The First-Meeting Gate: Mandate and Manager Fit for Institutional LPs

Which verified manager attributes move a limited partner beyond initial screening

Authored by Chennakeshav Adya

Independent Researcher

October 2026

Abstract

A manager that wins a first meeting with an institutional limited partner has not yet reached the decision that shapes the fundraise: whether the investor opens formal diligence or closes the file. This paper asks which verified manager attributes determine that advancement. It proposes a two-part decision rule. Written minimums on strategy, fund size, regulated status and attributable record act as conjunctive gates; a compensatory score on performance, team, alignment and terms then applies, with investor-specific noise. Published pension and sovereign-investor policies, regulator guidance and the academic literature support the categories. None supplies a threshold or weight for a Gulf investor.

An authored experiment constructs 6,000 managers and 30,000 manager and investor screening pairs across four investor types, using stated inputs and a fixed seed. Thresholds are set at a base case and then frozen. The experiment ranks repairs by the share of rejected pairs each would convert, compares an oracle ceiling and a random floor with three analyst models, runs four stress cases, builds an approached cohort to show selection bias, and replicates random study cohorts at five sizes to show sampling spread. It reports no observed result.

Under the stated inputs, 11.9% of pairs advance, and the overall rate is almost the same for first-time (11.8%) and established (11.1%) managers, because opposite experiences cancel: first-time managers advance with 28.3% of family offices and 0.2% of public pensions. Strategy fit is the first failed gate for 56.0% of rejections and fund size for 21.4%, while 13.8% of rejections pass every gate. Re-targeting converts 12.3% of rejected pairs alone; the full ladder converts 86.8%. Manager attributes without investor-specific fit reach an AUC of 0.657, against 0.960 for the gate-aware model.

The findings support a gate-first reading of the first-meeting decision and a preparation order that tests investor-specific fit before manager documentation. They do not show that any gate, weight or threshold exists for a real Gulf investor, because the outcome is generated from the same structure the models use. Selection into the approached cohort overstates the size gap by 11.2 points before weighting and 7.3 points after it. A real study would need consented, timestamped screening records and a random audit sample of rejected pairs.

JEL Classification: G23, G24, D81, L14

Keywords: fund placement, limited partners, private equity, investor screening, mandate fit, due diligence, emerging managers, selection bias, verification, institutional investors

1. The gate that matters is the one applied before formal diligence opens

A fund manager that obtains a first meeting with an institutional limited partner has not yet reached the decision that determines the fundraise. After the meeting the investor decides whether to open formal diligence, which commits analyst time, legal review and a place in an investment committee calendar, or to close the file. This paper studies that decision for a manager marketing a private fund. The question is which verified manager attributes determine advancement, and how much of the answer depends on the investor rather than on the manager.

The decision has two parts that behave differently. Investor policies and regulation set written minimums on size, strategy, record and status, and a manager below any one of them is not scored at all. Among managers that pass, a compensatory judgement on performance, team, alignment and terms decides, with noise. We propose treating the first part as conjunctive gates and the second as a score, and we test what that structure implies for outreach and for diligence preparation by constructing a population of managers and investor rules.

The sources support the structure at the level of categories. Two US public-pension policies describe fund-size, concentration or track-record criteria [3, 4]. The voluntary Santiago Principles discuss manager selection and the attribution of prior performance [5]. A survey of investors reports that diligence effort and selection criteria vary with investor size [8]. None of these sources supplies a gate level or a weight for a Gulf investor, and the experiment does not either. Its results describe how the proposed rule behaves under stated inputs.

2. Count advancement as a formal diligence stage opened within ninety days

We define the outcome as the investor opening a formal diligence stage after the first meeting: issuing a diligence questionnaire, requesting data-room access or assigning a named reviewer. The stage is counted when the first of those actions occurs on or before day ninety after the first meeting. A second meeting alone does not count, because a courtesy meeting can follow any first meeting without a change in the investor's file. The ninety-day window is an assumption of this paper, not a convention found in the sources.

The unit of analysis is the manager and investor pair. One manager is screened by several investors, and the same manager can advance with one and be rejected by another for reasons that have nothing to do with quality. Each constructed manager in the experiment is screened by five investors drawn from four types. The outcome is binary and counted once per pair. In a field study, the denominator would be pairs that held a first meeting and have an ascertainable outcome by the cutoff, including pairs with an earlier recorded qualifying action. Report silence at the cutoff separately from explicit rejection; both count as non-advancement if no qualifying action occurred. Report incomplete follow-up and pairs without a meeting separately. The simulation has no time dimension.

The outcome has to be fixed before the instrument because a claim that an attribute determines advancement is a claim about a counterfactual: whether the same pair would have advanced had that attribute alone been different. The experiment computes this counterfactual directly by repairing one attribute and holding the investor's noise fixed. A real dataset cannot observe it, which is why a later section separates what the model can know from what a cohort of screenings can estimate. A real study would record the investor's stated reason for a rejection and the pair's documentation state, which the model has as exact variables.

3. Published investor policies state size, concentration and record gates in numbers

The Los Angeles Fire and Police Pension System publishes an emerging managers policy, revised through December 2023. For private equity, its specialised-manager programme covers first-, second- or third-time institutional funds that are new to the programme with a stated target size of approximately USD 500 million or less, and follow-on commitments to second- or third-time funds of managers already in the programme with a target size of USD 1 billion or less. Managers have an opportunity to graduate to the core private-equity portfolio at their fourth fund. The Board may approve a commitment to an earlier fund for that portfolio. For real estate, selected criteria include less than USD 3 billion in assets under management and at least USD 100 million currently managed in the strategy. If there are no current assets under management in that strategy, a realised record of implementing it to the same degree is acceptable. A firm with a record shorter than five years may use prior-firm performance clearly attributable to its key individuals or the team involved. The allocation is capped at the lower of 35% of the firm's total assets under management after allocation and 20% of the capital raised for a specific commingled vehicle. No client may represent more than 50% of the firm's total assets under management [3].

CalPERS provides a documented example of concentration limits. Its posted Total Fund Investment Policy, dated effective 10 June 2024, states that a commitment cannot exceed 35% of any one fund. Aggregate net commitments to a single general partner are limited to 15% of the private equity programme's net committed capital, with exceptions to the general partner limit requiring specific committee approval [4]. They describe one investor's policy and do not supply the thresholds in this paper's simulation. The November 2021 ILPA questionnaire is a voluntary template [2].

Table 1 summarises the documents. These documents establish that two public investors express gates numerically, and that a size gate is a derived constraint: with a ticket T and a concentration cap c, the fund must be at least T divided by c. They do not establish how many Gulf investors publish such gates, whether a gate is applied at the first screen or later, or how often an exception is granted. The experiment therefore draws tickets and caps from stated ranges. These sources justify the gate inputs as categories. None supplies a population distribution.

Table 1. Written investor criteria that the first-meeting decision can apply

Source	Criterion stated	Status and limit
Emerging managers policy, private equity (LAFPP)	First- to third-time funds new to the programme with target size about USD 500 million or less; second- or third-time follow-on funds up to USD 1 billion	Board policy of one US public pension; revised to December 2023
Emerging managers policy, real estate (LAFPP)	Selected criteria: AUM below USD 3 billion; at least USD 100 million currently managed in the strategy or, if none, a realised record of equivalent implementation; allocation capped at 35% of post-allocation firm AUM and 20% of capital raised for a specific commingled vehicle; no client above 50% of firm AUM	Same policy; one investor
Total Fund Investment Policy (CalPERS)	One fund: at most 35% of the fund. One GP: at most 15% of programme net committed capital, subject to committee-approved exception.	One US public pension; policy version effective 10 June 2024
Santiago Principles, GAPP 18.2	External managers selected and monitored under written agreements, including an undertaking on the team to whom performance may be attributed	Voluntary; October 2008
ILPA Principles 3.0	Substantial GP equity contributed in cash; disclosure of management company ownership; key person notification	Voluntary best practice; 2019
ILPA DDQ 2.0	Twenty topics, with regulator and registration fields	Voluntary template; not required; November 2021

MatchPoint summary of the documents named in the first column, as read in this study; depth of reading is recorded in the source record. The table describes what each document states and makes no statement about market prevalence. [1, 2, 3, 4, 5]

4. Sovereign principles, regulators and performance rules fix the evidence a manager must show

The Santiago Principles are a voluntary set of practices for sovereign wealth funds, published in October 2008. The commentary to GAPP 18.2 says that funds using external managers carefully select reputable and creditworthy managers and enter written investment management agreements, which where appropriate include an undertaking on any promised methodology or a team to whom historical performance may have been attributed [5]. The text states practice and sets no threshold. It makes the attribution of a record, and not only its length, an item of evidence.

The Dubai Financial Services Authority describes a funds regime with Public, Exempt and Qualified Investor Fund categories, and states that DFSA-licensed firms may market and sell a wide range of foreign funds in or from the DIFC [6]. The rulebook adds a disclosure duty: a DFSA-authorised firm that offers a unit of a foreign fund must make available a current prospectus that names the fund's home regulator and describes the regulatory status accorded to the fund [14]. These texts establish that regulated status is a documented item when a fund is offered in that centre. They do not establish what an investor elsewhere requires. Regulated and administered status is therefore treated as a gate whose content must be confirmed for each investor and each marketing route.

US staff answers on the marketing rule address how a private fund records its performance. A February 2024 answer states that net performance must be calculated over the same period and with the same type of return and methodology as gross performance, and that excluding subscription facilities from gross internal rate of return while including them in net would violate the rule. A March 2025 answer states that the staff would not recommend enforcement action over gross-only extracted performance where four stated conditions are met. The page states that

these are staff views without legal force [7]. They apply to registered advisers in the United States. Their relevance to a Gulf manager depends on its registration and audience.

These sources establish that the form of a track record is itself checkable at the first pack: its basis, its attribution and its consistency between gross and net. They do not establish that any investor performs the check. They justify four gates as categories: strategy within the investor's mandate, a fund size inside the investor's window, regulated and administered status, and an attributable record. None supplies a threshold for a Gulf investor. Where an investor's own written criteria are not available, the gate should be recorded as unknown and not inferred from the practice of another investor.

5. Investor surveys show that diligence depth and the weight on size and track record vary with investor size

Da Rin and Phalippou surveyed investors listed in a 2008 directory, contacting 1,723 and obtaining 249 responses from 30 countries, a response rate of 14.4% [8]. Their central result is that the size of the private equity portfolio is the only consistently significant dimension along which diligence practices differ. Large investors spent more than twice as much time evaluating a fund proposal as small investors, and the difference was greatest for first-time funds, where uncertainty is highest. The sample is limited to investors listed in the directory that chose to respond, and the authors compare it with the directory on four characteristics.

On the track record, 77% of respondents report always benchmarking a manager's record, 52% always compute their own measure of past performance and 24% re-evaluate net asset values [8]. Past performance, measured as internal rate of return or multiple, received the largest score among selection criteria. Smaller investors gave more weight to commitments by other investors and to the opinion of an adviser or gatekeeper, and larger investors gave more weight to fund size and to the valuation of unrealised investments.

The survey is self-reported and was collected in 2009, and the authors treat the relationships as conditional correlations that they cannot prove causal. It reports no advancement outcome. Lerner, Schoar and Wong document large differences in performance across investor types, and find that investors with higher returns invest in older funds and have fewer managers in their own geography [9]. In the working-paper version, Hochberg and Rauh find that public pensions overweight home-state funds by 9.7% of the private equity portfolio on average, and that these own-state investments underperform their out-of-state investments by 3 to 4 points of net IRR a year [12]. The consequence for the design is four investor types with separate rules and no common weight.

6. The evidence on manager signals is weaker than the weight investors place on them

Kaplan and Schoar found strong persistence across funds of one partnership in a Venture Economics sample, and found that better performing funds were more likely to raise follow-on funds and larger funds [11]. Their sample suggested that first-time funds performed less well than later funds. Harris, Jenkinson, Kaplan and Stucke revisit the question with cash-flow data from investors' records through June 2019 [10]. Using final performance they confirm persistence. Using only what an investor could know at fundraising, they find little or no persistence for buyouts, overall and after 2000, while venture capital persistence holds.

Sorted by the prior fund's quartile at the time of fundraising, current buyout performance is statistically indistinguishable across quartiles, with investment multiples between 1.73 and 1.78, and first-time funds perform at least as well as any quartile [10]. Only when the prior fund is ranked on its final performance does the average public market equivalent separate, at 1.30 after a top-quartile fund against 1.01 after a bottom-quartile fund, with 1.26 for first-time funds [10]. Barber and Yasuda report in a working paper that interim performance has large effects on fundraising, most for low-reputation managers and when backed by exits, and that managers time fundraising to peaks [13]. An exit is verifiable and a valuation is an estimate.

These sources justify performance as a score component and verification of its basis as an item an investor can test. They do not show that a first-meeting decision rule that rewards performance selects better managers. The constructed investor rule in this paper weights performance by assumption, and the weight describes investor behaviour and not manager quality. The experiment's outcome is an investor decision. It makes no statement about fund returns. A finding that investors reward performance more than it predicts would be a statement about investor behaviour and would need its own evidence, which this paper does not supply.

7. Two mandate-fit conditions and six manager attributes divide into four gates and a score

The rule uses eight attributes. Two describe fit between the fund and the investor: whether the fund's strategy lies within the investor's written mandate, and whether the fund's target size lies inside the window set by the investor's ticket and concentration cap. These depend on the pair. Six describe the manager: regulated and administered status, attributable record, verified performance, team continuity, GP commitment and fund terms. Table 2 sets out the evidence an investor can verify for each and the role the attribute plays in the rule.

An attribute is verified when documentary evidence reaches the investor in the first pack and the investor can check it. A manager's true level and its documented level differ: a manager can have a stable team and not document it. The rule credits an undocumented gate item as a pass with a probability set by the investor's strictness, and enters an undocumented score item at the neutral value. The treatment is an assumption about investor behaviour. Sources [1, 2, 5, 7, 13, 14] support the evidence column and none supports the strictness values.

Table 2. Manager evidence records and their role in the first-meeting rule

Attribute	Evidence the investor can verify	Role in the rule	Basis
Strategy within investor mandate	Fund strategy against the investor's written mandate	Gate, pair specific	[3]
Fund size inside investor window	Target and hard cap against ticket and concentration cap	Gate, pair specific	[3, 4]
Regulated and administered status	Regulator and registration number; administrator and auditor	Gate for three investor types	[2, 6, 14]
Attributable record	Deal list with named team and prior-employer attribution	Gate for three investor types	[3, 5]
Verified performance	Net and gross on one basis; exits separated from valuations	Score	[7, 8, 13]
Team continuity	Biographies, tenure, succession and key person terms	Score	[1, 2]
GP commitment	Cash commitment as a share of the fund	Score	[1]
Fund terms	Fee, carry and governance terms against the ILPA Principles	Score	[1, 2]

MatchPoint design. Evidence examples are drawn from the documents cited in the last column and describe items an investor may request; they are not requirements for any named investor.

8. State the decision rule so that it can be shown to be wrong

The gate indicator for manager i and investor j is the product of four pass indicators. Each pass indicator takes the manager's true level when the item is documented, and a credited pass drawn from the investor's strictness when it is not. Equation 1 states this. The gates are strategy, size, status and record. The first two are specific to the pair and the last two depend on the manager's evidence. Writing the gate as a product makes the conjunction explicit: a single zero closes the file, and no strength elsewhere can offset it.

$$A_{ij} = \prod_{k \in G} a_{ijk} \quad (1a)$$

$$a_{ijk} = v_{ik} 1\{x_{ik} \geq c_{jk}\} + (1 - v_{ik}) \xi_{ijk}, Pr(\xi_{ijk} = 1) = 1 - h_j \quad (1b)$$

In Equation 1, A_ij is the gate indicator, G the four gates, a_ijk the pass indicator for gate k, x_ik the manager's true level, c_jk the investor's minimum, v_ik an indicator that the item is documented, xi_ijk a Bernoulli draw that credits an undocumented item, and h_j the investor's strictness, the probability that an undocumented item is treated as a failure. Strategy and size are always documented because the fund's terms state them. A pair with an undocumented item is therefore random with respect to the item's truth, and that randomness is the source of both lenient credit and strict rejection.

$$S_{ij} = \sum_{k \in K} w_k z_{ijk} + \varepsilon_{ij} \quad (2a)$$

$$Y_{ij} = A_{ij} 1\{S_{ij} \geq \tau_j\} \quad (2b)$$

In Equation 2, K is the five score items (performance, team, GP commitment, terms and depth of strategy fit), w_k their weights of 0.25, 0.20, 0.20, 0.20 and 0.15, z_ijk the standardised item with undocumented items entering at zero, and epsilon_ij investor-specific noise with a standard deviation that depends on the investor type. The threshold tau_j is set once at the base case so that a stated share of gate passers advances, and is then frozen. Y_ij is the outcome. A finding about the rule is falsifiable by a record of pairs that advanced despite failing a verified gate.

Three features of the rule are deliberate. First, the gates are conjunctive: a pair that fails any one gate cannot advance whatever its score, because the model treats such limits as first-meeting gates. The cited investor policies describe eligibility and investment limits. They do not establish that failing a limit during the first meeting always overrides the score [3, 4]. Second, documentation enters only through the verification indicator, so the same manager can pass with one investor and fail with another according to what each is shown or assumes. Third, the noise term stands for everything the analyst cannot record, including committee composition, timing and the analyst's own judgement, and its size differs by investor type.

9. Construct a population of managers and investor rules from stated inputs

The experiment constructs 6,000 managers in three types, first-time, second or third fund, and established, in shares of 35%, 30% and 35%. Each manager is screened by five investors drawn from four types, public pensions, sovereign investors, gatekeepers and family offices, in shares of 25%, 15%, 20% and 40%. This gives 30,000 screening pairs. Target fund size is lognormal with a median that rises with manager maturity. True attributes are drawn independently by manager

type, and the probability that an item is documented in the first pack rises with maturity. Table 3 and Table 4 give the main inputs, and the score weights and noise are stated with Equation 2.

Each investor draws a ticket, which with its concentration cap sets the smallest fund it can hold, and gatekeepers and family offices also set a largest fund. Pensions and sovereign investors require a core strategy fit, regulated and administered status and a minimum attributable record. Gatekeepers accept an adjacent strategy and a smaller record. Family offices require neither status nor record. Strictness, the chance that an undocumented gate item is treated as a failure, falls from 90% for sovereign investors to 25% for family offices. Thresholds are set at base so that 35% to 55% of gate passers advance, then frozen.

No input is an observed distribution. The directions follow the sources: written gates for public and sovereign investors, a lower diligence effort and a greater reliance on others at smaller investors, and thinner documentation for younger managers [3, 5, 8]. The magnitudes are authored. The outcome is generated by the same structure the analyst models later use, so the experiment shows how the proposed rule behaves when its premise holds. It does not show that the premise holds. Results should be read as conditional statements: if investors screen in this way, these patterns follow.

Table 3. Investor rule inputs

Investor	Share	Ticket	Cap	Largest fund	Fit needed	Status	Deals	Strict	Advance
Public pension	25%	30	20%	none	core	Yes	6	85%	40%
Sovereign	15%	100	20%	none	core	Yes	12	90%	35%
Gatekeeper	20%	20	15%	1,500	core or adjacent	Yes	4	60%	45%
Family office	40%	5	25%	400 median	core or adjacent	No	None	25%	55%

MatchPoint analysis; stated model inputs, not observed values. Ticket and largest fund in USD million. Cap is the largest share of a fund the investor holds. Status means regulated and administered status is required. Deals is the minimum attributable realised deals. Strict is the probability that an undocumented gate item is treated as a failure. Advance is the share of gate passers that advance at base.

Table 4. Manager inputs

Manager	Share	Median fund	Regulated	Deals	Stable team	GP share	Documented: status / record / performance / team / GP / terms (%)
First-time	35%	150	55%	3	75%	1.5%	60 / 45 / 35 / 80 / 60 / 50
Second or third	30%	300	80%	9	80%	2.0%	80 / 70 / 75 / 85 / 75 / 70
Established	35%	900	95%	24	85%	2.5%	95 / 90 / 90 / 90 / 85 / 85

MatchPoint analysis; stated model inputs, not observed values. Median fund in USD million. Deals is the mean attributable realised deals. The last column gives the probability that each item is in the first pack.

10. The overall advance rate hides opposite experiences for first-time and established managers

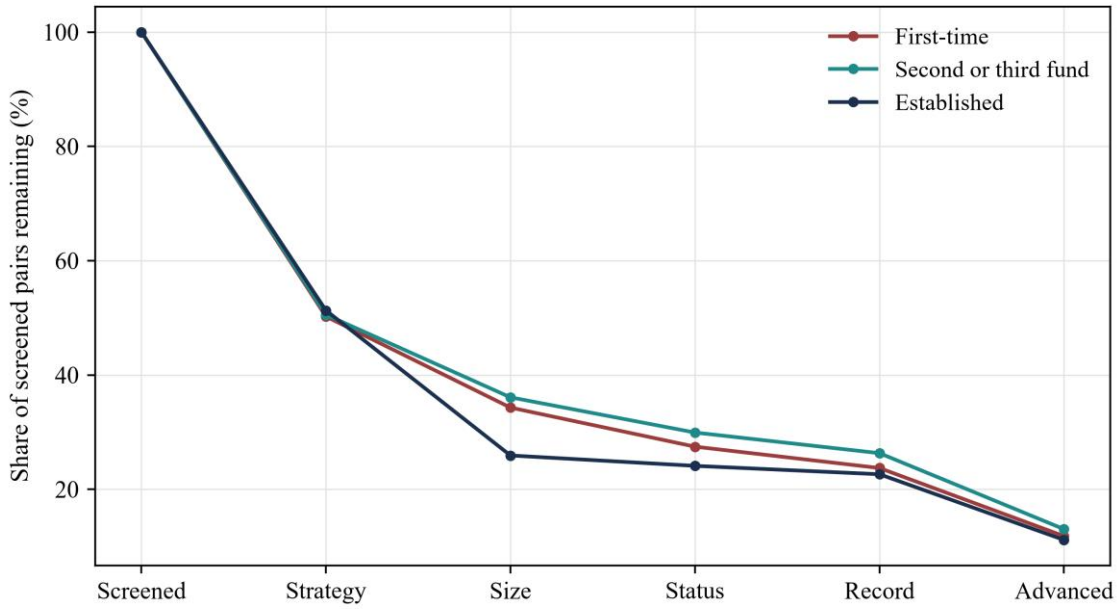
Under the stated inputs, 11.9% of the 30,000 pairs advance, or 3,581 pairs. By manager type the rates are 11.8% for first-time managers, 13.0% for second or third funds and 11.1% for established managers. The ratio of the established rate to the first-time rate is 0.9. Table 5 and Figure 1 follow the funnel. Strategy fit passes for 50.7% of pairs, the size window leaves 31.9%, status leaves 27.0% and the record gate leaves 24.1%. The score then retains 49.5% of the gate passers.

Table 5. Funnel from screened pairs to advancement by manager type

Stage	First-time	Second or third fund	Established	All managers
Screened pairs	100.0%	100.0%	100.0%	100.0%
Strategy within LP mandate	50.2%	50.6%	51.3%	50.7%
Fund size inside LP window	34.3%	36.1%	25.9%	31.9%
Regulatory and operating status	27.5%	29.9%	24.1%	27.0%
Attributable track record	23.7%	26.3%	22.6%	24.1%
Advanced after score	11.8%	13.0%	11.1%	11.9%

MatchPoint analysis of the constructed population. Model output under stated inputs; not an observed result. Gates are applied in the order shown.

Figure 1. Share of screened pairs remaining at each stage



MatchPoint analysis of the constructed population. Model output under stated inputs; not an observed result.

The similarity of the aggregate rates is produced by composition. Table 6 shows that first-time managers advance with 28.3% of family offices, 0.2% of public pensions and 0.0% of sovereign investors. Established managers advance with 5.6% of sovereign investors but with only 5.4% of family offices, because their funds exceed the family offices' largest fund. An aggregate comparison would conclude that track record is irrelevant. The pair-level table shows that it is decisive for some investor types and irrelevant for others.

Table 6. Advance rate by investor type and manager type

LP type	First-time	Second or third fund	Established	All managers
Public pension	0.2%	4.0%	10.6%	5.0%
Sovereign investor	0.0%	0.2%	5.6%	2.0%
Gatekeeper	2.4%	15.9%	27.2%	15.3%
Family office	28.3%	22.2%	5.4%	18.4%

MatchPoint analysis of the constructed population. Model output under stated inputs; not an observed result.

The consequence is plain. The advance rate a manager experiences depends on which investors are on its list, and a list is a choice of gates. A rate measured across a mixed list says little about the manager. A first-time manager approaching only family offices and a first-time manager approaching only sovereign investors face rates of 28.3% against 0.0% in this construction. The same reasoning applies to a placement firm's own list-building rules, such as including or

excluding family offices or sovereign investors, which fix the advance rate it will observe before any manager attribute is considered.

11. Most rejections fail a mandate-fit gate before any manager attribute is read

Table 7 classifies rejected pairs. 86.2% fail at least one gate and 13.8% pass every gate and miss the score threshold. When the gates are tested in the stated order, strategy is the first failed gate for 56.0% of rejections, the size window for 21.4%, status for 5.5% and record for 3.3%. Order matters for this attribution, so the last column counts failures of each gate in any order: strategy 56.0%, size 43.8%, record 29.8% and status 22.1%. These shares are properties of the constructed rule and the stated mix of investor types, and they would change with a different mix.

Failures overlap. 44.4% of rejections fail exactly one gate and 41.8% fail two or more. A manager that repairs one gate in the second group remains rejected, which the next section measures. The two pair-specific gates lead the table because they depend on the investor and not on the manager. A manager with perfect documentation still fails when its strategy lies outside the mandate or its fund lies outside the window. The share that fail two or more gates is the reason the repairs in the next section compound.

Table 7. Rejected pairs by failed gate

Obstacle	First failed gate	Sole failed gate	Fails this gate (any order)
Strategy within LP mandate	56.0%	23.4%	56.0%
Fund size inside LP window	21.4%	14.7%	43.8%
Regulatory and operating status	5.5%	2.9%	22.1%
Attributable track record	3.3%	3.3%	29.8%
Passes every gate, misses score threshold	13.8%	13.8%	n/a
Fails two or more gates	n/a	41.8%	n/a

MatchPoint analysis of the constructed population. Share of rejected pairs. Model output under stated inputs; not an observed result.

The implication for preparation is direct. A manager that opens with the data room prepares documents for investors that cannot advance it. Fit to the investor's mandate and window should be tested first, and manager documentation should follow for the investors that pass. This ordering is a result of the constructed rule and of the assumption that the investor applies written mandates strictly. It would reverse if investors waived fit gates more often than they waived record gates. The assumption is stated in the input register and can be replaced with an investor's own record.

12. A single repair rarely converts a rejection, because gates fail together and repairs compound

The experiment measures what each attribute would change by repairing one attribute at a time for the rejected pairs. A repair sets the attribute to its best verified state and holds the investor's noise fixed, so the pair is the same pair with one input changed. The conversion for repair k is the share of rejected pairs that then advance. The ladder applies the repairs cumulatively, in order of their single conversion. Equation 3 defines both. The measure is a counterfactual that the model can compute and that a cohort of observed pairs cannot.

$$C_k = \frac{1}{|\mathcal{R}|} \sum_{i \in \mathcal{R}} 1\{Y_i^{(k)} = 1\} \quad (3a)$$

$$L_m = \frac{1}{|\mathcal{R}|} \sum_{i \in \mathcal{R}} 1\{Y_i^{(1,\ldots,m)} = 1\} \quad (3b)$$

In Equation 3, R is the set of rejected pairs, Y_i^(k) the outcome of pair i after repair k, C_k the single conversion, and L_m the ladder after the first m repairs in order of C_k. A structural repair cannot remove an advance, which the controls confirm. Table 8 gives the results. Re-targeting to a core strategy fit converts 12.3% of rejected pairs and resizing converts 8.6%. Terms convert 4.1%, performance 3.5%, GP commitment 2.7%, team 1.4%, status 1.1% and record 1.1%.

The conversions differ by manager type in ways that follow from the gates. Resizing converts 1.5% of rejected first-time pairs and 17.3% of rejected established pairs, because established managers are mostly rejected on a window they exceed while first-time managers are rejected on several gates at once. Re-targeting converts 10.3% for first-time managers and 14.0% for established managers. Verifying the record converts 1.1% of first-time rejections alone, and the ladder in the table shows that it is the last gate to bind for these managers.

Table 8. Single repair and cumulative ladder for rejected pairs

Repair, in order of single-repair conversion	Converted alone	Cumulative ladder, all pairs	Cumulative, first-time	Cumulative, established
Re-target to an LP with core strategy fit	12.3%	12.3%	10.3%	14.0%
Resize the fund into the LP window	8.6%	27.4%	13.1%	44.3%
Verify fund terms at 0.85 alignment	4.1%	37.7%	21.5%	56.2%
Verify net performance 1 SD above benchmark	3.5%	46.5%	26.2%	68.7%
Verify a cash GP commitment of 3%	2.7%	51.9%	30.9%	74.9%
Verify a stable named team	1.4%	54.6%	32.8%	78.2%
Verify regulated and administered status	1.1%	61.3%	38.5%	82.8%
Verify attributable record at LP minimum	1.1%	86.8%	86.0%	87.5%

MatchPoint analysis of the constructed population. Share of rejected pairs converted to advancement with the investor's noise unchanged. Repairs set the attribute to its best verified state: core fit, size inside the window, status and record verified at the investor minimum, performance one standard deviation above benchmark, a stable team, a 3% cash GP commitment and terms at 0.85 alignment. Model output under stated inputs; not an observed result.

The ladder shows the compounding. After strategy and size the ladder stands at 27.4%, and after all eight repairs it reaches 86.8%. For first-time managers the ladder stands at 38.5% before the record repair and rises to 86.0% only when the record is added, because the strict investors reject first-time managers on record alone. A single-repair ranking understates the attribute that is last to bind. A manager that fixes weak points one at a time sees little change in its advance rate until the final gate falls, so a readiness plan should be sequenced against the investor list and not scored attribute by attribute.

13. Documentation converts few rejections and removes some advances, because disclosure reveals the truth

The repairs above set an attribute to its best state and verify it. Documentation alone is a weaker action: it reveals the manager's true level, which may be above or below the gate. When every item is documented in the first pack, 3.1% of rejected pairs convert, and 4.2% of the pairs that advanced no longer do. Documenting regulated status alone converts 0.44% of rejected pairs and

documenting the attributable record alone converts 0.67%. The comparison is between the same manager documented and undocumented, with the investor and its noise unchanged.

The mechanism is the credit for undocumented items. Lenient investors credit an undocumented gate item as a pass with a probability of 40% to 75% in this construction, and full documentation removes the credit where the true level is below the gate. Strict investors treat an undocumented item as a failure with a probability of 85% to 90%, so documentation helps only where the true level passes. Which effect dominates depends on the true levels the manager holds. This is consistent in direction with the finding that interim performance moves fundraising most when it is backed by exits [13], but that study concerns fundraising outcomes and not first screening.

The implication is for the order of work and not for disclosure policy. A manager should know its true position against each investor's gate before it opens the pack, because the pack converts or removes advances according to that position. Nothing in the model supports withholding material information or overstating a record. It shows only that documentation is a measurement of the manager against the gate, and that the measurement should precede the approach. The decision to document becomes a decision about which investors to approach with which pack.

14. Manager attributes alone rank advancement poorly, and investor-specific fit restores a ceiling-level ranking

The ranking test compares three analyst models with a ceiling and two floors. The ceiling is an oracle that knows the investor's gates and score and sees what the investor sees, with the probability in Equation 4. The gate-aware model applies the published knock-outs, setting the probability to zero for a verified failure, and fits a logit to the remaining pairs. The main-effects logit uses every observed attribute, including the window and strategy flags. The manager-only logit uses documented manager attributes, with no investor type, strategy fit or window. The floors are verified performance alone and random order.

$$P_{ij} = [\prod_{k \in G} p_{ijk}][1 - \Phi((\tau_j - \mu_{ij})/\sigma_j)] \quad (4a)$$

$$\text{AUC} = Pr(\hat{p}_i > \hat{p}_l \mid Y_i = 1, Y_l = 0) + \tfrac{1}{2} Pr(\hat{p}_i = \hat{p}_l \mid Y_i = 1, Y_l = 0) \quad (4b)$$

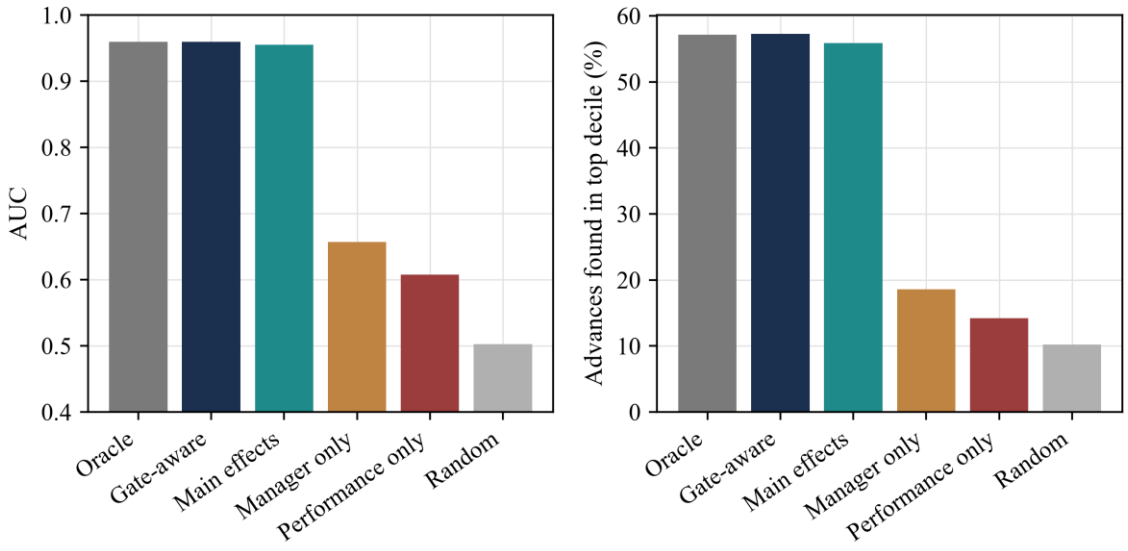
In Equation 4, P_ij is the probability that pair ij advances, p_ijk the probability that gate k is passed given its documentation state (zero or one when documented, one minus strictness when not), Phi the standard normal distribution function, mu_ij the mean score, sigma_j the investor noise and tau_j the frozen threshold. The area under the curve, AUC, is the probability that a randomly chosen advancing pair is ranked above a randomly chosen rejected pair, with ties at half. Table 9 and Figure 2 report AUC and the share of advances found in the top decile.

Table 9. Ranking of screening pairs against ceiling and floor

Ranking of screening pairs	AUC	Share of advances in top decile	Mean predicted rate
Ceiling: oracle with the LP's own gates and score	0.960	57.2%	11.7%
Gate-aware: published knock-outs, then logit	0.960	57.2%	11.9%
Main-effects logit on observed attributes	0.955	55.9%	11.9%
Manager attributes only, no LP-specific fit	0.657	18.6%	11.9%
Verified performance alone	0.608	14.2%	n/a
Floor: random order	0.503	10.2%	n/a

MatchPoint analysis of the constructed population. Models are fitted on the base population and evaluated on the same pairs. Model output under stated inputs; not an observed result.

Figure 2. Ranking performance of three analyst models against ceiling and floors



MatchPoint analysis of the constructed population. AUC and share of advances in the top decile. Model output under stated inputs; not an observed result.

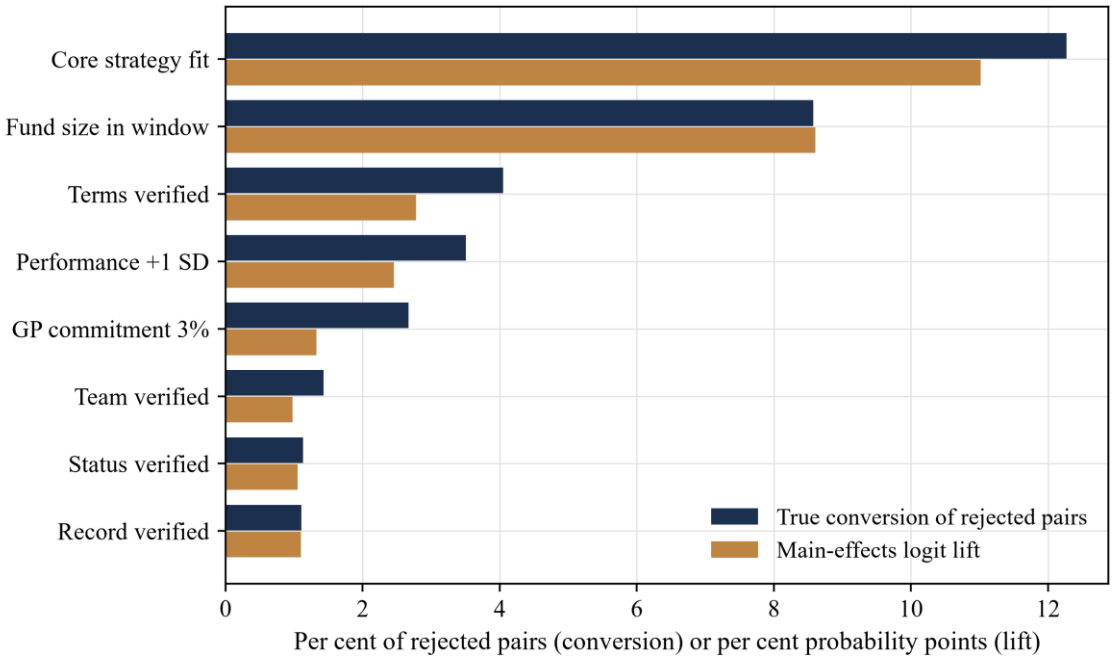
The oracle reaches an AUC of 0.960 and the gate-aware model 0.960. The main-effects logit reaches 0.955, which is consistent with an additive log-odds model with large negative coefficients on failed-gate flags approximating a conjunction. The manager-only model reaches 0.657, and verified performance alone reaches 0.608, against 0.503 for random order. The top decile holds 57.2% of advances under the gate-aware ranking, 18.6% under the manager-only ranking and 14.2% under performance alone. The investor-specific flags carry the ranking. Verified performance, the attribute investors in the survey weight most [8], barely lifts it.

The gate-aware model is also calibrated at base, as it should be when the premise holds. In its top decile the predicted advance rate is 68.1% and the simulated rate is 68.3%. Calibration matters for the next sections, because a stress that moves the investor's rule leaves the ranking nearly intact and breaks the forecast, as the stress cases show. A team that reports a ranking without a calibration check cannot tell the two situations apart. The check is inexpensive and should be reported with every ranking.

Figure 3 compares the logit's lift with the true conversion. The logit's lift tracks the true conversion in rank, with a Spearman correlation of 0.90 and 3 of the top three repairs in

agreement, but it understates the soft attributes. For performance it implies a lift of 2.5% against a true conversion of 3.5%, which matters for a team that uses a fitted model to decide where to spend effort on documentation. The finding that matters for the decision is plain: in this construction a manager-quality score cannot determine advancement by itself, because it cannot see the investor's mandate or window. The near equality of the gate-aware and main-effects models shows that the gain comes from investor-specific information and not from the conjunctive form of the rule.

Figure 3. True conversion of rejected pairs against the lift implied by the main-effects logit



MatchPoint analysis of the constructed population. Conversion as defined in Equation 3; lift is the mean change in the logit's predicted probability for the same repair. Model output under stated inputs; not an observed result.

15. The fund-size window is a numerical constraint set by investor tickets and caps

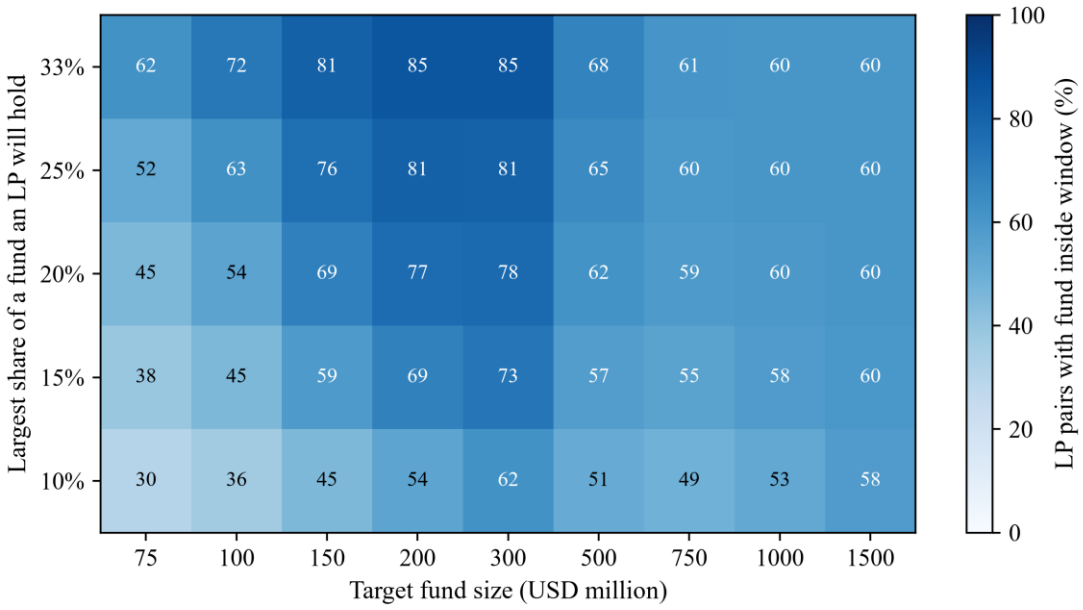
For a ticket T and a concentration cap c, the smallest fund the investor can hold is T divided by c, and some investors also set a largest fund F_max. The fund size F_i passes when it falls between the two, as Equation 5 states. The first-close target is therefore part of investor fit and not only a fundraising choice. The window is computed for each pair from the investor's own ticket, so one fund can pass for one investor and fail for another. A target below the minimum of an investor with a large ticket cannot be rescued by documentation.

$$F_i \geq \frac{T_j}{c_j} \quad (5a)$$

$$Pr(\text{window}) = Pr(\frac{T_j}{c_j} \leq F_i \leq F_j^{\max}) \quad (5b)$$

In Equation 5, T_j is the investor's ticket, c_j its largest share of a fund and F_j^max its largest fund where one is set. Figure 4 gives the share of investor pairs whose window contains a target size for five caps. At a cap of 20%, a USD 150 million fund fits 69% of pairs and a USD 75 million fund fits 45%. At a cap of 10% the USD 150 million fund fits 45%. A USD 500 million fund fits 62% at a 20% cap and a USD 1,500 million fund fits 60%.

Figure 4. Share of investor pairs whose window contains the target fund size



MatchPoint analysis of the constructed investor population. Cells give per cent of pairs. Model output under stated inputs; not an observed result.

The share is not monotone in fund size. Small funds fail large tickets, and large funds exceed the largest fund set by gatekeepers and family offices. No single size is right in this construction, and the best size depends on the list. CalPERS's cited 2024 policy sets a 35% cap on commitments to one fund. LAFPP's cited real-estate emerging-manager criteria separately cap an allocation at 20% of the capital raised for a specific commingled vehicle [3, 4]. The model uses its own assumed caps and tickets. The decision for a manager is to choose the target size against the list it intends to approach.

16. Four stress cases leave three signatures that a manager can tell apart

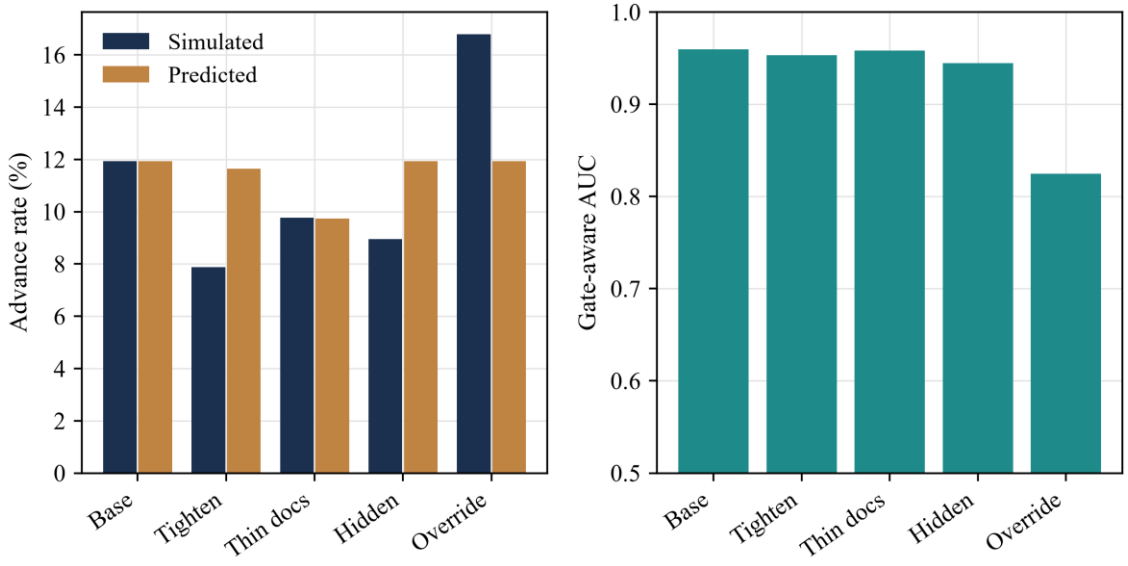
The forecast model retains its base-case coefficients. Each case changes one aspect of the simulated decision process. In the first case investors tighten: concentration caps fall by 20%, so the smallest fund rises by 25%, strictness rises by five points and score thresholds rise by 0.35 score units. In the second, documentation thins: a quarter of documented items drop out, and the share of undocumented items rises from 27.3% to 45.3%. In the third, a hidden investor screen rejects 25% of otherwise advancing pairs on a criterion absent from the analyst's data. In the fourth, warm relationships override the screen: each pair with a prior relationship has a 50% chance of being selected for an exception. A selected pair advances if its score is at least the threshold minus 0.5 score units, regardless of the screening criteria.

Table 10. Stress cases with an unchanged forecast model

Case	Simulated advance rate	Gate-aware predicted rate	Gate-aware AUC	Advances among knock-out pairs	Undocumented items	Resize repair converts	Re-target repair converts
Base	11.9%	11.9%	0.960	0.00%	27.3%	8.6%	12.3%
LPs tighten	7.9%	11.6%	0.953	0.00%	27.3%	6.3%	7.9%
Thinner documentation	9.8%	9.7%	0.958	0.00%	45.3%	7.4%	9.2%
Hidden LP screen	8.9%	11.9%	0.945	0.00%	27.3%	6.2%	8.9%
Relationship override	16.8%	11.9%	0.825	5.61%	27.3%	8.2%	12.7%

MatchPoint analysis of the constructed population. Gate-aware model fitted at base. Repair columns show the share of rejected pairs converted under each stress. Model output under stated inputs; not an observed result.

Figure 5. Simulated and predicted advance rates, and ranking, under four stress cases



MatchPoint analysis of the constructed population. Gate-aware model fitted at base and frozen. Model output under stated inputs; not an observed result.

When investors tighten, the simulated rate falls to 7.9% against a predicted 11.6% while the AUC holds at 0.953: the ranking is intact and the forecast overshoots. When documentation thins, the simulated rate falls to 9.8% and the prediction falls with it to 9.7%, with an AUC of 0.958, because the model reads the documentation state. A hidden screen leaves the simulated rate at 8.9% against a prediction of 11.9%, a slightly weaker AUC of 0.945 and lower repair conversions. A relationship override raises the simulated rate to 16.8% above the prediction of 11.9%, cuts the AUC to 0.825 and lets 5.61% of knock-out pairs advance.

Table 10 and Figure 5 summarise the cases. The simulated patterns suggest where to investigate. A forecast that overshoots while the ranking nearly holds indicates either moved thresholds or a criterion the manager cannot see: here the two cases differ by less than 0.01 in AUC (0.953 against 0.945) and both lower the repair conversions, so advance data alone cannot separate them and the investor's stated reason is needed. In the simulation, a decline that tracks documentation coverage points to fewer verified items. Advances after failure of a verified criterion flag possible exceptions. A relationship is one possible explanation to check with the investor.

17. Selection into the approached cohort overstates every gap, and weighting repairs only part of it

A manager approaches the investors it expects to respond. The model draws the approach from a logit with a constant of -2.4, a coefficient of 1.2 for a warm relationship, 0.9 for core fit, 0.7 for a size window that fits and 0.8 per standard deviation of the investor's unrecorded appetite. The cohort contains 23.8% of the pairs. Its advance rate is 21.5% against 11.9% for the whole longlist, and it contains 29.8% size failures against 38.6% on the longlist.

$$\hat{\Delta}_w = \frac{\sum_i w_i y_i a_i}{\sum_i w_i a_i} - \frac{\sum_i w_i y_i b_i}{\sum_i w_i b_i} \quad (6a)$$

$$w_i = 1/\hat{\pi}_i, \hat{\pi}_i = \text{logit}^{-1}(\hat{\beta}_0 + \hat{\beta}_1 \text{warm}_i + \hat{\beta}_2 \text{core}_i + \hat{\beta}_3 \text{window}_i) \quad (6b)$$

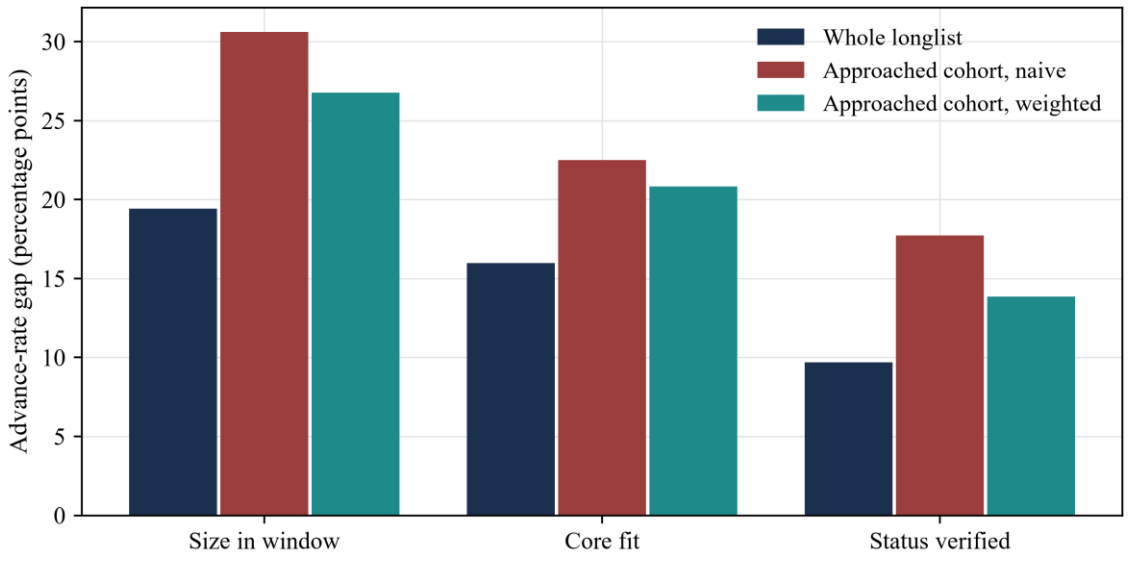
In Equation 6, y_i is the outcome, a_i and b_i indicate the two groups being contrasted, w_i an inverse-probability weight, and pi_i the probability of being approached estimated from the recorded covariates over the whole longlist. Appetite is not recorded, so the weights cannot correct for it. Table 11 and Figure 6 give three contrasts. The size gap is 19.4 points on the longlist, 30.6 in the naive cohort and 26.8 after weighting, which repairs 34% of the bias. The strategy gap moves from 16.0 to 22.5 and 20.8, repairing 26%, and the status gap from 9.7 to 17.7 and 13.9, repairing 48%.

Table 11. Advance-rate gaps on the whole longlist and in an approached cohort

Contrast (advance rate, percentage points)	Whole longlist	Approached cohort, naive	Approached cohort, weighted	Contrast-group pairs in cohort
Fund size inside versus outside LP window	19.4	30.6	26.8	2,132
Core versus non-core strategy fit	16.0	22.5	20.8	3,865
Regulated status verified versus undocumented	9.7	17.7	13.9	1,146

MatchPoint analysis of the constructed population. Gaps in percentage points. Weights use recorded covariates only; the unrecorded investor appetite remains. Model output under stated inputs; not an observed result.

Figure 6. Gaps on the longlist, in the naive cohort and after weighting



MatchPoint analysis of the constructed population. Model output under stated inputs; not an observed result.

The residual bias is the investor appetite that led the manager to approach those investors. In this construction every gap measured on a cohort the manager chose therefore overstates the association measured in the full simulated population. Weighting on recorded covariates reduces the discrepancy. To estimate an association less sensitive to this selection, include a subset of pairs approached without regard to expected response. Selection also thins the cohort of the pairs the question most concerns: a cohort of approached pairs holds few of the size failures.

18. Several hundred screenings can detect the documentation gap, but cannot rank the attributes

A real team would hold a small number of screenings. The experiment draws random cohorts of 100, 250, 500, 1,000 and 2,500 pairs from the population, as Table 12 shows, 1,000 replications each, and measures the gap in advance rate between pairs with verified regulated status and pairs with undocumented status. The gap in the population is 9.7 points. At 100 screenings the middle 95% of estimates runs from -2.3 to 22.6 points and the gap is detected in 5% of replications. At 500 the range is 4.5 to 15.3 points and detection is 84%. At 2,500 the range is 7.4 to 11.9 points.

Table 12. Sampling spread of a random study cohort

Random screenings in the study	95% range of advance rate (%)	95% range of verified-versus-undocumented gap (points)	Replications detecting the gap
100	6.0 to 19.0	-2.3 to 22.6	5%
250	8.0 to 16.0	1.8 to 16.9	44%
500	9.2 to 14.8	4.5 to 15.3	84%
1,000	10.0 to 14.0	6.0 to 13.8	100%
2,500	10.7 to 13.1	7.4 to 11.9	100%

MatchPoint analysis of the constructed population. Random cohorts drawn without replacement; detection is a two-sided two-proportion z-test at the 5% level, counted only where the estimated gap has the correct sign. Model output under stated inputs; not an observed result.

The advance rate itself is also uncertain: at 100 screenings its middle 95% runs from 6.0% to 19.0%, and at 500 from 9.2% to 14.8%. The table covers one contrast. Ranking several attributes whose gaps overlap, and estimating the repair conversions of Equation 3, needs more screenings and, for conversions, an experiment, because conversions are counterfactual and no cohort of observed pairs contains them. Wider ranges at smaller sizes mean that a point estimate from a small cohort should not be read as a finding.

A single placement may screen far fewer pairs than the 500 at which detection reaches 84%, and the table illustrates one modelled contrast rather than setting a field-study minimum. The practical reading is that a team should not report an attribute ranking from one mandate. It can report the gate failures it observed, and it can use comparable screenings pooled across mandates to estimate associations, subject to permission and lawful processing. Estimating an effect requires a separately specified causal design. A team can still record the first failed gate of each rejection, which is informative without a large sample.

19. The calculations reconcile, and the controls establish internal consistency only

The model asserts 57 control identities. They include the pair count, probability sums, the identity linking ticket, cap and minimum fund size, threshold targets, and funnel reconciliation. The first-failure shares plus score-only failures sum to one. The sole-failure shares plus multiple-criterion and score-only failures also sum to one. Other checks cover the ladder's monotonicity, the property that a structural repair cannot remove an advance, convergence of the logit score equations, the equality of fitted and simulated means, the ordering of the oracle and analyst AUC values, the agreement of the oracle mean probability with the simulated rate, regeneration with the same seed, and a null stress that reproduces the base case.

Table 13. Reproducibility controls

Control	Result
Screening pairs	30,000 pairs from 6,000 managers and 5 LPs each
Control identities asserted	57, all passed
Seed and regeneration	Fixed seed 1386; second generation reproduces every decision
Thresholds	Calibrated to base targets; rise only in the investor-tightening scenario
Logit fits	Score equations below 0.00001; fitted mean equals simulated mean
Repairs	Same LP noise for the pair; structural repairs never remove an advance
Cohort replications	1,000 random cohorts at each of 5 sizes
Null stress	A stress with no change reproduces the base advance rate exactly
Observed data used	None. All results are model output under stated inputs

MatchPoint analysis. Controls test identities, monotonicity and reproducibility. They do not test the assumptions behind the inputs.

Table 13 lists the controls. They do not include an independent second implementation of the whole experiment. A separate closed-form calculation of the size-window shares in Figure 4 from the stated ticket and cap inputs agrees with the simulated shares to within one percentage point, and a binomial calculation reproduces the advance-rate ranges in Table 12. They establish that the results follow from the stated inputs. They do not establish that the inputs describe any investor. The outcome is generated from the structure the analyst models use, which is circular by design, and the figures show how the method behaves when its premise holds. They are not evidence that the premise holds.

20. A consented study needs timestamped screening records and a random audit sample of rejections

A real study would take the manager and investor pair as its unit and record the dates of the first meeting and of any later action, so that the outcome in Section 2 can be counted. It would record the first-pack documentation state of each attribute and the investor's written criteria where available, and it would separate silence from an explicit rejection. Table 14 states the elements. Agree participation and confidentiality terms with investors and managers. If the study processes identifiable representatives' personal data, document an applicable lawful basis before collection. Where DIFC law applies and consent is that basis, the individual must freely give clear affirmative consent, and the controller must be able to demonstrate it. Participant institutions' agreement does not itself establish consent from individuals. Other jurisdictions and duties require separate assessment [15].

Table 14. Evidence design for a study of first-meeting advancement

Element	Requirement	Reason
Unit and outcome	Manager–investor pairs with a first meeting and an ascertainable outcome by day ninety	Matches the decision in the rule
Timestamps	First meeting, pack, first qualifying diligence action, rejection and follow-up cutoff dates	Separates delay from rejection
Attribute state	Documented or not, and the documented level, at the date of the first pack	Distinguishes true level from evidence
Investor criteria	Written gates, caps and tickets, with the date of the document	Allows the gate model to be tested
Two sampling controls	Random approach order among eligible pairs; separate random audit of explicit rejections	Limits order selection; tests recorded reasons
Sample size	Set target from baseline rate, material difference, clustering and missing outcomes	Justifies the field sample; model sizes are illustrative
Analysis plan	Fixed before data: risk differences with intervals, gate-level audit, no attribute ranking from one mandate	Prevents choosing the gap after the fact
Consent and data	Participation terms; confidentiality and a documented lawful basis where required	Required before any collection

MatchPoint design. Each row states what a real study would collect and why. No such dataset is held or used in this paper.

Use two separate controls. Randomly ordering outreach among comparable eligible pairs can reduce selection caused by prioritising particular pairs first. Independently, sample rejected pairs at random and ask investors to identify the criterion they applied; record unanswered audits and report the response rate. That sample can reveal discrepancies between recorded and stated reasons. Neither design identifies the effect of changing a manager attribute. One possible intervention would enrol one eligible pair per manager immediately after its first meeting, before any qualifying diligence action. Randomly offer additional documentation support at enrolment. Pairs that have already reached the endpoint remain in the observational cohort and are ineligible for this intervention. Compare all assigned pairs by the offer they received, whether or not they used it. Apply the same ninety-day endpoint and prespecify how to report and analyse incomplete follow-up. This estimates the effect of offering support for the enrolled population; a change in any individual attribute remains a separate question. Fix the analysis in advance and report uncertainty, incomplete responses and the population each estimate describes.

21. Report the result to the client and the investment committee as a fit test, an order of work and a stated uncertainty

A client deciding whether to approach an investor needs to know which gate it fails, whether the failure is repairable and in what order. We propose a one-page report with three parts. The first lists the investors on the list with their written gates and the client's position against each. The second gives the repair ladder for the failing pairs, in the order the model supports. The third states the uncertainty: how many screenings the statement rests on and which investor behaviour is assumed.

An investment committee of a placement firm or a manager reviewing its fundraise needs a different statement. It needs the share of its list that passes the fit gates, the expected advance rate under the stated assumptions with a range, and the stress signatures it should watch. A report that quotes one advance rate without the composition of the list invites the error in Section 10. A report that ranks attributes from one mandate invites the error in Section 18. Each report should say that the results of this paper are constructed and not observed.

The report should avoid language that implies a prediction about fund returns or a commitment to an investor. Advancement is a decision about whether to examine a manager. It is not a decision to invest, and this paper reports nothing about the quality or performance of any manager. A pair rejected for failing a gate has failed a written criterion of that investor. A pair that advances has entered formal diligence through one of the events defined in Section 2. The report should also name the investors for which the gates are unknown.

22. Data and AI can keep the gate inputs current, but they do not replace the investor's own statement

The gate inputs are mostly public or documentary: published policies, regulator registers, concentration limits and the manager's pack. Software can read these documents, extract the stated criteria, date them and flag changes. That reduces the cost of keeping the fit table current across a long list, and it makes visible which investors have published no criteria at all, so that their gates are recorded as unknown rather than inferred. The extraction should be checked against the document by a person, with the source and date retained for every criterion, because a missed revision of a policy changes the gate.

A language model can also draft the manager's position against each gate, and a classifier can estimate the probability of advancement from past pairs. The constructed results show the limits. A classifier trained on a chosen cohort inherits the selection bias in Section 17 and cannot see an unrecorded criterion, as the hidden-screen case in Section 16 shows. We propose that any model used for this decision be evaluated on an independently sampled set of eligible first-meeting pairs, report calibration against observed advances, and be monitored for advances among pairs documented as failing a gate.

Human decision rights should remain explicit. The investor decides whether to advance a manager, and the manager or its adviser decides whether to approach an investor. A model can inform the second decision. It should not be presented to an investor as a statement of the investor's own policy unless the investor has provided it. This paper reports no model trained on observed screening pairs, no measured accuracy on observed pairs and no improvement in real outcomes. Any tool used should record the version of each document it has read and the date it was last checked.

23. A decision rule for a live fund placement

For a live placement we propose five steps. First, list the investors and record each written gate, with its source and date, and mark every gate that cannot be sourced as unknown. Second, test the fund's strategy and target size against each list entry and remove or defer investors that fail a gate the manager cannot repair. Third, for the remaining investors, compare the manager's true attribute levels with the gates before opening the pack. Fourth, document the items that pass and prepare the repair ladder for those that fail. Fifth, where appropriate, randomise approach order among comparable eligible investors and log every approach, whether it led to a meeting, and the diligence outcome where ascertainable.

The rule has a stopping condition. If fewer than a stated share of the list passes the fit gates, the issue is the fund's design or the list, and more documentation will not change it. If the fit share is adequate and advancement remains low, the stress signatures in Section 16 indicate where to look. The thresholds for these tests are for the team to set from its own list and its own tolerance.

This paper supplies no universal value, and the constructed numbers must not be used as one. The rule is a sequence of questions and not a score, so that a reviewer can see at which step each investor was removed or kept.

The rule makes no promise that an investor will advance a manager or commit capital. It orders the work so that effort goes first to the gates that decide the outcome in the constructed rule. It records the assumptions on which the order rests. Where a real investor does not behave as the rule assumes, the record of exceptions will show it, and the order should then be revised for that investor. A placement team that keeps this record across mandates will, over time, hold the evidence that this paper does not have.

24. Limitations and the next decision

The experiment is a construction. Its investors apply written gates and a noisy score because we made them do so. Its tickets, caps, strictness values, documentation probabilities and thresholds are authored, its outcome is an investor decision and not a fund return, and its population has four investor types and eight attributes. It has no time dimension and no feedback from one pair to the next. The sources support the categories of gate and score. None supports a magnitude, and none reports an advancement rate for any investor or any region.

Several omissions matter. Investors differ within a type, and the constructed types are averages. Managers adapt their pack and their fund size after rejection. Investors share information with each other, so the pairs are not independent. The regulator and policy documents reviewed are two US public-pension policy documents, a voluntary sovereign document, a regulator summary page and rulebook provisions, and staff views. The source set includes no Gulf investor's own policy. The surveys and working papers are dated, and several are working-paper versions.

The next decision is empirical. A team that holds screening records across mandates under agreed participation terms and an applicable lawful basis where personal data are processed should test whether advancement follows the gate structure, whether the first failed gate is the investor's stated reason, and whether reasons stated in an audit of randomly sampled explicit rejections differ in distribution from the reasons already recorded. Until then, the result to carry forward is conditional: where investors apply written gates before scoring, a manager's preparation should follow the order of the gates. A negative finding would be as useful as a positive one, because it would show which gate investors do not apply.

Frequently Asked Questions

Does this paper show which manager attributes real investors use to decide?

It does not. The investor decisions are generated from a stated two-stage rule, and the analysis then asks how that rule behaves. The results show what follows if investors apply written gates before a noisy score. They do not show that any investor does so, and no observed screening record is used. A real study of consented, timestamped screenings would be needed to test the structure and to estimate any gate level or weight.

Does a stronger track record raise the chance that an investor advances a manager?

In the construction, performance is a score item with an assumed weight, and verified performance alone ranks pairs with an AUC of 0.608, nearer to random order at 0.503 than to the 0.960 reached with investor-specific fit. The published evidence cited here finds little persistence

from a prior fund's rank for post-2000 buyouts when only information available at fundraising is used [10]. The paper tests no relationship between performance and fund returns.

Should a first-time manager avoid public pensions and sovereign investors?

The construction gives first-time managers 0.2% advancement with public pensions and 0.0% with sovereign investors, because those investors are assumed to require a minimum attributable record. That is an assumption. One public pension publishes a programme for first- to third-time funds of about USD 500 million or less [3]. A manager should read each investor's written criteria and test its own position against them before choosing a list.

Is it advisable to document every item in the first pack?

The model shows that documentation reveals the true level of each item. Documenting everything converts 3.1% of rejected pairs and removes 4.2% of advances that depended on credit for undocumented items. The result is an argument for knowing the position against each gate before the approach. It supports no withholding of material information, and a manager must give accurate and complete information to any investor.

How many screenings are needed to say that an attribute matters?

For one contrast, the verified-versus-undocumented status gap, random cohorts of 100 screenings detect the gap in 5% of replications and cohorts of 500 in 84%. Ranking associations across several attributes needs a larger, representative sample. Estimating whether a repair causes advancement requires a separate intervention or another credible causal design. Pool records across consenting mandates where the design and applicable law permit.

Does advancement mean the investor will commit?

It does not. Advancement here means that a formal diligence stage opens within the stated window. The investor can still close the file at any later stage. The paper measures no commitment, no fund return and no time to first close, and it makes no statement about the quality of any manager or fund.

References

[1] Institutional Limited Partners Association. ILPA Principles 3.0: Fostering Transparency, Governance and Alignment of Interests for General and Limited Partners. 2019. Voluntary best practice; GP commitment and key person sections read. https://ilpa.org/wp-content/uploads/2019/06/ILPA-Principles-3.0_2019.pdf

[2] Institutional Limited Partners Association. ILPA Due Diligence Questionnaire 2.0. Updated November 2021. Voluntary template, not a required document; contents list and selected items read. https://ilpa.org/wp-content/uploads/2021/11/ILPA-DDQ-2.0.pdf

[3] Los Angeles Fire and Police Pension System. Emerging Managers Policy (6.0), revision history to 21 December 2023. Policy of one US public pension. https://content.lafpp.lacity.gov/wp-content/uploads/2023/09/investments-emergingmanagerspolicy.pdf

[4] California Public Employees' Retirement System. Total Fund Investment Policy, Private Equity Program Policy, concentration and fund investment limitations; posted policy, accessed 24 November 2026. https://www.calpers.ca.gov/sites/default/files/spf/docs/total-fund-investment-policy.pdf

[5] International Working Group of Sovereign Wealth Funds. Sovereign Wealth Funds: Generally Accepted Principles and Practices, the Santiago Principles. October 2008, GAPP 18 and 18.2 commentary. Voluntary. https://www.ifswf.org/sites/default/files/santiagoprinciples_0_0.pdf

[6] Dubai Financial Services Authority. Collective Investment Funds, key features of the DFSA funds regime. Regulator summary page, undated, accessed 14 October 2026. https://dfsa.ae/what-we-do/collective-investment-funds

[7] U.S. Securities and Exchange Commission, Division of Investment Management. Marketing Compliance Frequently Asked Questions. Page reviewed 15 January 2026; private fund IRR answer of 6 February 2024 and extracted performance answer of 19 March 2025; staff views without legal force. https://www.sec.gov/rules-regulations/staff-guidance/division-investment-management-frequently-asked-questions/marketing-compliance-frequently-asked-questions

[8] Da Rin, M. and Phalippou, L. The Importance of Size in Private Equity: Evidence from a Survey of Limited Partners. ECGI Finance Working Paper 408/2014, revised July 2016; published in Journal of Financial Intermediation, 2017. https://ecgi.global/sites/default/files/2026-10/fin-408-2014_the_importance_of_size_in_private_equity_evidence_from_a_sur.pdf

[9] Lerner, J., Schoar, A. and Wong, W. Smart Institutions, Foolish Choices? The Limited Partner Performance Puzzle. NBER Working Paper 11136, February 2005. https://www.nber.org/system/files/working_papers/w11136/w11136.pdf

[10] Harris, R. S., Jenkinson, T., Kaplan, S. N. and Stucke, R. Has Persistence Persisted in Private Equity? Evidence from Buyout and Venture Capital Funds. NBER Working Paper 28109, November 2020. https://www.nber.org/system/files/working_papers/w28109/w28109.pdf

[11] Kaplan, S. N. and Schoar, A. Private Equity Performance: Returns, Persistence and Capital Flows. NBER Working Paper 9807, June 2003; published version not reviewed. https://www.nber.org/system/files/working_papers/w9807/w9807.pdf

[12] Hochberg, Y. V. and Rauh, J. D. Local Overweighting and Underperformance: Evidence from Limited Partner Private Equity Investments. NBER Working Paper 17122, June 2011; abstract read; published version not reviewed. https://www.nber.org/system/files/working_papers/w17122/w17122.pdf

[13] Barber, B. M. and Yasuda, A. Interim Fund Performance and Fundraising in Private Equity. Working paper dated 29 October 2015; published in Journal of Financial Economics, 2017. https://conference.nber.org/conf_papers/f82479/f82479.pdf

[14] Dubai Financial Services Authority. DFSA Rulebook, Collective Investment Rules (CIR), version VER40/01-26, Rules 15.1.2 and 15.1.3 on prospectus disclosure for foreign funds; binding rules, current at access on 14 October 2026. https://dfsaen.thomsonreuters.com/rulebook/cir-1513

[15] Dubai International Financial Centre. Data Protection Law, DIFC Law No. 5 of 2020, July 2025 consolidated text, Articles 9, 10 and 12. Applies within its specified scope; accessed 24 November 2026. https://assets.difc.com/v1/media/edge/images/dubaiintern0078-difcexperie96c5-production-3253/media/project/difcexperiences/difc/difcwebsite/documents/laws--regulations/data-protection-law.pdf

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

https://www.linkedin.com/in/ckadya/

https://www.matchpoint-partners.com/team/ck-adya.html

Disclaimer

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.