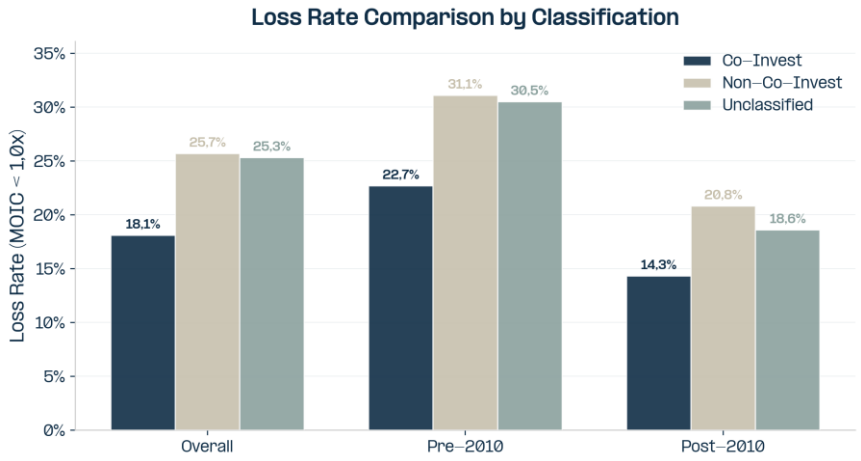

Figure 3: Loss rates across subgroups. Co–investments (dark blue) consistently exhibit lower loss rates.

Quantile regression analysis reveals where the co–investment advantage operates in the return distribution. The effect is concentrated at the **25th percentile**: co–investments earn +0.28x more than non–co–investments at Q25 ($p = 0.036$). At the median and above, the differences are smaller and not statistically significant.

This pattern tells a clear story: co–investments protect against the worst outcomes but do not generate significantly more home runs. It is consistent with GPs offering deals where they have genuine conviction: fewer speculative bets, fewer write–offs.

In practical terms: For a 50–deal co–investment program, the expected number of losses would be approximately 9, compared to roughly 13 for an equivalent portfolio of fund–only deals. Four fewer losses per portfolio is a meaningful improvement in risk–adjusted returns.

The Quality Advantage Has Persisted

Across every vintage period in our dataset, co–investments exhibit no adverse selection and lower loss rates. The quality advantage has been a consistent feature of the data, not a temporary phenomenon.

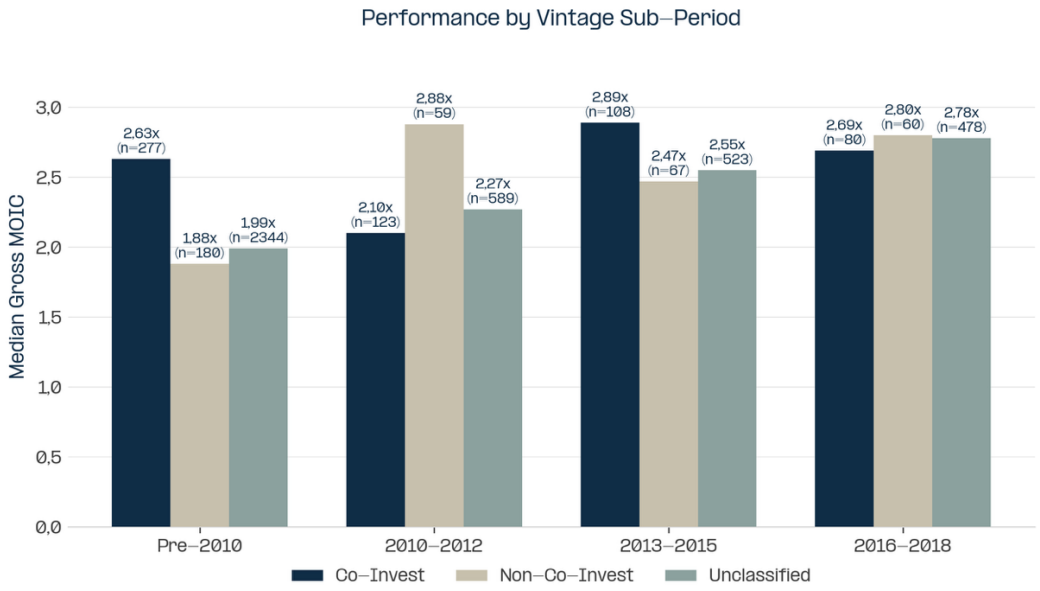

Figure 4: Median MOIC by vintage period. Pre–2010 shows a significant co–investment premium. The 2010–2012 cohort is an anomaly where NCI temporarily outperformed, driven by exceptional NCI manager concentration rather than CI deterioration. From 2013 onwards, performance converges.

The gross return differential was larger in pre–2010 vintages and narrower in more recent years. The 2010–2012 period can be described as a vintage anomaly. During this post–GFC period, a small number of NCI managers with exceptional returns drove the apparent CI underperformance. Co–investment returns themselves did not deteriorate; NCI returns were temporarily elevated. From 2013 onwards, co–investments and non–co–investments perform on par. Post–2010 PSM confirms this: ATT = +0.07x (p = 0.650).

The durable case for co–investment rests on three pillars:

1. **No adverse selection**, validated across all time periods and PSM specifications
2. **Lower loss rates** of 18.1% vs 25.7% overall, robust across subgroups
3. **Fee savings** of 200–400 bps annually (Fang, Ivashina, and Lerner, 2015), increasingly important as PE return premiums over public markets compress

On the third point: a co–investment returning 2.58x gross with zero fees and carry compares to a fund investment returning 2.33x gross that nets down further after a standard 2/20 fee structure. The combined gross performance advantage plus fee savings makes the net return differential substantial.

Unrealized Portfolio Confirmation

Our primary analysis focuses on realized deals, where performance is final. But does the pattern hold in the current, unrealized portfolio? To answer this, we examine unrealized deals with holding periods exceeding four years (mature enough that interim valuations are informative) and exclude stale valuations reported before 2023 to ensure comparability.

	Co–Investment	Non–Co–Investment	Unclassified
N	120	71	795
Median MOIC	1.95x	1.70x	1.78x
Loss rate (< 1.0x)	15.8%	22.5%	21.6%
Home run rate (≥ 3.0x)	18.3%	22.5%	20.3%

Table 6: Unrealized portfolio for deals with holding periods exceeding four years and valuations from 2023 onwards. Differences are directionally consistent with realized findings but do not reach statistical significance.

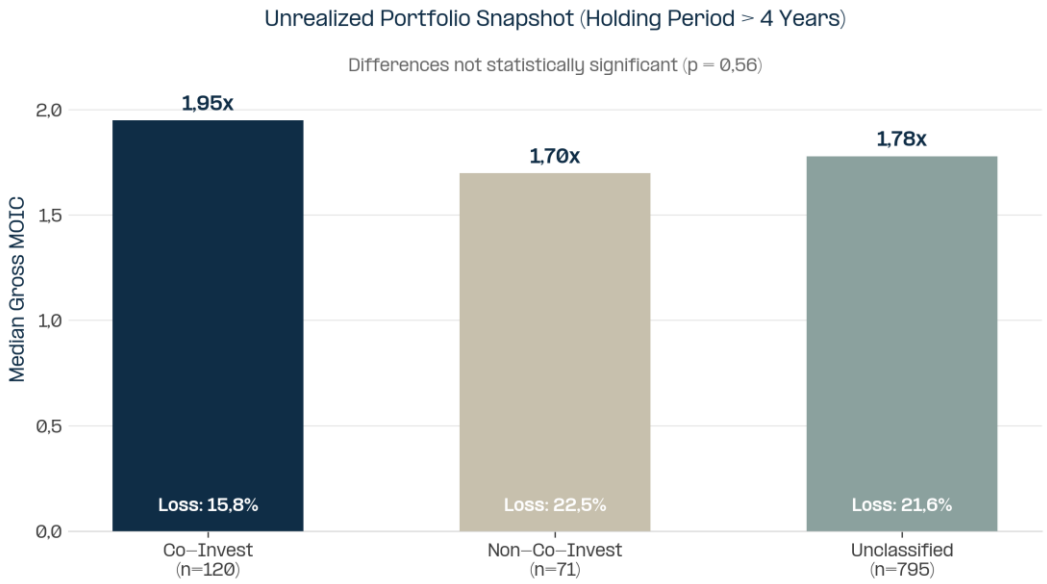

Figure 5: Unrealized deal performance (holding period > 4 years, valuations from 2023+). Co–investments show directionally favorable metrics, consistent with the realized findings.

None of the pairwise differences reach statistical significance (CI vs NCI p = 0.56; CI vs Unclassified p = 0.39; loss rate chi–squared p = 0.33), which is expected given the smaller samples and inherent valuation uncertainty. The key observation is **directional consistency**: co–investments carry lower current loss rates and marginally higher current valuations, matching the realized pattern. There is no indication that the current portfolio will reverse the findings from realized deals.

The unrealized portfolio also serves as a check against survivorship concerns. If co–investment positive selection were driven by selective realization (GPs exiting their

best CI deals first), we would expect the unrealized CI portfolio to look worse than NCI. It does not.

Quintile Analysis: Entry Discipline Matters

The previous sections answer whether co–investments are safe. This one answers what makes them work. In a fund, a weak deal is diluted across a portfolio. In a co–investment, the LP takes concentrated single–deal exposure, so entry characteristics have a more direct impact on returns. To test this, we rank each deal on a composite of three entry characteristics: company size (smaller is better), entry EV/EBITDA multiple (lower is better), and EBITDA margin (higher is better). The individual rankings are averaged and split into quintiles from Q1 (most attractive entry profile) to Q5 (least attractive).

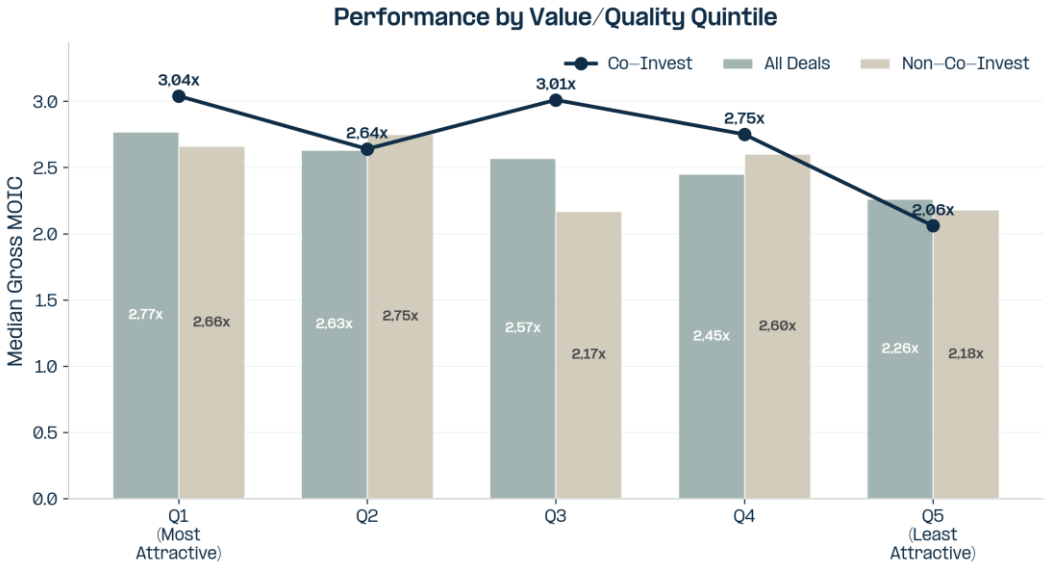

Figure 6: Median MOIC by value/quality quintile. The gradient from Q1 (most attractive) to Q5 (least attractive) is steepest for co–investments.

The quintile analysis reinforces a practical point: entry characteristics matter at least as much as co–investment status. Across all groups, smaller deals acquired at lower multiples with higher margins produce better outcomes. For co–investments specifically, the Q1–Q5 spread is nearly 1.0x (3.04x vs 2.06x), compared to 0.48x for NCI and 0.48x for unclassified. Co–investment returns are more sensitive to entry discipline because the LP takes concentrated exposure to a single deal rather than having the outcome diluted across a diversified fund portfolio, so overpaying has a more direct impact on returns.

The pattern holds in the middle of the distribution: Q3 co–investments (3.01x) outperformed Q3 non–co–investments (2.17x) by a wide margin, while Q4 co–investments (2.75x) still outpaced their non–co–investment counterparts (2.60x). Entry discipline matters throughout, not only at the extremes.

The steep quality gradient underscores the value of active deal assessment. The ability to evaluate entry pricing, business quality, and deal structure independently, and to selectively participate in co–investment opportunities rather than accepting all offers, can meaningfully influence program–level outcomes.

EVLI

Discussion and Practical Implications

The co–investment debate in academic literature has produced conflicting conclusions. Fang, Ivashina, and Lerner (2015), in the most widely cited study, found that co–investments underperformed fund investments and attributed the result to adverse selection by GPs. Braun, Jenkinson, and Schemmerl (2020) challenged that conclusion with a larger dataset and found no underperformance on a gross basis. More recently, Adams Street Partners (2024) found no correlation between co–investment quality and GP performance across 270+ funds.

Our findings are most consistent with Braun et al. (2020) and extend their work in three ways. First, the three–way comparison design validates the comparison group itself: because unclassified deals perform almost identically to confirmed non–co–investments, we can be confident that co–investment positive selection is not an artifact of how the control group was constructed. Second, the progressive PSM design shows how the estimated treatment effect evolves as more deal characteristics are controlled for, and no specification produces evidence of adverse selection. Third, the loss rate finding (18.1% vs 25.7%, $p < 0.001$) is more robust than any single MOIC comparison and persists across every subgroup we tested. This pattern is consistent with a GP conviction mechanism: GPs offer co–investment on deals where they have genuine confidence, resulting in fewer speculative bets and fewer write–offs.

Taken together, the weight of evidence from this study and the recent literature has shifted decisively against the adverse selection hypothesis. The practical question for allocators is no longer whether co–investments are safe, but how to design programs that capture their advantages.

Note: the following observations are drawn from the data; they are not investment recommendations.

1. **Co–investment deal quality is robust.** Across every analytical cut in this study, co–investments perform at least in line with fund–only deals, gross of total fees, and typically better. The concern that GPs use co–investment to offload weaker deals finds no support in our data. For allocators evaluating whether to build or expand a co–investment program, the quality–of–deal–flow question should not be the barrier. Co–investment outperformance is not concentrated in a single sector: among the four GICS sectors with sufficient data for comparison, co–investments exhibit lower loss rates in all four.
2. **Downside protection is real.** Co–investments lose money less often: roughly 18% versus 26%. For a 50–deal portfolio, that translates to approximately four fewer losses. This benefit persists across geographies and time periods, and represents portfolio construction value independent of any return premium.
3. **The quality advantage supports long–term allocation.** The co–investment quality signal (no adverse selection, lower loss rates) has been consistent across vintage periods. While the gross return differential was larger in earlier vintages, the underlying quality advantage has not deteriorated. Combined with fee savings of 200–400 bps annually, this makes co–investment a structurally attractive component of a PE allocation, not a tactical bet on a temporary anomaly.

4. **Entry discipline is the differentiator.** Active deal selection, entry pricing discipline, and deep GP relationships are the factors that separate strong co-investment outcomes from average ones. The quintile analysis shows that co-investment returns are especially sensitive to entry characteristics; independent deal assessment capability and willingness to pass on overpriced opportunities are important program design features.

While this study resolves the adverse selection debate, it opens questions about what drives the cross-section of co-investment returns. Fang, Ivashina, and Lerner (2015) raised the concern that co-investments offered later in a fund's life may reflect capital recycling pressure rather than deal quality; our data contains the fund lifecycle information needed to test this formally. Separately, preliminary patterns in our dataset suggest GPs commit more capital to co-investment deals than to comparable fund-only deals, raising the question of whether this reflects a conviction sizing mechanism. More broadly, the combination of deal-level entry characteristics, GP behavior variables, and realized outcomes in our dataset may help shed light on the variation in co-investment performance, a question with direct implications for LP program design and one that the data is well suited to explore.

EVLI

Methodology Note

Dataset. Proprietary fund–of–funds database from Evli Plc. 8229 total deals; 5138 realized deals (613 CI, 382 NCI, 4143 unclassified). Mid–1990s to 2024, $1.25 trillion in deployed capital.

Performance metric. Gross MOIC. Gross returns are the correct measure for testing adverse selection (whether GPs offer worse deals), as net returns conflate deal quality with fee structure.

Statistical methods. Mann–Whitney U tests for distribution comparison; bootstrap confidence intervals (10000 iterations); propensity score matching in four progressive specifications; quantile regression at Q10–Q90; within–manager fixed effects for GPs with both deal types; Benjamini–Hochberg FDR correction for multiple testing.

Three–way design. The core methodological contribution. By comparing CI to both NCI and unclassified deals separately, we validate the comparison group and demonstrate that findings are robust to comparison group specification.

Unrealized robustness. Unrealized deals with holding periods exceeding four years and reporting dates from 2023 onward are examined separately. The four–year threshold ensures deals are mature enough for interim MOIC to be informative; the 2023 reporting filter excludes stale valuations that would distort the comparison. Directional patterns are consistent with realized findings.

Limitations. This study relies on a single–source dataset from one LP platform, which may not be representative of the broader co–investment universe. All returns are gross; co–investments may carry hidden costs (legal fees, due diligence expenses, operational overhead) not captured in MOIC. MOIC does not account for the time value of money, and while IRR results are directionally consistent, we use MOIC as the primary metric due to broader data availability. Classification relies on pattern matching free–text fields, introducing potential misclassification. Subgroup analyses have limited statistical power in recent vintages due to smaller realized sample sizes. Finally, our dataset reflects Evli's own deal selection: the co–investments analyzed are those Evli chose to participate in, which means the results may partly reflect LP selection skill in addition to GP deal offering quality.

For the full academic paper with detailed regression tables, literature review, and additional robustness checks, please contact Evli.