
The Price of "Free"

Why Co-Investing Costs Family Offices More Than They Think



Dr. Ralph Welpel | Chairman & CEO, Mirai Capital Partners

May 2026 | Mirai Strategic Insights Series | Issue No. 5

ralph.welpel@miraicapitalpartners.com | www.miraicapitalpartners.com

In the spring of 2019, a family I have known for two decades wrote a roughly EUR 40 million cheque into a co-investment alongside a top-quartile European GP whose prior funds had all been top-decile. The deal was a software roll-up in the Nordics. The family completed diligence in three weeks. By 2024, the GP's fund had returned around 2.1x net to its LPs. The co-invest, the same asset class, the same GP, the same vintage, was marked at roughly 0.7x. The family had paid no fees and no carry, and was sitting on an unrealised loss of about EUR 12 million against a fund position that, dollar for dollar, would have more than doubled their money.

Co-investing has a serious academic case behind it. Goldman, Adams Street and BlackRock have all published in defence of co-invest programmes, and a peer-reviewed study in the *Journal of Financial Economics* by Braun, Jenkinson and Schemmerl in 2020 found no statistically significant evidence of adverse selection in the deals offered to LPs. The aggregate message from this body of work is that co-invest programmes can outperform fund commitments net of fees. And yet, in the hands of many family offices, the strategy quietly disappoints. These two facts are not contradictory, and the distance between them is one of the most expensive misunderstandings I see in this part of the market. I want to use this memo to explain why.

I am not arguing that families should stop co-investing. The structural advantages are real, and the families I have seen take it seriously have done well. I am arguing that the version of co-investing in the academic studies and the version many families actually practise are two different activities. The first works. The second usually does not. And what separates them is knowable before any cheque is written.

The reason sits in plain sight. The fee savings on a co-invest exist because the sponsor needs to syndicate the deal. GPs syndicate for several legitimate reasons: the cheque exceeds the fund's single-name concentration limit, the GP wants to lead a transaction structurally bigger than its average deal, or it wants to reward important LPs. The very feature that draws families to co-investing, no management fee, no carry, is, more often than people admit, a signal about the deal itself. Many families read the signal as a discount. It is closer to a warning.

What the Data Actually Says

The academic record is worth walking through carefully, because the disagreement inside it is real, and it matters more than the marketing on either side suggests.

The institutional case

Goldman Sachs Asset Management, in a 2023 study of vintage years 1998 to 2018, found that multi-manager co-investment funds outperformed primary, secondary, and fund-of-funds private equity strategies in aggregate, with materially lower dispersion in returns.¹ Their loss ratio, the proportion of investments returning less than capital invested, was below 5% for co-investments versus roughly 10% for buyout funds. The same Goldman data also shows the distinction between vehicle types matters: multi-manager co-investment funds returned a pooled net IRR of 15.4% for vintages 1998 to 2018, while single-GP co-investment funds, structurally closer to what a family writing direct cheques experiences, returned 9.9% over the same period.

Braun, Jenkinson and Schemmerl, writing in the *Journal of Financial Economics* in 2020, examined a large sample of more than a thousand buyout and venture capital co-investments drawn from the CapitalIQ deal universe between 1981 and 2011.² They tested directly for adverse selection. They found no statistically significant evidence of it. Co-invests performed in line with their corresponding fund deals on a gross basis, and the no-fee/no-carry economics produced a meaningful net-of-fees outperformance.

Adams Street Partners, drawing on more than 270 funds and a co-investment programme dating to 2006, reported in early 2025 that the median outperformance of their co-investment opportunities, measured across vintage years against the same GPs' broader deal universe, was approximately 10%, with no evidence of

adverse selection in the deals they were offered.³ At the individual deal level across the same dataset, performance was distributed around parity. BlackRock's institutional white paper makes the same broad case from a different sample.⁴

Each of these studies is rigorous. None is wrong. The question is whether the conditions they describe bear any resemblance to the conditions under which many family offices actually co-invest. They do not.

The contrary case

Lily Fang, Victoria Ivashina and Josh Lerner, writing in the same journal in 2015, looked at twenty years of direct private equity investments by seven large institutional investors.⁵ They found that solo direct investments outperformed market benchmarks. They found that co-investments materially underperformed the corresponding funds with which they co-invested. The authors interpreted this as evidence of adverse selection in the deals offered to those investors. They attributed the gap to two things. The deals offered as co-invests were substantially larger than the average sponsor deal. And the performance gap widened sharply when co-invests were done at the hottest moments of the cycle.

A more recent and, for family offices, more uncomfortable study comes from Lerner, Schoar, Mao and Zhang, published in the Journal of Financial Economics in 2022.¹⁰ They obtained State Street's custodial records for 108 institutional LPs and were able to see, for the first time in the academic literature, not just the main funds these LPs invested in but also the side vehicles they invested in alongside those funds: co-investments, parallel funds and separately managed accounts. The sample covered 1,467 such vehicles.

Three findings stand out. First, on average, these side vehicles underperformed the main funds they sat next to. The headline message that side vehicles are a fee-efficient way to earn a similar return to the main fund is, in their data, not supported. Second, the average hid wide variation. Some LPs did much better than the main fund in their side vehicles, and others did much worse. Third, the variation was systematic. Top-tier LPs, meaning large, well-resourced institutions with strong outside options, earned materially better returns in their side vehicles than lower-tier LPs investing in vehicles with the same GPs. The gap was widest when both groups were investing alongside the best-performing GPs.

Their interpretation is straightforward. Side vehicles, including co-investments, are not a passive, equal-access discount on fund economics. They are an active negotiation. GPs use them to reward the LPs they most want to keep, and the LPs who extract genuine outperformance from them are the ones with the scale, the team and the alternatives to decline what does not suit them. The finding is harder for family offices than it looks. The institutional return data can be replicated, but only by investors who look like the top tier of this sample: with the resources to evaluate each deal, and with the leverage to say no without consequence. A family office that takes most of what it is shown is, on the evidence of this paper, not in that group.

The two sets of findings are not contradictory. They describe two different kinds of investor.

In co-investing, the question is not what the family pays for the deal. It is which deals the family gets to see.

The first set of studies describes an institutional LP with a dedicated co-invest team, seeing high deal volume, applying a repeatable underwriting process, and building a diversified portfolio of dozens of co-invests over a vintage. The second describes an LP making a smaller number of concentrated, opportunistic co-invests, often in the largest and most exuberant deals of a cycle.

The first looks like a multi-manager co-invest fund. The second, in my experience, looks like many family offices.

The Asymmetry No One Quantifies

There is a thought experiment I have put to families more than once. Imagine the same GP, the same fund, two principals at the same close. The first commits EUR 40 million to the fund. The second commits EUR 20 million to the fund and EUR 20 million to a co-invest in one of the fund's deals. The cheque sizes here are illustrative; the asymmetry holds at any scale. Both principals have written the same total cheque. Both believe in the same manager.

The first owns a diversified slice of ten or fifteen portfolio companies. The second owns half that slice, plus a concentrated single-asset position in whichever deal the GP chose to syndicate. The fee economics look better for the second on paper. The risk economics look very different.

If the deal is a winner:

the second principal does better in absolute terms on that specific deal. Their EUR 20 million co-invest cheque buys roughly ten times the exposure to that single asset compared with what their fund commitment alone would have given them, and the co-invest slice earns no fund-level fees and no carry. On an outlier success, the additional gain on the concentrated position can more than offset the lost diversification benefit elsewhere in the fund, and the co-investor's total proceeds exceed the fund-only LP's, often meaningfully.

If the deal is a loser:

the second principal absorbs the same diluted loss across the fund's other holdings as the fund-only LP, but with an additional concentrated, undiluted loss on the co-invest cheque. The asset was syndicated for a reason. Whether that reason is benign (a concentration limit, a deal larger than the fund's average, a wish to reward important LPs) or less benign (a deal the GP wanted to share the downside on) is rarely disclosed, and in a loss outcome the distinction matters less than the size of the position itself.

The asymmetry is not in the per-deal mathematics in the simple sense. On any single deal, the co-investor's payoff is higher on winners and lower on losers. That is variance, not free return. The deeper asymmetry sits one level up, in the population of deals being syndicated. The fund-only LP is exposed to the fund's full distribution of outcomes. The co-investor is exposed to the same distribution through their fund commitment, and on top of it takes a concentrated, undiluted bet on one deal drawn from the sub-sample the GP chose to syndicate. The institutional evidence reviewed above is not unanimous on the character of that sub-sample, but it tilts in one direction: on the available LP-level data, the syndicated population is on average somewhat worse than the fund's overall deal book, with wide variation around the average. The families best placed to do well in this strategy are the ones that look least like the average family office.

Cambridge Associates' research on manager dispersion is useful here, though it speaks to a different level. Looking at average annual manager returns from January 2008 to June 2023, the spread between top- and bottom-quartile managers in global private equity is materially wider than in US equities, global ex-US equities, emerging market equities, or hedge funds.⁶ The same logic applies, with greater force, at the deal level inside any one fund. Deal-level dispersion within a fund is typically wider than manager-level dispersion across funds, because the manager number averages over the deals. A fund commitment averages across that deal-level dispersion. A single-asset co-invest does not. The cost of the concentration may or may not be paid in worse conditional expected returns on the syndicated population. It is always paid in higher variance on the single-asset bet.

Three Problems That Recur

When I look at how the family office co-invest book actually gets built, three problems recur. They are connected, and they reinforce each other. I rank them by how much damage they do.

First: selection bias on what gets syndicated.

Whether or not adverse selection exists in the average GP's average co-invest, Fang, Ivashina and Lerner identified specifically that co-invests done in the largest deals and during market peaks underperformed.⁵ These are precisely the deals families are most likely to be offered, because these are the deals that need additional capital. The deals that do not need to be syndicated rarely are. Adams Street's own data acknowledges this: 60% of the co-investments they reviewed were among the top 25% largest investments in the lead GP's fund.³ A family that takes every co-invest its top GPs offer has, by construction, accepted a non-random sample skewed toward size and toward cycle peaks.

Second: concentration without modelling.

The Goldman case for co-investing, mathematically, requires a diversified portfolio. A family doing five or ten co-invests across the life of its programme is not running the institutional strategy. It is running five or ten concentrated single-asset bets. In my experience, this book is rarely modelled the way the same family would

model any other concentrated direct holding. Scenarios for two of the five failing are uncommon. Stress tests against a cycle that looks like 2007 rather than 2014 are uncommon. The line items get sized as fund commitments even when the risk profile is direct.

Third: diligence compression.

A primary fund commitment is typically diligenced over three to six months. A co-invest, in many family offices, is diligenced over two to four weeks. The teaser arrives, the GP signals a hard close in three weeks, and the family that wants to maintain the relationship moves at the GP's pace. Bain & Company's 2026 Global Private Equity Report frames the underwriting bar starkly. As Rebecca Burack, head of Bain's Global Private Equity Practice, puts it, "12 is the new 5", meaning today's deals require 10 to 12% annual EBITDA growth to generate the same 2.5x return that 5% growth produced a decade ago, because the era in which multiple expansion did the work for sponsors is over.⁷ That is not a judgement many families can make in three weeks.

The Costs That Are Not Called Fees

The economic argument for co-investing assumes the headline economics tell the whole story. They do not. There are three frictions, well documented in the legal and consulting literature, that erode the headline savings in ways many families never quantify.

Portfolio-company fees.

Many co-investments are marketed as no-fee, no-carry. Sponsors routinely extract fees at the portfolio-company level instead. Transaction fees are typically 1 to 3% of the purchase price. Monitoring fees are commonly 1 to 3.5% of EBITDA, or a fixed amount of USD 1 to 10 million per year.⁸ Inside a primary fund commitment, ILPA Principles 3.0 set the industry standard that portfolio-company fees should be 100% offset against management fees paid by the LP, and the trend in recent fund vintages has been toward higher offset percentages, with many large funds now negotiating 100%. To the extent the family's primary fund LPA contains a full offset, the fees collected at the portfolio-company level flow back to the LP as a reduction in management fees paid.⁹ Inside a co-investment vehicle, there is no management fee to offset against, which means these fees are simply borne by the co-investor. The headline 'no fee' is true at the vehicle level. It is materially less true at the deal level.

Warehousing carry.

When a GP or its main fund temporarily holds a portfolio investment before transferring it into the co-investment vehicle, co-investors are typically asked to bear their pro rata cost of the carry the main fund accrued during the warehousing period, plus interest. The reference rate is generally the main fund's preferred return rate, a spread over the prime rate, or the cost of borrowing on the main fund's subscription line. These costs accrue daily and are uncapped.⁸ On a deal warehoused for several months, they can quietly remove a meaningful percentage of the headline 'no fee' economics before the family is even economically in the deal.

Continuation vehicles and the disappearing discount.

Continuation vehicles have grown rapidly. Torsys reports they represented 13% of private equity exits by count in 2024,⁸ and Bain's 2026 Global PE Report puts CVs at under 10% of total PE exit value, with GP-led volume growing roughly 37% annually since 2022.⁷ When an asset that families originally co-invested into is rolled into a continuation vehicle, the no-fee/no-carry economics typically do not survive. The co-investor is offered a binary choice: cash out at a price set by the sponsor, or roll into a fee-bearing and carry-charging vehicle on terms negotiated by main-fund LPs. The feature that justified the original commitment, the absence of fund-level fees, evaporates at the moment the asset most needs additional time to compound.

Two activities, one name

It is worth being explicit about how different the institutional discipline and the typical family office practice actually are. The same word covers two activities.

Institutional Co-Invest Discipline	Family Office Co-Invest Practice
Diversified portfolios across dozens of deals per vintage	Five to ten opportunistic deals across the programme's life
Dedicated underwriting team with repeatable process	Same investment team that runs primaries, secondaries, and direct deals
Months of pattern recognition across hundreds of teasers	Two to four weeks of diligence per individual deal
Performance measured deal by deal against the lead GP's fund	Performance measured by aggregate fee savings
Co-invest book sized as concentration risk and modelled	Co-invest book treated as a fee-efficient extension of the fund programme

The institutional case rests on portfolio construction, deal volume, and underwriting discipline. The family office case, in practice, rests on fee savings. These are not the same thesis. They should not be expected to produce the same results.

What the Disciplined Families Do Differently

I am not an absolutist. The families I see doing this well share three habits. They are not complicated. They are uncomfortable, which is why many families do not adopt them.

They write the programme down before the teasers arrive.

Three or four pages, signed off by whoever functions as the family's investment committee. Which GPs the family will co-invest with. Roughly how much capital per vintage. The maximum single-deal size as a percentage of risk capital. Sectors and geographies in and out of scope. Minimum diligence period. The document is not a contract. It is a forcing function. When a teaser arrives with a three-week deadline, the family is not deciding whether to co-invest in the moment. It is deciding whether this specific deal fits a programme it already designed.

They run a co-invest committee on a defined cadence.

Monthly, or at least quarterly, regardless of whether deals are pending. The committee reviews pipeline, recent declines, and how the existing book is performing against the lead GPs' main funds. Reactive decision-making, where the family only convenes when a deal lands, structurally tilts the book toward whatever the GPs are actively syndicating, which is structurally tilted toward the largest deals at the hottest moments. A defined cadence introduces the option of saying no without breaking the relationship, because the discipline is institutional rather than personal.

They benchmark against the right thing.

Deal by deal, against the same GP's main fund. Not against the family's overall private equity benchmark. Not against a public-market equivalent. Against the specific fund the deal was offered out of. Because that is the only benchmark that tests the thesis. If the co-invest book consistently underperforms the funds it sits next to, the headline fee savings are not savings. They are a transfer from the family to the GP, paid in concentration and selection risk that the family did not properly price.

There is a quieter point underneath. Co-investing has become, for many families, a relationship product. A way of signalling to a GP that the family is serious, sophisticated, and worth keeping in the deal flow. I understand the instinct. The problem is that GPs rarely lose access to good LPs because of a missed co-invest. They lose access because of slow capital, sloppy paperwork, or unreliable communication. A family that says no to two co-invests but is clean, fast, and constructive on every primary commitment is a better LP than a family that says yes to everything and brings nothing else. This matches what I hear back when I ask GPs directly which LP behaviours actually determine access.

A Final Thought

The headline economics of a co-invest book are easy to describe. Zero management fee, zero carry. The actual economics, the ones that determine whether the family is better off, depend on six things: the diligence the family did, the concentration it accepted, the size and timing of the deals it was offered, the portfolio-company fees and warehousing carry it did not negotiate away, and how each asset performed against the same GP's primary fund. The first set of numbers is given to the family. The second has to be earned.

I do not think families should stop co-investing. I think they should stop pretending that "no fees" is the same thing as "no cost." The cost shows up in concentration risk, in compressed diligence, in selection effects, in portfolio-company fees that are not offset, in warehousing carry that accrues daily, and in the slow drift of the co-invest book away from the performance of the funds it was meant to enhance.

I could be wrong about parts of this. The institutional return data is real, and a family with the right team, deal flow, and discipline can plausibly replicate a meaningful share of it. What I am more confident about is the structure of the bet. On any single co-invest the family is doubling down on a specific asset and taking the fund's full deal risk on it, in exchange for a fee saving and the chance of a larger absolute return if that asset is one of the fund's winners. The question is not whether that trade can pay off. It can. The question is whether, across the population of deals the GP actually syndicates, it pays off often enough to compensate for the concentration. That is a trade I want my families to make with their eyes open.

The mirage in the title is not the co-investment opportunity itself. The mirage is the belief that what is free is also safe.

-
1. Goldman Sachs Asset Management, "The Case for Co-Investments" (2023), <https://am.gs.com/cms-assets/gsam-app/documents/insights/en/2023/The-Case-for-Co-Investments.pdf>
 2. Braun, R., Jenkinson, T., Schemmerl, C., "Adverse Selection and the Performance of Private Equity Co-Investments," *Journal of Financial Economics* 136(1), April 2020, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2871458
 3. Adams Street Partners, "Co-Investments: Dispelling the Myth of Adverse Selection" (January 2025), https://www.adamsstreetpartners.com/wp-content/uploads/2025/01/2025-01-15-Co-Investments-Adverse-Selection_P.pdf
 4. BlackRock Private Equity Partners, "The Advantages of Co-Investments" (white paper, via BVAI), https://www.bvai.de/fileadmin/Themenschwerpunkte/Co-Investments_und_Secondaries/Leitfaeden_Whitepapers_Studien/BlackRock_PE_Advantages_of_Co-Investing_White_Paper.pdf
 5. Fang, L., Ivashina, V., Lerner, J., "The Disintermediation of Financial Markets: Direct Investing in Private Equity," *Journal of Financial Economics* 116(1), April 2015, <https://www.hbs.edu/faculty/Pages/item.aspx?num=47928>
 6. Cambridge Associates, "Better Alternative(s): Private Investments May Improve Outcomes for DC Plan Participants" (March 2024), <https://www.cambridgeassociates.com/wp-content/uploads/2024/03/2024-03-Better-Alternatives-PI-May-Improve-Outcomes-for-DC-Plan-Participants.pdf>
 7. Bain & Company, Global Private Equity Report 2026 press release, '12 is the new 5' (23 February 2026), <https://www.prnewswire.com/news-releases/private-equity-resurgence-gathers-steam-as-new-era-challenges-firms-to-enhance-value-creation-bain--company-global-pe-report-302693957.html>
 8. Torys LLP, "Co-investment trends in 2025: continuation vehicles, warehousing costs and fees," *Torys Quarterly* Q3 2025, <https://www.torys.com/our-latest-thinking/torys-quarterly/q3-2025/co-investment-trends-in-2025>
 9. Institutional Limited Partners Association, ILPA Principles 3.0 (June 2019), section on Fee Income Beyond the Management Fee, https://ilpa.org/wp-content/uploads/2019/06/ILPA-Principles-3.0_2019.pdf
 10. Lerner, J., Schoar, A., Mao, J., Zhang, N., "Investing Outside the Box: Evidence from Alternative Vehicles in Private Equity," *Journal of Financial Economics* 143(1), 2022, <https://www.sciencedirect.com/science/article/abs/pii/S0304405X21002270>

About the Author

Dr. Ralph Welpe is Chairman and CEO of Mirai Capital Partners, a London-based boutique advisory firm working with some of the world's most influential families and entrepreneurs. With over 25 years of experience across investment banking, private banking, and alternative investments in Europe, Asia, and the Middle East, he advises family principals on strategy, governance, and cross-border investment structuring.

Contact: ralph.welpe@miraicapitalpartners.com | www.miraicapitalpartners.com | London, United Kingdom

Disclaimer: This document is provided for informational and discussion purposes only. It does not constitute investment advice, legal advice, or a solicitation of any kind. The opening illustration is a composite drawn from the author's experience. All data cited is from publicly available sources as referenced.