

Keystone: *commit or pass* on a \$2.5B mid-market buyout Fund V.

A pension plan's private-equity program is offered a \$75M commitment to Fund V of an incumbent North American lower-mid-market buyout manager (industrials and business services). Three priors are doing the work: a 56% fund-size step-up, a return stack still leaning on multiple expansion, and a largely-unrealized Fund IV. Keystone underwrites the manager, not the marketing: track-record decomposition, PME benchmarking, and value-creation attribution, delivered as a commit/pass recommendation.


Author S. Ize-Iyamu **Audience** Investment Committee **Length** 3 pages **Status** Diligence memo
Sleeve Funds & Co-Investments

The Question

A re-up is the default; the diligence job is to falsify it. The GP has compounded a credible lower-mid-market franchise across three prior vintages, but the Fund V pitch carries three things an LP must price independently. **Size step-up:** Fund V targets \$2.5B vs Fund IV's \$1.6B (a 56% jump); relative (percentage) returns tend to compress as funds scale, and absorbing the capital means either more deals or larger checks, the latter implying style drift up-market into more efficiently-priced auctions. The size-return link is partly reversion-to-the-mean, so it warrants a deployment plan, not an automatic pass. **Return composition:** roughly 40% of realized Funds II-III value came from multiple expansion (in line with Bain's buyout decompositions), a tailwind that cannot be relied on in a higher-rate regime; the durable question is whether organic EBITDA growth can carry Fund V. **Maturity:** Fund IV (2019) is 0.4x DPI, mostly unrealized, so its 1.7x TVPI is a mark, not cash. The recommendation must hold under all three.

FIGURE 1 · TRACK RECORD · NET TVPI BY VINTAGE, REALIZED VS UNREALIZED

$$\text{NET TVPI} = \text{REALIZED DPI} + \text{UNREALIZED RVPI} \cdot \text{NET IRR} \cdot \text{KS-PME}$$



Returns compress vintage-over-vintage (2.4x → 2.1x → 1.7x marked), and the realized share thins toward the recent funds. The underwrite holds Fund V to 2.0x net only if organic EBITDA growth replaces the multiple-expansion tailwind.

Sizing the commitment

The ask is **\$75M into a \$2.5B fund** (a 3% LP stake), inside a private-equity sleeve pacing roughly **\$300M/yr** against a 12% private-markets target allocation. At a re-up of this size the position is a top-ten GP relationship: the LP is buying access to the next two vintages and the co-investment flow that travels with them. Base-case underwrite is **2.0x net TVPI / 16-18% net IRR**; the bear case is fund-size dilution plus a one-turn exit-multiple contraction, landing near **1.5x net**.

Benchmarks per Cambridge Associates / Burgiss pooled medians.

<p>COMMITMENT REQUESTED</p> <p>\$75M</p> <p>3% of a \$2.5B fund</p>
<p>BASE-CASE NET TVPI</p> <p>2.0x</p> <p>16-18% net IRR</p>
<p>FUND-SIZE STEP-UP</p> <p>+56%</p> <p>\$1.6B → \$2.5B</p>


THE UNDERWRITE

Decompose the franchise into **four tests, each with a commit/pass trigger**. **(1) PME**: benchmark realized cash flows against a public small-cap index (Kaplan-Schoar PME and direct alpha) to separate manager skill from market beta. **(2) Value-creation attribution**: bridge gross value into revenue growth, margin expansion, multiple change, and deleveraging, deal by deal, to see what the GP actually does. **(3) Loss-ratio and dispersion**: realized loss rate and the spread between top- and bottom-quartile deals, to test repeatability versus a few heroes. **(4) Team and terms**: partner tenure, carry distribution to the next generation, GP commitment, and the LPA economics. A re-up clears all four or it is re-sized.

What drives the returns

FIGURE 2 · VALUE-CREATION ATTRIBUTION · REALIZED FUNDS II-III VS FUND V UNDERWRITE

SHARE OF GROSS VALUE CREATED · 100% STACKED



The commit case requires operational value (revenue + margin = 65%) to replace multiple expansion (40% → 20%). If the GP cannot evidence the path, the re-up is re-sized, not approved.

Prior realized returns lean on multiple expansion (40%); the Fund V underwrite only works if operational levers carry two-thirds of value. This is the single most important diligence finding and the hinge of the recommendation.

WORKED EXAMPLE · SIMULATED FINDING

Against a Russell 2000 benchmark, realized Funds II-III post a **KS-PME of 1.32** and **direct alpha of +550 bps**: real skill, but narrowing from Fund II (PME 1.45) to Fund III (PME 1.22). Realized loss ratio is **8% of invested capital** below 1.0x (healthy for the strategy), and top-quartile-deal concentration is moderate: the top three deals drive 45% of gross profit, not 70%+. Read together: a genuine, repeatable lower-mid-market engine whose edge is compressing as size rises. Commit, but at a size that respects the trend rather than extrapolating Fund II.

Diligence plan (6-week underwrite)

WEEK	WORKSTREAM	FORCING-FUNCTION DELIVERABLE	COMMIT/PASS TRIGGER
W1-2 Quant base	Rebuild deal-by-deal cash flows from the data room; compute PME, direct alpha, loss ratio, gross-to-net bridge	Independent track-record model (not the GP's)	PME > 1.15 and direct alpha > +300 bps on realized
W2-3 Attribution	Decompose each realized deal into revenue, margin, multiple, leverage; map to named operating playbooks	Value-creation bridge by deal and by fund	Operational share (rev + margin) > 50% and rising
W3-4 References	10-12 calls: management teams (incl. one loss), former LPs, lenders, departed partners	Reference log with attributed quotes	No unresolved integrity or team-stability flag
W5-6 Terms · IC	LPA and economics review; GP commit, key-person, succession, fee/carry, co-invest rights; synthesis	Commit/pass memo + side-letter ask list	IC votes on a sized recommendation

The plan is gated: each phase carries a commit/pass trigger, so a re-up that fails the PME floor or the operational-attribution test is stopped before reference and terms work begins, defended by a model the LP built rather than the GP's marketing.

Underwriting scorecard

DIMENSION	METRIC	UNDERWRITE	WHY IT MATTERS
Skill vs beta	KS-PME / direct alpha (realized)	1.32 / +550 bps	Separates manager skill from a rising market
Composition	Operational share of value	Target > 50% (was 45%)	Multiple expansion cannot be relied on in this rate regime
Repeatability	Realized loss ratio	8% of capital < 1.0x	Tests downside discipline, not just winners
Concentration	Top-3-deal share of gross profit	45%	Above 70% the record is a few heroes
Alignment	GP commitment to fund	3% (above 2% norm)	Skin in the game on the same terms as the LP
Continuity	Senior-partner tenure / carry to next gen	Avg 11 yrs / 35% carry	Succession risk is the quiet killer of re-ups

Risks & mitigations

- HIGH** **Fund-size decay and style drift: a 56% step-up pushes the GP up-market into larger, more efficiently-priced auctions where its lower-mid-market sourcing edge does not travel.**
Mitigation: size the commitment to the trend, not Fund II; require a written deal-size and deal-count plan; condition on a portfolio-construction covenant (max % in deals above a stated EV) and quarterly deployment reporting against it.
- HIGH** **Return composition reverses: with 40% of prior value from multiple expansion, a flat-to-lower exit-multiple regime alone can pull Fund V from 2.0x toward 1.5x net.**
Mitigation: underwrite to a one-turn multiple contraction in the base case; require evidence of named operational playbooks (procurement, pricing, bolt-on integration) with realized EBITDA-growth attribution before committing.
- MED** **Succession and key-person: returns are associated with two founding partners; carry distribution to the next generation is thinner than the franchise narrative implies.**
Mitigation: key-person and no-fault-divorce provisions in a side letter; review the carry-vesting schedule for the rising partners; stage the relationship (commit to Fund V, reassess Fund VI on succession evidence).

30 / 60 / 90: post-commitment

QUARTER 1	YEAR 1	RE-UP
Close & onboard <ul style="list-style-type: none"> › Side letter executed (key-person, MFN, co-invest rights) › Capital-call schedule and pacing modeled › Baseline KPIs loaded to the monitoring system 	Monitor & co-invest <ul style="list-style-type: none"> › Deployment tracked against the portfolio-construction covenant › Co-investment pipeline screened (fee/carry-free) › Quarterly value-creation attribution refreshed 	Fund VI decision <ul style="list-style-type: none"> › Realized DPI vs the J-curve underwrite › Operational-share-of-value trend confirmed › Succession evidence reviewed before re-commit

DECISION ASKED

Approve a **\$75M commitment to Fund V**, conditional on a side letter securing **key-person and no-fault-divorce provisions, MFN, fee/carry-free co-investment rights, and a portfolio-construction covenant** capping exposure to above-mandate deal sizes. **Pass** if the GP declines the key-person or co-invest terms, or if the W2-3 attribution cannot evidence operational value above 50%. **Success:** the IC votes on a sized recommendation backed by an independent track-record model, not the GP's marked TVPI.

Sources: Bain & Co., *Global Private Equity Report (2025)*; Harris, Jenkinson, Kaplan & Stucke, "Has Persistence Persisted in Private Equity?" (*Journal of Corporate Finance*, 2023); Braun et al. (*Journal of Finance*, 2026) and Rossi (2019) on fund size and reversion-to-the-mean; Cambridge Associates / Burgiss benchmarks; Kaplan-Schoar PME (2005).