A Manager Selection Scorecard for Private Market Allocators

A structured framework for choosing private fund managers, and for telling skill apart from luck

Authored by CK Adya
Independent Researcher, London Business School
May 2026

Abstract

In private markets the choice of manager matters more than in any other part of investing, because the dispersion between the best and worst managers is enormous and, unlike in public markets, persistent. Yet many allocators choose managers on the strength of a headline track record and a polished presentation, neither of which reliably predicts future performance. This paper sets out a structured scorecard for manager selection built around six weighted dimensions: team and organisation, track record, strategy and edge, process and risk, terms and alignment, and operations and controls. It shows how to decompose a reported track record into market beta, leverage, timing and true skill, how to test for performance persistence, and how to recognise the red flags that precede manager failure. Using a worked evaluation of four managers, it demonstrates how the scorecard converts a mass of qualitative impressions into a defensible ranking, and it provides a due diligence process and monitoring framework that family offices and institutional allocators can apply directly.

Keywords: manager selection, due diligence, private markets, alpha, persistence, family office, fund investing, scorecard

1. Introduction

In public equity markets the choice of manager is, for many investors, a question that can be sidestepped: a low-cost index fund captures the market return, and the dispersion between active managers, while real, is modest and rarely persistent. In private markets none of this holds. There is no index to buy, the gap between the best and worst managers in a given vintage can exceed twenty percentage points of annual return, and, crucially, that gap shows genuine persistence, with strong managers tending to remain strong across successive funds. The choice of manager is therefore not a refinement but the single most important decision a private market allocator makes.

This places a heavy burden on the selection process, and it is a burden that many allocators do not meet. The typical selection rests on a headline track record, a referral from a trusted source and a confident presentation by an articulate team. Each of these is informative, but none is reliable on its own. A headline track record can be the product of market beta, leverage and timing rather than skill; a referral reflects another investor's judgement, which may be no better than one's own; and presentation skill correlates poorly with investment skill. An allocator who selects on these alone is, in effect, guessing with conviction.

This paper offers a structured alternative: a weighted scorecard that forces the allocator to assess a manager across the dimensions that actually predict future performance, and a process that decomposes the track record, tests for persistence and surfaces the red flags that precede failure. The aim is not to replace judgement with a formula, since judgement remains central, but to discipline judgement so that it is applied consistently, documented defensibly and protected against the cognitive biases that make manager selection so error-prone. The framework is built for the family office and institutional allocator choosing private fund managers across private equity, credit, real estate and venture.

2. Why Manager Selection Dominates Outcomes

The case for treating manager selection as the central decision rests on two empirical features of private markets: dispersion and persistence. Dispersion refers to the spread of outcomes across managers in the same strategy and vintage. In private equity and venture in particular, this spread is vast, far wider than in public markets, so that the difference between selecting a top-quartile and a bottom-quartile manager dwarfs the difference between asset classes or vintages. An allocator who gets the asset allocation right but the manager wrong can easily underperform one who does the reverse.

Persistence refers to the tendency of performance to repeat across a manager’s successive funds. In public markets, past performance is famously a poor guide to future returns. In private markets the evidence is more encouraging: strong managers, particularly in venture and certain niches, tend to remain strong, because their advantages, in sourcing, in operational value-add and in access to the best opportunities, are durable and self-reinforcing. This persistence is what makes manager selection worthwhile: if outcomes were random, no amount of diligence would help, but because skill persists, identifying it pays.

Together, dispersion and persistence mean that the returns to good manager selection are both large and capturable. But persistence is not universal, and herein lies the difficulty: it is strong in some strategies and weak in others, and even a genuinely skilled manager can be ruined by growing too large, losing key people or drifting from its edge. The scorecard that follows is designed to identify not just managers who have performed, but managers whose performance is likely to persist, which is a harder and more valuable judgement.

3. The Scorecard at a Glance

The scorecard assesses a manager across six dimensions, each scored and then combined into a weighted composite. The dimensions and their weights reflect what the evidence and practitioner experience suggest actually predict persistent performance, rather than what is easiest to observe.

Figure 1. The Six Dimensions and Their Weights

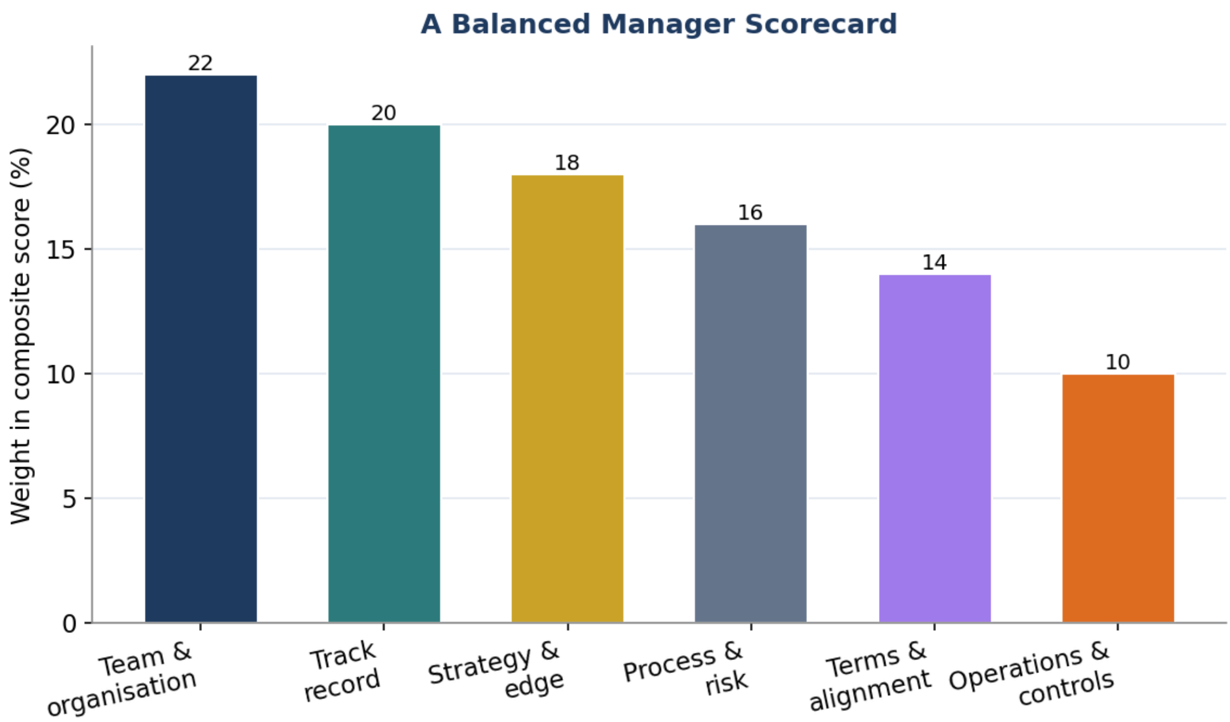

Weights reflect each dimension's contribution to predicting persistent performance.

Team and organisation carries the highest weight, because in private markets returns are produced by people, and the stability, depth and alignment of the team predict future performance better than any single number. Track record follows closely, but only once decomposed into its true skill component. Strategy and edge assesses whether the manager has a genuine, durable advantage. Process and risk examines how decisions are made and how losses are controlled. Terms and alignment captures whether the economics align the manager with the investor. Operations and controls, though weighted lowest, is a gate as much as a score: a failure here can disqualify a manager regardless of brilliance elsewhere. The sections that follow take each in turn.

4. Dimension One: Team and Organisation

The team is the engine of a private market manager, and its assessment is the most important and the most qualitative part of the scorecard. The allocator is trying to judge whether the people who produced the track record will remain, remain motivated, and remain capable of repeating their success. This requires looking beyond the founder to the depth of the organisation, the stability of the senior team, the clarity of succession and the way the economics are shared among the people who actually generate returns.

Stability and depth. A manager dependent on one or two individuals carries key-person risk that can destroy future performance if they leave. Depth, with multiple capable investors and a pipeline of talent, signals durability. The allocator should examine tenure, turnover and how the team has weathered previous departures and difficult periods.

Alignment within the team. How carry and ownership are distributed matters as much as how much the firm earns. A team in which junior partners who generate returns share meaningfully in the upside is more stable and motivated than one in which the founder retains almost everything. Lopsided economics predict departures and decline.

The team assessment is also where the allocator forms a judgement about culture and integrity that no document captures. Spending time with the team, observing how they discuss their failures as well as their successes, and probing how they handled past conflicts reveals character that the data cannot. Because team quality carries the highest weight, this qualitative work is not a soft adjunct to diligence but its core, and it should be conducted with the seriousness that its importance warrants.

5. Dimension Two: Decomposing the Track Record

A reported track record is the most cited and least understood piece of evidence in manager selection. A headline internal rate of return tells the allocator what happened, but not why, and the why is what predicts the future. The essential analytical task is to decompose the reported return into its sources, separating the part attributable to the manager's skill from the part attributable to the market, to leverage and to timing.

Figure 2. Decomposing a Reported Track Record

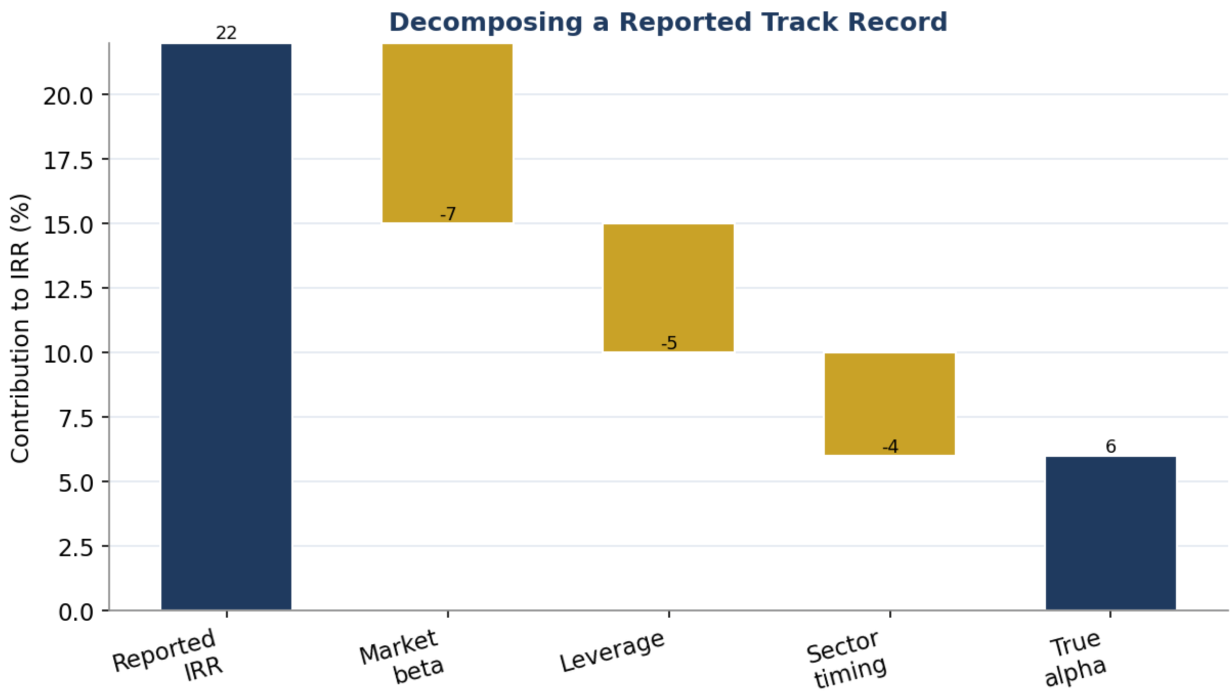

A headline return decomposed into market beta, leverage, timing and the residual true skill.

The decomposition often humbles a headline number. A manager reporting an impressive return may, on analysis, owe much of it to a rising market that lifted all assets, to leverage that amplified that market return, and to fortunate timing of entry and exit, leaving a true skill component that is far smaller than the headline suggests. Conversely, a manager with a modest headline return achieved in a difficult market, without excessive leverage, may display more genuine skill. The allocator who takes the headline at face value will systematically prefer the lucky and the leveraged over the skilled.

Doing this properly requires deal-level data, not just fund-level returns. The allocator should examine the dispersion of outcomes across the portfolio, since a return driven by one or two outsized winners is more fragile than one earned consistently across many deals. It should examine the use of leverage and subscription lines, which can flatter the reported internal rate of return without adding value. And it should benchmark against the relevant market over the same period, so that beta is stripped out. Only the residual, the return the manager added through skill, deserves weight in the scorecard.

6. Testing for Persistence

Because persistence is what makes selection worthwhile, the allocator should test for it directly rather than assume it. The question is whether a manager's strong performance reflects a repeatable process or a single fortunate fund, and the pattern across successive funds is the best available evidence.

Figure 3. Performance Persistence Across Successive Funds

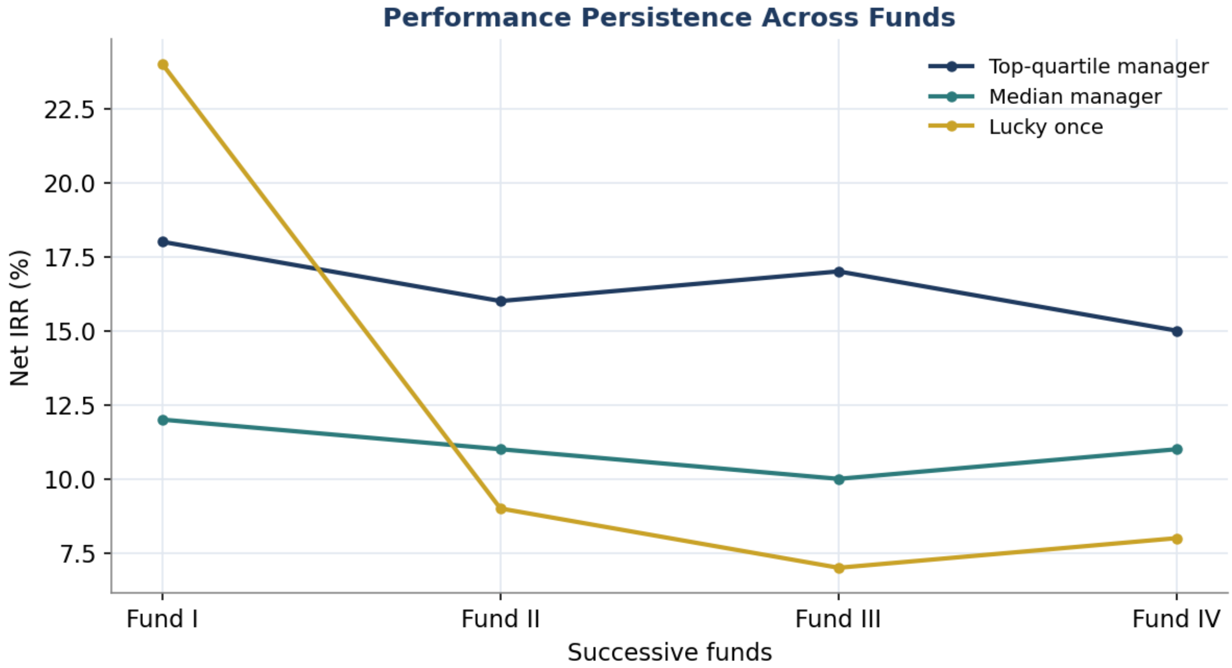

Net returns across a manager's successive funds for three patterns: genuine skill, the median, and a single lucky fund.

The chart illustrates the three patterns an allocator encounters. The genuinely skilled manager delivers strong returns consistently across funds, with the dispersion that any honest investor shows but a clear central tendency above the median. The median manager clusters near the middle fund after fund. The dangerous case is the manager whose first fund was exceptional and whose subsequent funds reverted sharply: this is the signature of luck or of a one-time opportunity that has closed, dressed up by the marketing of the single great fund. An allocator should weight consistency across funds far more heavily than a single spectacular result, and should be especially wary of a first-time fund whose track record consists of a handful of deals from a prior role.

7. Dimension Three: Strategy and Edge

A manager can have a strong team and a clean track record and still be a poor investment if its strategy lacks a durable edge. The strategy-and-edge dimension asks a deceptively simple question: why does this manager earn excess returns, and why will that continue? A satisfactory answer identifies a specific, defensible advantage and explains why competitors cannot easily erode it.

Durable edges in private markets take a few recognisable forms. Some managers have proprietary sourcing, seeing opportunities others do not through networks, reputation or specialisation. Some have genuine operational value-add, improving the businesses they own in ways a financial buyer cannot. Some have structural advantages in a niche too small or too complex for larger competitors. The common feature of a durable edge is that it is hard to replicate; an edge that consists merely of being willing to pay more, or of access to cheap leverage, is not an edge at all and will not persist.

The allocator should also watch for the erosion of a previously real edge. A strategy that worked when capital was scarce may stop working when the niche fills with imitators; an operational edge may fade as the team that provided it ages or departs; a sourcing advantage may weaken as the manager grows too large to pursue the smaller deals where it applied. The strategy-and-edge assessment is therefore forward-looking: not whether the manager had an edge, but whether it still has one and can defend it against the competition and the scale that success inevitably brings.

8. Dimension Four: Process and Risk

Where the strategy dimension asks what the manager does, the process dimension asks how it does it, and a disciplined, repeatable process is one of the strongest predictors that past performance will continue. The allocator is looking for evidence that returns are produced systematically rather than by inspiration, because a system can be repeated while inspiration cannot. This means examining how opportunities are sourced and screened, how decisions are made and challenged, and how the manager thinks about and controls risk.

Decision-making. A robust investment process has clear stages, genuine debate and a mechanism for challenge, so that decisions are tested rather than rubber-stamped. The allocator should ask to see the investment committee process, understand who can say no, and probe how the manager handles internal disagreement. A process dominated by a single unchallengeable individual is fragile, however talented that individual may be.

Risk control. Strong managers think as hard about avoiding losses as about generating gains. The allocator should understand how the manager sizes positions, manages concentration, uses leverage and monitors its portfolio for deterioration. A manager that cannot articulate how it controls downside, or whose track record shows large losses alongside its wins, carries a risk profile that the headline return conceals.

Process also reveals how a manager will behave in conditions it has not yet faced. A manager whose process is well documented and consistently applied is more likely to navigate a downturn or a strategy stress without panic or drift. The allocator cannot observe the future, but it can observe whether the manager has built the disciplined machinery that handles adversity well, and this is a meaningful part of judging whether good performance will persist into harder times.

9. Dimension Five: Terms and Alignment

Even a skilled manager is a poor investment if the terms transfer too much of the value to the manager or misalign its incentives with the investor's. The terms-and-alignment dimension assesses whether the economics are fair and whether they point the manager toward the outcomes the investor wants. This connects manager selection to the broader question of net-to-investor returns, since the most skilled manager on a punitive fee structure can deliver less to the investor than a slightly less skilled manager on fair terms.

The central alignment question is whether the manager makes its money primarily from performance or from fees. A meaningful general partner commitment, funded in cash, ensures the manager loses alongside the investor. A genuine hurdle and a whole-fund waterfall defer the manager's carry until the investor has been made whole. A management fee sized to run the business rather than to enrich the partners signals that the manager expects to earn through results. The allocator should read these terms as a statement of the manager's confidence and intent, and should be wary of a manager whose economics are structured to pay handsomely regardless of performance.

Alignment extends beyond the headline economics to the smaller terms that reveal a manager's posture toward its investors. How are expenses allocated between the fund and the manager? Are transaction and monitoring fees offset against the management fee? Is there a most-favoured-nation clause? How does the manager treat co-investment, and does it allocate the best opportunities fairly? These details, often buried in the fund documents, distinguish managers who see investors as partners from those who see them as a source of fees, and they belong in the scorecard.

10. Dimension Six: Operations and Controls

The final dimension is weighted lowest but functions as much as a gate as a score, because a failure of operations or controls can destroy an investment regardless of how brilliant the investment team is. This dimension covers the manager's back office, its valuation policy, its custody and fund administration, its compliance and its governance. Many of the most damaging losses investors have suffered in private markets came not from bad investments but from operational failures and, in the worst cases, fraud that proper operational diligence would have surfaced.

The allocator should insist on independent administration, third-party custody and an external audit, and should scrutinise the manager's valuation policy, since the freedom to mark illiquid assets is a freedom that can be abused. It should understand the manager's compliance function, its cyber and operational resilience, and its governance, including the role of any independent directors or advisory board. A dedicated operational due diligence review, conducted separately from the investment assessment and ideally by people whose job is to find problems, is the appropriate way to address this dimension.

Because operations function as a gate, the scoring here is asymmetric. A strong operational setup adds modestly to the composite, but a serious operational red flag, such as a manager marking its own illiquid book with no independent check, or a history of valuation disputes, should override a high score on every other dimension. The allocator should be willing to walk away from a brilliant investment team whose operations cannot be trusted, because the history of private markets is littered with skilled investors whose investors lost everything to an operational or governance failure.

11. The Red Flags That Precede Failure

Beyond scoring the dimensions positively, the allocator should actively hunt for the warning signs that tend to precede manager underperformance and failure. These red flags recur with enough regularity that their presence should materially lower a manager's score and, in combination, should stop a commitment entirely.

Figure 4. Red Flags Among Managers That Subsequently Underperformed

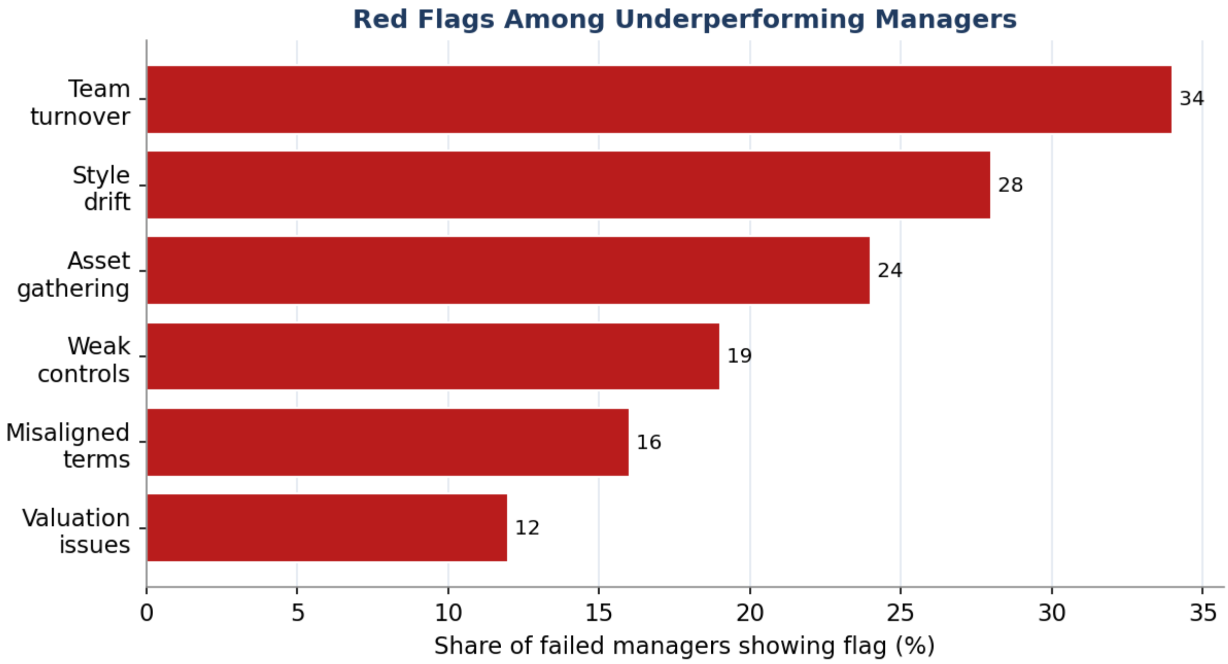

Share of underperforming managers exhibiting each warning sign in the period before decline.

Team turnover heads the list, because the departure of the people who produced the track record undermines the basis for expecting it to continue. Style drift, where a manager moves away from the strategy it is good at into adjacent areas it is not, is the second most common, and it often accompanies the third, asset gathering, where a manager raises ever-larger funds that exceed the capacity of its strategy and force it to compromise. Weak controls, misaligned terms and valuation issues complete the list. None of these is automatically disqualifying in isolation, but each should prompt deeper investigation, and their accumulation is a clear signal to decline.

The value of an explicit red-flag list is that it counteracts the optimism that infects manager selection. An allocator who has spent weeks with an impressive team, who likes the people and wants the deal to work, is psychologically primed to explain away warning signs. A written list of red flags, reviewed deliberately before the decision, forces the allocator to confront evidence it might otherwise rationalise, and so protects the process from the very enthusiasm that good managers are skilled at generating.

12. Applying the Scorecard: A Worked Evaluation

To show the scorecard in use, consider four managers under evaluation for the same mandate. Each is scored from one to five on each of the six dimensions, and the scores are combined using the weights from Figure 1 to produce a composite. The exercise demonstrates how the scorecard converts a mass of qualitative impressions into a defensible, comparable ranking.

Figure 5. Manager Scores by Dimension

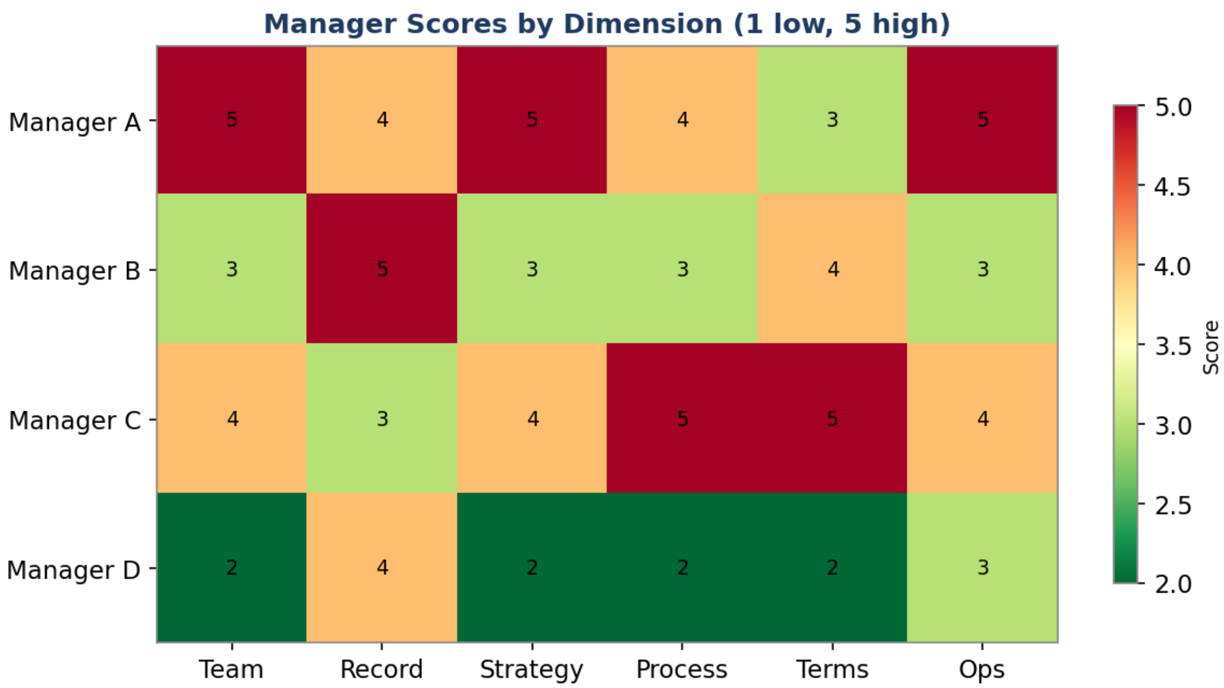

Each manager scored one to five across the six dimensions before weighting.

The raw scores already tell a story. Manager A is strong across the board, with particular strength in team and strategy. Manager B has an excellent track record but weaker team depth and process, the classic profile of a manager whose past may not persist. Manager C is balanced and strong on process and terms. Manager D is weak on the dimensions that matter most, despite a respectable track record, and should give the allocator pause. Applying the weights converts these impressions into a single comparable number.

Figure 6. Composite Manager Ranking After Weighting

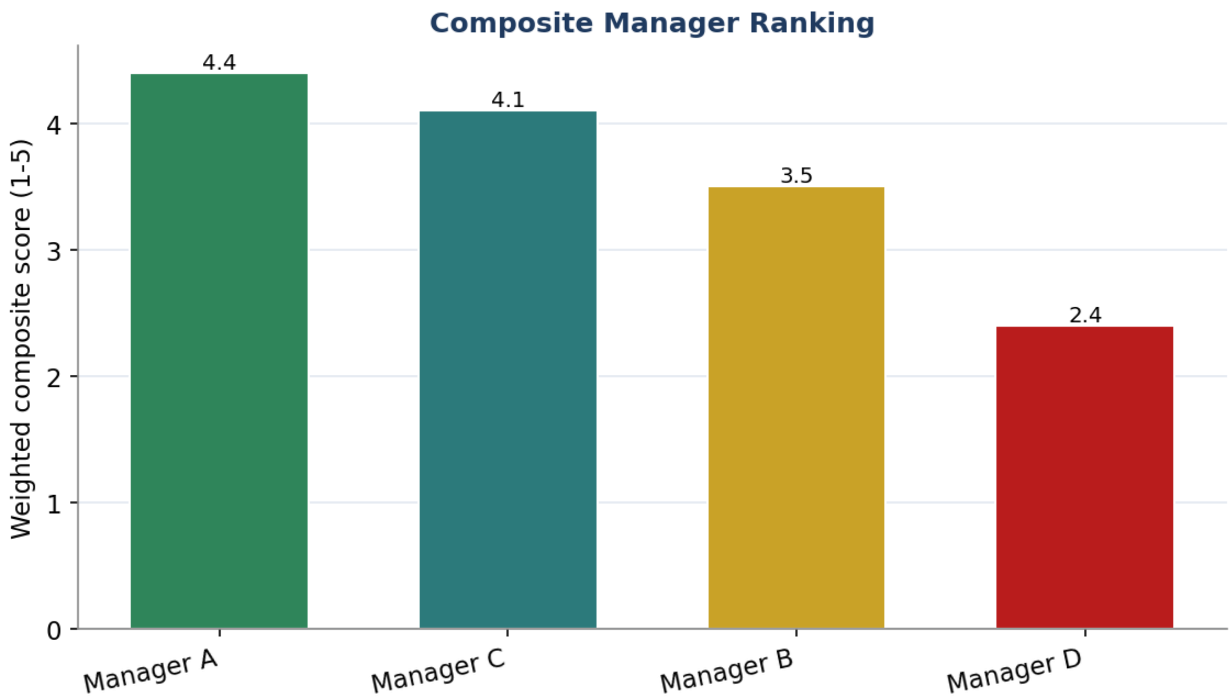

Weighted composite scores producing a defensible ranking across the four managers.

Table 1. Scorecard Summary

Manager	Top strengths	Key weaknesses	Composite	Decision
A	Team, strategy, operations	Slightly weaker terms	**4.4**	Commit
C	Process, terms, strategy	Average team depth	**4.1**	Commit / shortlist
B	Track record, terms	Team, process, strategy	**3.5**	Decline or monitor
D	Track record	Team, strategy, process	**2.4**	Decline

Composite on a one-to-five scale using the weights in Figure 1. The ranking differs from a track-record-only ranking.

The crucial result is that the scorecard ranking differs from the ranking a naive track-record comparison would produce. Manager B, with the strongest headline record, falls to third once its weak team and process are weighted in, while Manager A, whose record is good but not the best, ranks first on the strength of the dimensions that predict persistence. This is the scorecard doing its job: directing capital toward the managers most likely to perform in future, rather than those who have performed in the past.

13. The Due Diligence Process

The scorecard is the output of a disciplined process, and the quality of the output depends on the rigour of the process that feeds it. A sound manager due diligence process moves through defined stages, each designed to gather the evidence that a particular part of the scorecard requires, and it separates the people who advocate for a manager from those whose job is to find reasons to decline.

Figure 7. The Manager Due Diligence Process

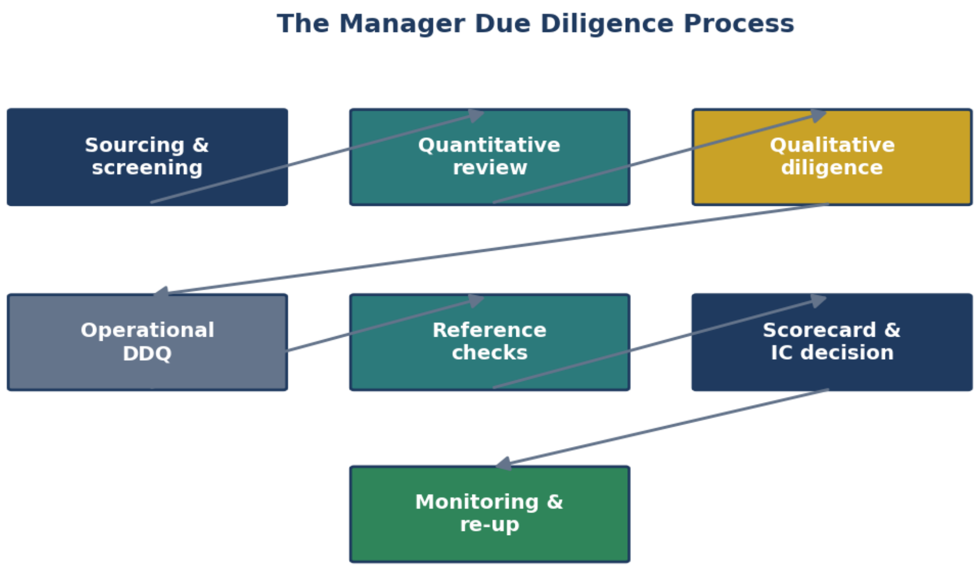

A staged process feeding the scorecard and the investment committee decision.

The process begins with sourcing and screening, which narrows a wide universe to a manageable shortlist using basic criteria. Quantitative review then decomposes the track record and tests for persistence. Qualitative diligence, the most time-intensive stage, assesses team, strategy and process through meetings, document review and observation. Operational due diligence, conducted independently, addresses the operations-and-controls gate. Reference checks, including off-list references the manager did not provide, test the picture against the views of those who have worked with or invested alongside the manager. The scorecard then synthesises all of this for the investment committee decision, and monitoring continues through the life of the commitment.

A feature of good process is that it is consistent across managers, so that they are evaluated on the same dimensions with the same rigour and the results are genuinely comparable. Allocators that diligence each manager differently, going deep where they are enthusiastic and shallow where they are not, produce scorecards that reflect their own enthusiasm rather than the managers' quality. Consistency is what makes the composite scores comparable and the ranking meaningful, and it is the discipline that most distinguishes a professional selection process from an ad hoc one.

14. Sensitivity to Weighting Choices

Because the composite depends on the weights assigned to each dimension, a thoughtful allocator should test how sensitive the ranking is to those choices. If small changes in weighting reorder the managers, the ranking is fragile and the allocator should be cautious; if the ranking is stable across reasonable weightings, it is robust.

Figure 8. Sensitivity of the Composite Score to Weighting Choices

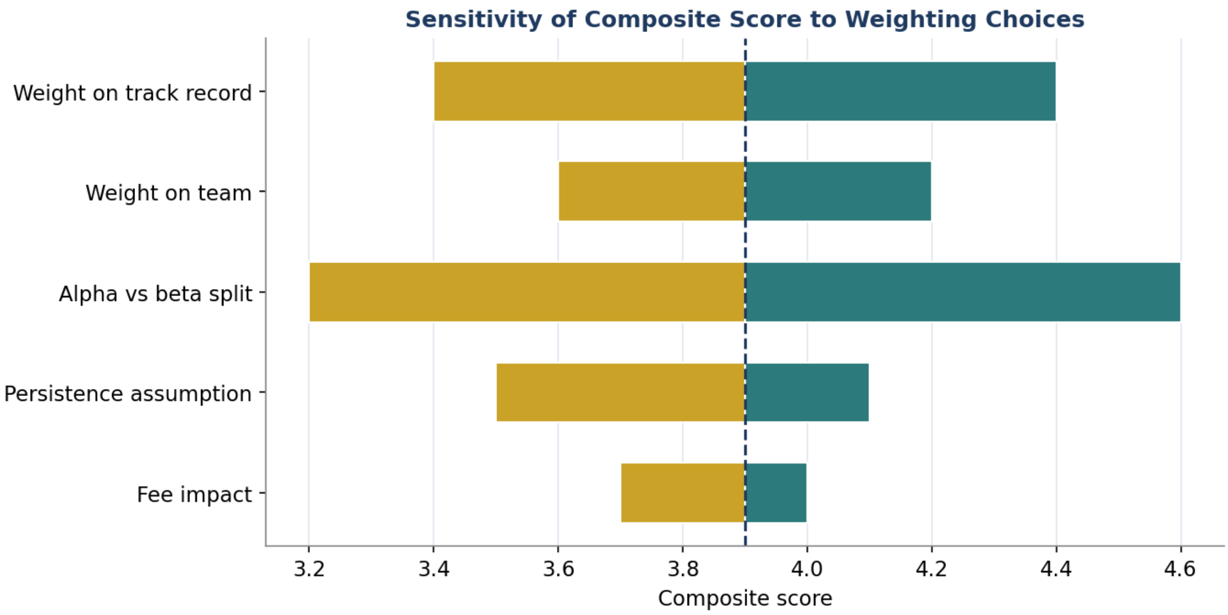

Each bar shows the swing in the composite from varying one weighting assumption around the base case.

In the worked example, the ranking proves robust: across a wide range of reasonable weightings, Manager A and Manager C remain ahead of Managers B and D, because their strength on the highly weighted team and process dimensions is not an artefact of a particular weighting choice. The sensitivity analysis also clarifies which judgements matter most. The split between alpha and beta in the track record, and the weight on team quality, move the composite most, which tells the allocator where to concentrate its analytical effort and where to be most careful about its assumptions.

15. Cognitive Biases in Manager Selection

Manager selection is unusually vulnerable to cognitive bias, because it involves charismatic people, persuasive narratives and high stakes, and because feedback on whether a decision was right arrives years later, too late to correct the process. An allocator who does not actively guard against bias will find that its careful framework is quietly subverted by the very human tendencies the framework was meant to discipline.

- **Halo effect.** being swayed by a polished presentation and an articulate founder, which correlate poorly with investment skill.
- **Narrative seduction.** accepting the manager's narrative of its track record without independent decomposition, because the story is compelling.
- **Confirmation bias.** gathering evidence that supports a manager the allocator already likes while discounting contrary signals.
- **Recency and salience.** over-weighting a single spectacular fund and treating it as proof of skill rather than possible luck.
- **Herding.** committing because respected peers have, substituting their judgement for one's own diligence.

The scorecard and the structured process are, in large part, devices for resisting these biases. By forcing the allocator to score each dimension explicitly, to decompose the track record independently, to review a written red-flag list and to test the ranking's robustness, the framework interrupts the intuitive leaps that bias exploits. It cannot eliminate judgement, nor should it, but it channels judgement through a structure that makes the biases visible and therefore resistible.

16. Monitoring After Commitment

Selection does not end at the commitment; it continues through monitoring, because the conditions that justified the commitment can change, and the decision to re-up into a manager's next fund is itself a selection decision informed by what monitoring reveals. An allocator that diligences carefully at entry but monitors loosely thereafter forfeits much of the value of its initial work and risks re-upping into a manager that has quietly deteriorated.

Effective monitoring tracks the same dimensions the scorecard assessed, watching for change. Has the team remained stable, or have key people left? Has the strategy held, or has the manager drifted or grown beyond its capacity? Is the process still disciplined, and are the terms of the next fund still aligned? Monitoring also tracks the realisation of the current fund against expectations, distinguishing temporary marks from genuine impairment. The allocator should treat each annual meeting and each re-up as an opportunity to re-score the manager, not as a formality.

The re-up decision deserves particular discipline because it is where relationship and inertia most threaten judgement. An allocator that has committed to a manager, met its team for years and values the relationship is psychologically inclined to re-up even when the evidence has turned. Applying the scorecard afresh at each re-up, with the same rigour as at first commitment, protects against this inertia and ensures that capital continues to flow to managers on the strength of their current quality rather than the comfort of an existing relationship.

17. Special Cases: First-Time Funds and Emerging Managers

First-time funds and emerging managers pose a particular challenge, because the track record that the scorecard relies on is absent or attached to a prior firm. Yet emerging managers can be among the most attractive opportunities, often hungry, aligned and pursuing strategies before they become crowded, so an allocator that excludes them entirely forfeits real returns. The framework adapts rather than breaks for these cases.

For a first-time fund, the allocator leans more heavily on the dimensions that do not require a fund track record: the team's individual track records from prior roles, carefully attributed to avoid claiming credit the individual did not earn; the clarity and durability of the strategy and edge; the quality of the process the team has built; and the alignment of the terms. Operational diligence becomes more important, not less, because a new firm's back office is unproven. Reference checks carry extra weight, since they substitute for the missing fund history. The scorecard accommodates this by allowing the team and strategy dimensions to bear more of the assessment where the track-record dimension is necessarily thin.

18. Common Errors

- **Selecting on the headline.** choosing managers on headline returns without decomposing them into skill, beta, leverage and timing.
- **Mistaking luck for skill.** over-weighting a single spectacular fund and ignoring the absence of persistence across a manager's history.
- **Skipping operations.** neglecting operational due diligence, leaving the portfolio exposed to the failures that cause the largest losses.
- **Re-upping on autopilot.** allowing relationship and inertia to drive re-ups into managers that have deteriorated.
- **Inconsistent process.** diligencing managers inconsistently, so that scores reflect the allocator's enthusiasm rather than the managers' quality.
- **Overlooking alignment.** ignoring fee and alignment terms, so that a skilled manager on punitive terms delivers little to the investor.

19. An Implementation Roadmap

1. Define the scorecard dimensions and weights in advance, tailored to the strategy, and commit to them before evaluating any manager.
2. Build a consistent due diligence process with defined stages, separating investment advocacy from operational scrutiny.
3. Decompose every track record into skill, beta, leverage and timing using deal-level data, and test for persistence across funds.
4. Conduct independent operational due diligence as a gate, with the power to override a high investment score.
5. Score each manager explicitly on every dimension, review the written red-flag list, and test the ranking's sensitivity to weights.

6. Document the decision and its rationale, so that monitoring and future re-ups can be judged against the original thesis.
7. Re-score managers at each annual review and re-up with the same rigour as at first commitment.

20. Co-Investment and the Selection Question

Co-investment, where an allocator invests directly alongside a manager in a specific deal, usually with reduced or no fees, has become a central feature of private markets and adds a dimension to the selection question. The opportunity is attractive: co-investment lowers the blended cost of a manager relationship and lets the allocator concentrate capital in deals it finds especially compelling. But it also shifts part of the selection burden from the manager to the allocator, because the allocator is now choosing individual deals, not merely a manager, and must have the capability to evaluate them.

The scorecard framework helps here in two ways. First, a manager that offers co-investment generously and fairly, allocating attractive deals to its investors rather than hoarding them, scores well on alignment and signals a partnership posture. A manager that offers co-investment only in its weakest deals, using investors as a dumping ground for risk it does not want, reveals the opposite, and the pattern of what a manager offers for co-investment is itself a diligence signal. Second, the allocator must honestly assess its own capacity to evaluate co-investments quickly, since these opportunities often come with short deadlines, and an allocator without that capacity should be cautious about a programme that assumes it.

The deeper point is that co-investment changes the risk profile of a manager relationship. A diversified fund commitment spreads risk across many deals chosen by the manager; a concentrated co-investment programme adds single-deal risk chosen partly by the allocator. An allocator should size its co-investment activity against its genuine deal-evaluation capability and its tolerance for concentration, and should treat the decision to build a co-investment programme as a strategic choice about how much of the selection burden it wishes to carry itself rather than delegate to its managers.

21. Benchmarking and Vintage Diversification

A manager cannot be judged in isolation; its performance must be compared against the right benchmark and seen in the context of its vintage. Benchmarking in private markets is more difficult than in public markets, because there is no investable index and because reported peer benchmarks suffer from survivorship and selection biases that flatter the universe. The allocator should benchmark a manager against the relevant strategy and vintage peer set, while treating the precise quartile ranking with appropriate scepticism given the data's limitations.

Vintage diversification is the portfolio-level counterpart to manager selection. Because private market returns vary considerably by the year a fund begins investing, concentrating commitments in a single vintage exposes the allocator to the luck of that year's entry environment. Committing steadily across vintages, year after year, diversifies this timing risk and is one of the most reliable disciplines in private market investing. It also imposes a useful discipline on selection, since a steady commitment pace forces the allocator to evaluate managers continuously rather than in occasional bursts, building the institutional capability that good selection requires.

The interaction between selection and vintage diversification is important. An allocator committed to steady pacing across vintages will inevitably commit in less attractive years as well as attractive ones, and the discipline of the scorecard becomes its protection: in a frothy vintage, when capital is abundant and standards slip, a rigorous scorecard prevents the allocator from committing to weak managers simply to maintain its pace. The combination of steady vintage diversification and uncompromising selection standards is what allows an allocator to be consistently invested without being indiscriminately invested.

22. A Strategic Perspective

The allocators who compound capital most successfully in private markets are not those with access to the most managers or the most capital, but those with the most disciplined selection process. Because dispersion is vast and skill persists, the returns to getting selection right are larger in private markets than anywhere else in investing, and they are available to any allocator willing to do the work. The scorecard is not a guarantee of good outcomes, since even the best process will sometimes back a manager that disappoints, but it tilts the odds decisively and, just as importantly, it produces decisions that can be defended, learned from and improved over time.

For family offices in particular, a structured scorecard offers a further benefit: it institutionalises a capability that might otherwise reside in a single individual's instinct. By making the dimensions, weights and process explicit, the family builds a selection discipline that survives the departure of any one person and that the next generation can learn and apply. Manager selection thus becomes a durable institutional asset of the family, not merely a series of individual bets, and that durability is itself a source of long-term advantage.

23. Conclusion

In private markets the manager is the investment, and the choice of manager dominates outcomes to a degree unmatched elsewhere. Yet the typical selection rests on the least reliable evidence: a headline track record, a referral and a polished presentation. This paper has set out a structured alternative, a weighted scorecard across six dimensions, supported by a decomposition of the track record, a test for persistence, an explicit red-flag list and a disciplined process that guards against the biases manager selection invites. The worked evaluation showed the scorecard reordering managers away from a naive track-record ranking and toward the managers most likely to perform in future.

The framework will not make selection easy, because judgement remains central and the future remains uncertain. But it will make selection better: more consistent, more defensible, more resistant to bias, and more likely to direct capital to genuine and persistent skill. For the family office and institutional allocator, that improvement is among the highest-return uses of effort available, because in private markets, more than anywhere else, choosing well is the whole of the game. The discipline to choose well, applied consistently over time, is what separates the allocators who capture the promise of private markets from those who merely pay its fees.

24. Limitations

This paper presents a framework and uses modelled figures to illustrate its application rather than to describe any specific manager or fund. The weights and scores are indicative and should be adapted to the strategy, the allocator's objectives and current evidence. No framework can guarantee that a selected manager will perform, since outcomes depend on future conditions that cannot be known. Nothing here constitutes investment advice. Allocators should conduct their own diligence and obtain advice appropriate to their circumstances before committing capital.

Appendix A. The Scorecard Dimensions and Weights

Table 2. Scorecard Reference

Dimension	Weight	What it assesses
Team and organisation	22%	Stability, depth, succession, internal alignment
Track record	20%	Skill after stripping out beta, leverage and timing
Strategy and edge	18%	Durable, defensible source of excess return
Process and risk	16%	Repeatable decision-making and loss control
Terms and alignment	14%	Fair economics aligning manager with investor
Operations and controls	10%	Administration, valuation, custody, governance (gate)

Weights are a starting point and should be tailored to the strategy and the allocator's priorities.

Appendix B. A Manager Diligence Checklist

- Has the track record been decomposed into skill, beta, leverage and timing using deal-level data?
- Is performance persistent across funds, or driven by a single fortunate fund?
- Is the team deep, stable and aligned, with manageable key-person risk?
- Does the manager have a durable, defensible edge that survives its growth?
- Has independent operational due diligence cleared the operations-and-controls gate?
- Are the terms fair and aligned, with a meaningful GP commitment and a genuine hurdle?
- Have off-list references been taken, and has the written red-flag list been reviewed?

Appendix C. Glossary

Alpha. The portion of return attributable to manager skill, after stripping out market and other systematic factors.

Beta. The portion of return attributable to the market's movement rather than to skill.

Dispersion. The spread of outcomes across managers in the same strategy and vintage.

Operational due diligence. Independent review of a manager's back office, valuation, custody, compliance and governance.

Persistence. The tendency of a manager's relative performance to repeat across successive funds.

Re-up. A commitment to a manager's successor fund, itself a selection decision.

Vintage. The year a fund begins investing, used to compare managers on a like-for-like basis.

References

[1] Kaplan, S., and Schoar, A. (2005). Private Equity Performance: Returns, Persistence, and Capital Flows. Journal of Finance, 60(4), 1791-1823.

[2] Harris, R., Jenkinson, T., Kaplan, S., and Stucke, R. (2023). Has Persistence Persisted in Private Equity? Journal of Corporate Finance.

[3] Korteweg, A., and Sorensen, M. (2017). Skill and Luck in Private Equity Performance. Journal of Financial Economics.

[4] Institutional Limited Partners Association (2023). Due Diligence Questionnaire and Operational Due Diligence Guidance.

[5] Cambridge Associates (2024). Manager Selection in Private Investments: Process and Evidence.

[6] Lerner, J., Schoar, A., and Wongsunwai, W. (2007). Smart Institutions, Foolish Choices: The Limited Partner Performance Puzzle. Journal of Finance.

[7] Preqin (2024). Private Capital Manager Selection and Performance Persistence.

[8] CFA Institute (2023). Manager Selection and Operational Due Diligence Standards.

About the Author

Chennakeshav Adya is an independent researcher and transaction advisor with more than two decades of experience across strategy, corporate finance and private markets, spanning the United Arab Emirates, the United Kingdom, India and beyond. He advises institutional and family-office investors on manager selection, alternatives and capital structuring, and is a guest speaker at London Business School. His work focuses on translating practitioner experience into rigorous, accessible frameworks for allocators in the Gulf and the United Kingdom.