

WORKWISE SOLUTIONS

WHITE PAPER

The AI-Native Fund

Organizational Design After the Analyst Pyramid

Dr. Leigh Coney

Founder & Principal Consultant, WorkWise Solutions

Q2 2026

workwisesolutions.org

leigh@workwisesolutions.org

Abstract

The Skill Erosion Paradox (Q1 2026) examined what AI-heavy workflows do to the formation of analytical judgment. It was, by design, a defensive paper. This one is constructive. It answers the question that paper left open: if the analyst pyramid no longer makes economic sense, what should a deal team actually look like?

The argument runs in three parts. First, the pyramid was an information-processing architecture rather than a training scheme, and AI has now absorbed the processing layer that justified its wide base. The evidence is no longer speculative; it is visible in hiring data across banking, accounting, and private equity itself. Second, the obvious response, keeping the old shape and thinning the bottom, fails on its own terms. It saves the cost of juniors while destroying the mechanism that produced seniors. Third, a workable replacement exists, and this paper specifies it: a column-shaped team built around an agent core, with concrete staffing ratios, a redesigned associate pipeline that treats apprenticeship as a product rather than a byproduct, and a compensation model that converts deleted seats into a judgment premium, deeper carry participation, and an embedded engineering seat.

The paper is written for managing partners, chief operating officers, heads of talent, and the deal leaders who will run these teams. It draws on current labor market research, fund compensation data, and the early record of AI-native investment firms.

Keywords

AI-native private equity, deal team structure, analyst pyramid, staffing ratios, associate pipeline, apprenticeship, carried interest, compensation design, agentic AI, organizational design, talent strategy, deal engineer, verification, deliberate development

Contents

1. The Pyramid Was a Technology
2. The Evidence: The Base Is Already Shrinking
3. Why “Fewer Juniors” Is Not an Operating Model
4. The Shape That Replaces It
5. Rebuilding the Associate Pipeline
6. The New Compensation Economics
7. The Transition: Three Moves in Twenty-Four Months
8. Conclusion: The Fund That Trains on Purpose

References

About the Author

1. The Pyramid Was a Technology

Put a buyout firm's 2019 org chart next to one from 1989 and the shape is identical. A few partners at the top, a band of principals and vice presidents beneath them, and a wide base of associates and analysts doing the reading, the modeling, and the memo drafting. Banking, consulting, and law share the same silhouette.

That shape did not persist out of sentiment. It persisted because it solved a real problem: turning raw information into decision-ready analysis required enormous quantities of structured human labor. David Maister, in his classic study of professional service firms, described the economics plainly. Profitability rests on leverage, the ratio of junior fee-earners to the senior professionals whose judgment clients actually buy (Maister, 1993). Private equity imported the structure wholesale from its feeder industries, swapping billable hours for deal capacity.

The pyramid produced two outputs at once. The explicit product was processed information: comp tables, three-statement models, CIM summaries, diligence memos. The byproduct was judgment. Thousands of hours at the base of the pyramid were how the firm manufactured the partners of 2040, one redlined model at a time.

AI has now collapsed the price of the first output. Goldman Sachs CEO David Solomon told an audience in January 2025 that AI can draft 95 percent of an IPO prospectus in minutes, work that previously occupied a six-person team for two weeks. "The last 5% now matters," he said, "because the rest is now a commodity" (Fortune, 2025). Researchers at Chicago Booth found that GPT-4, given anonymized financial statements, predicted the direction of future earnings more accurately than professional analysts, 60 percent against 53 (Kim, Muhn, & Nikolaev, 2024).

The benchmarks keep moving. OpenAI's GDPval evaluation, built from real work products across 44 occupations in the nine sectors that contribute most to US GDP, found the best frontier model matching or beating experienced professionals on 47.6 percent of blind comparisons in September 2025 (OpenAI, 2025a). By December the leading model's figure was 70.9 percent, and on OpenAI's companion benchmark of junior investment banking analyst tasks, three-statement models and leveraged buyouts, it scored 68 percent against its predecessor's 59 (OpenAI, 2025b).

When the explicit product becomes a commodity, the economic case for the wide base disappears. The byproduct, judgment formation, does not stop being necessary. It just stops being free.

That is the design problem this paper addresses. Not whether the pyramid survives; it will not. The question is what replaces it, and whether the replacement is designed or improvised.

2. The Evidence: The Base Is Already Shrinking

The contraction at the bottom of professional hierarchies is measurable, current, and concentrated exactly where AI capability overlaps with junior work.

Start with the broadest evidence. Stanford economists analyzing payroll records from ADP, covering millions of workers, found that early-career employees aged 22 to 25 in the most AI-exposed occupations experienced a 13 percent relative decline in employment after generative AI reached mass adoption, while older workers in the same occupations held steady or grew (Brynjolfsson, Chandar, & Chen, 2025). The

declines concentrate where AI automates work rather than assisting it. As Brynjolfsson summarized the finding, what younger workers know overlaps with what large language models can replace.

The pattern shows up first in the institutions that have always trained finance talent. The *New York Times* reported in 2024 that major banks, Goldman Sachs and Morgan Stanley among them, were weighing cuts to incoming junior analyst classes of as much as two-thirds as AI absorbs entry-level work (Lake, 2025). The Big Four accounting firms, the other great apprenticeship engine of commercial life, have cut graduate intake outright: KPMG's UK intake fell from 1,399 to 942, Deloitte cut 18 percent, EY 11 percent, and PwC 6 percent, while UK accountancy graduate job listings dropped 44 percent year on year (Koopman, 2025).

The wider labor market confirms the squeeze. The New York Fed's tracker put unemployment for recent college graduates at 5.7 percent in early 2026, with underemployment above 41 percent, and the group's longstanding edge over the broader workforce has inverted (Federal Reserve Bank of New York, 2026). Anthropic CEO Dario Amodei warned that AI could eliminate half of entry-level white-collar jobs within five years, naming finance and consulting among the most exposed fields (Axios, 2025).

Private equity is not watching from a distance. Deloitte's 2025 study of M&A organizations found that 86 percent had integrated generative AI into deal workflows, and among adopters, 88 percent of private equity firms had invested at least \$1 million in the technology specifically for their deal teams (Deloitte, 2025). Bain's 2026 Global Private Equity Report describes leading firms building competitive moats through scale, expertise, technology, and AI, even as management fee revenue comes under pressure (Bain & Company, 2026).

The early AI-native firms show where this goes. OffDeal, a startup calling itself the world's first AI-native investment bank, runs sell-side processes with roughly ten total staff, three bankers per auction, and agents handling buyer discovery, document drafting, and data room management; it reported ten deals closed or near completion in its first year, and one of its 25-year-old bankers earned a \$2 million bonus (Quilty-Harper, 2025). Clarum, a diligence startup, is compressing the 20 to 25 analyst hours of a thesis-fit screen into roughly three. EQT partner Chris Litchford frames the substitution plainly: where competitors hire analysts at \$85,000 a year to cold call and gather information, tooling that delivers most of that coverage can be used across the entire group (Quilty-Harper, 2025).

The people running the largest firms are saying the quiet part at conference volume. Vista Equity Partners CEO Robert Smith told the more than 5,500 attendees at SuperReturn 2025 that within a year 40 percent of them would have AI agents and the other 60 percent would be looking for work (Bloomberg News, 2025).

Treat any single claim with skepticism; conference predictions are not data. But the direction is consistent across payroll records, hiring decisions, and capability benchmarks. The base of the pyramid is already in liquidation.

3. Why “Fewer Juniors” Is Not an Operating Model

The instinctive response, visible across the industry right now, is to keep the pyramid's shape and thin its base. Same structure, fewer analysts, AI filling the space. That is a cost decision dressed as an operating model, and it fails in three ways.

The first failure is the supervision deficit, the central mechanism of *The Skill Erosion Paradox*. Every AI-embedded workflow depends on humans who can catch machine errors, and error-catching skill is built by

doing the underlying work, not by reviewing someone else's. Expertise research is unambiguous on the point: high-level judgment comes from deliberate practice, effortful repetition at the edge of capability with fast feedback (Ericsson, Krampe, & Tesch-Römer, 1993). A firm that deletes the practice deletes, on a lag, the supervision.

The second failure is developmental. Matt Beane, who spent a decade studying how skill actually transfers in settings from surgical suites to warehouses, found that mastery depends on a working bond between experts and novices built on challenge, complexity, and connection. AI deployed for pure productivity severs exactly that bond, by routing the novice's learning material to the machine (Beane, 2024). The legal industry has reached the same diagnosis: in Thomson Reuters' 2025 Future of Professionals survey, 80 percent of law firm respondents expect AI to fundamentally alter how they conduct business, specifically how firms price, staff, and deliver legal work (Thomson Reuters Institute, 2025).

The third failure is arithmetic. The partners of 2036 enter the industry between now and 2028, or they do not enter it at all. Halting junior development does not remove the cost of judgment. The cost converts into an unfunded liability that comes due at succession, payable at whatever the market then charges for proven investment judgment, and trained supply is contracting upstream at the same moment.

That upstream contraction deserves its own line item. Private equity has historically outsourced foundational training to the banks' analyst programs: two years of reps, paid for by someone else, delivered through on-cycle recruiting. The programs are shrinking, and the recruiting machinery itself has wobbled. Apollo paused its 2027 associate hiring outright, with General Atlantic, TPG, and KKR following within days, after JPMorgan moved to terminate analysts who accept future-dated buy-side offers (Jensen Partners, 2025). The conveyor belt that let funds treat talent as a just-in-time input is no longer reliable at either end.

So the answer to the pyramid is a different shape, with the training function moved from byproduct to product. The rest of this paper specifies that shape.

4. The Shape That Replaces It

4.1 From pyramid to obelisk

Deal work separates into three layers. Production: reading, extraction, modeling, drafting, screening. Verification and calibration: checking machine output against sources, stress-testing assumptions, knowing where models fail. Judgment and relationships: investment selection, negotiation, management assessment, and the decision to walk away.

The pyramid assigned production to a wide human base. The AI-native fund assigns production to an agent core and staffs humans against the two layers machines do not hold. The resulting silhouette is a column rather than a triangle. Call it the obelisk: nearly straight sides, a slight taper, a hard mineral core running underneath.

Figure 1 shows the shift in concrete terms for a representative mid-market firm running two deal squads on a \$1.0 to \$2.5 billion fund. The legacy configuration carries sixteen investment professionals at roughly three junior seats per partner. The AI-native configuration carries eleven: the same senior bench, three associates instead of nine, one embedded deal engineer, and an agent core doing the production work of, conservatively, six to eight analysts.




Figure 1. Deal team structure for a representative mid-market fund. Headcounts are WorkWise design targets for a \$1.0 to \$2.5 billion fund running two concurrent deal squads, not survey data.

These ratios are design targets and they flex with strategy. A sourcing-intensive growth fund keeps more humans in origination. A complex carve-out shop keeps more in diligence. The invariant is the direction: the base compresses by half or more, the middle holds, and a technical seat appears inside the deal team rather than beside it.

4.2 Four seats, redefined

The associate becomes a verification owner

The job shifts from producing the first draft to interrogating machine drafts against primary sources, owning provenance, and finding the errors that matter. This is a harder job than the one it replaces, and it carries explicit, measured standards: error-detection rates on seeded tests, source verification depth, and the unassisted capability scores laid out in *The Skill Erosion Paradox*. The role also keeps protected manual production time, because verification skill decays without it.

The vice president becomes the orchestration layer

VPs run the agent fleet at deal level: what gets automated, what gets verified by whom, and what reaches the investment committee. The VP's leverage was once three analysts; it is now a production system plus the associates who check it. This is the seat where deal velocity is won or lost.

The partner's job barely changes, and expands

Judgment, relationships, conviction, and the willingness to kill a deal late remain the product LPs pay for. What changes is span. With compressed layers below, a partner carries more concurrent situations, and partner attention, rather than analyst hours, becomes the firm's binding constraint.

The deal engineer is the new seat

Someone has to build agent workflows, maintain evaluation harnesses, and keep the data plumbing honest, and that person belongs inside the deal team, attending Monday pipeline meetings, rather than in a separate innovation group. The biggest firms have staffed a version of this seat for years; EQT's long-running Motherbrain platform is the example other dealmakers point to as proof of how heavily the largest

players are already investing (Quilty-Harper, 2025). One engineer per two to three squads is a workable starting allocation.

4.3 The new constraint

Under the pyramid, deal capacity was rationed by analyst hours. Screening a hundred targets deeply meant something like two thousand junior hours at 20 to 25 hours per screen, so funnels stayed narrow and triage stayed crude. With screening compressed by a factor of five or more, the funnel widens at near-zero marginal cost, and the constraint migrates up the team: partner attention, investment committee time, verification hours.

This inverts the management problem. The question stops being “how do we feed the analysts” and becomes “what deserves judgment.” The metric that matters in the obelisk is judgment hours per closed deal, and the organizational design exists to spend those hours only where they change outcomes.

5. Rebuilding the Associate Pipeline

5.1 Hire for the second job

When the first job was modeling and memo production, firms screened for process stamina and technical polish. In the obelisk, the first job is verification and the second is judgment, so selection has to predict different things: reasoning under ambiguity, adversarial instinct, and the ability to articulate why an analysis is wrong rather than reproduce one that is right. Tool fluency still matters and should be assumed; it is table stakes, not a differentiator. Interviews should include unassisted work, because AI-assisted polish now masks exactly the variance a hiring process exists to detect.

Cohorts shrink. A firm that ran nine associates across two squads hires three or four, which roughly triples selectivity at constant applicant volume. The breakdown of on-cycle recruiting helps here rather than hurts: with the two-year conveyor gone, firms can recruit against demonstrated judgment instead of against a stopwatch (Jensen Partners, 2025).

5.2 Apprenticeship as product, not byproduct

The pyramid trained people by accident, as exhaust from production work. The obelisk produces no exhaust, so development has to be engineered into the job. *The Skill Erosion Paradox* supplied the components: graduated AI allocations by seniority, structured manual repetitions on live deals, error archaeology on past transactions, and capability assessment that measures the human rather than the human-plus-machine.

Beane’s three components translate directly into program design (Beane, 2024). Challenge: every associate owns first-pass production on a defined slice of live work each quarter, manually. Complexity: rotations across sourcing, diligence, and portfolio work, because pattern libraries do not form inside a single workflow. Connection: a named senior owner for each associate, whose own scorecard includes that associate’s measured progress.

Promotion gates replace tenure clocks. Time-based advancement made sense when every associate accumulated roughly the same repetitions per year. AI makes accumulated experience wildly uneven, so

the firm must test capability directly, and the capability dashboard from the previous paper stops being a defensive instrument and becomes the promotion system. Table 2 shows the structure.

Gate	Evidence required	Assessment source
Error detection	Sustained detection rate above 80 percent on seeded-error reviews	Monthly seeded exercises
Unassisted production	Build a complete model and memo from primary sources at IC quality	Quarterly unassisted exercises
Adversarial judgment	Bear cases that surface risks the agent stack did not flag	Semi-annual bear case reviews
Verification discipline	Documented source provenance on all AI-assisted work product	Random work audits

Table 2. Competency gates for advancement from associate to vice president. Thresholds calibrate to firm baselines established in the first assessment cycle.

5.3 The end of the free ride

For thirty years, private equity let investment banks pay for basic training, then bought the survivors. Section 2 described why that subsidy is ending: the banks are cutting the very programs that produced pre-trained associates. The associate pool of 2028 will arrive with fewer manual repetitions than any cohort in the industry's history.

Funds therefore need an internal foundation program: modeling from primary filings, full-cycle deal simulations, and supervised verification work with seeded errors. This is a real cost, on the order of a partial seat per cohort. It is also a recruiting weapon. In a market where most employers offer juniors a future of supervising machine output, the firm that credibly manufactures judgment will take its pick of a generation that has noticed what is happening to its prospects (Federal Reserve Bank of New York, 2026).

6. The New Compensation Economics

6.1 What seats cost today

Heidrick & Struggles' 2025 survey of 656 North American private equity investment professionals gives the baseline. At funds of \$1.0 to \$1.5 billion, mean total cash compensation for 2024 ran roughly \$244,000 for associates and senior associates, \$378,000 for vice presidents, \$689,000 for principals, and \$1.11 million for partners and managing directors (Heidrick & Struggles, 2025). Carry concentrates sharply above the associate band. Vice presidents at these funds averaged about \$2.4 million in carry dollars at work, the survey's measure of expected carry value at a net 2x fund outcome; where associates hold carry at all, the stakes are becoming real, averaging \$1.9 million among those reporting it.

Fully loaded, an associate seat costs \$300,000 to \$350,000 a year. The agent infrastructure replacing most of that seat's production output, with model usage, tooling, and incremental data included, runs at a small fraction of one seat even when priced generously.

6.2 Where the deleted seats go

Table 1 works the arithmetic for the representative firm in Figure 1, holding the senior bench constant and using the Heidrick means as anchors. The legacy configuration spends about \$7.7 million in cash compensation across sixteen professionals. The AI-native configuration spends about \$7.1 million across eleven, while paying associates a quarter more, funding a deal engineer, and carrying a generous agent infrastructure budget.

Role	Legacy seats	Legacy cost	Obelisk seats	Obelisk cost
Partners	3	3,330	3	3,330
Principals and vice presidents	4	2,140	4	2,140
Associates (verification owners)	9	2,205	3	915
Deal engineer	0	0	1	375
Agent infrastructure (models, tooling, data)	0	0	n/a	350
Total annual cash cost	16	7,675	11	7,110

Table 1. Illustrative annual cash cost of a two-squad investment team at a \$1.0 to \$1.5 billion fund, in USD thousands. Per-seat anchors round the 2024 mean total cash compensation figures from Heidrick & Struggles (2025); obelisk associates carry a 25 percent premium. Figures are design illustrations, not survey results.

Three things in this table deserve emphasis. The saving is modest by design: the obelisk exists to reallocate roughly the same spend toward judgment, verification, and engineering, not to shrink the bill. Capacity rises even as headcount falls, because the agent core does the production work of six to eight analysts. And the associate premium prices a harder job, buys retention in a structure where every departure removes a third of the junior bench rather than a ninth, and pays for measured verification responsibility.

6.3 Carry moves down

Fewer participants make deeper participation affordable. Extending real carry to associates, on schedules that typically take four to seven years to vest fully (Heidrick & Struggles, 2025), does three jobs at once. It aligns the people who own verification with fund outcomes, which is where earlier papers in this series argued accountability for AI-assisted analysis must ultimately sit. It supports the longer, gated development arc of Section 5, since carry vests on roughly the same clock as judgment forms. And it replaces the up-or-out tournament, which stops working once cohorts are small and regretted attrition is expensive.

6.4 Why the timing is forced

Bain's arithmetic for the current cycle is blunt. Deals that once cleared a 2.5x return with 5 percent annual EBITDA growth now need 10 to 12 percent, while fee income compresses and the cost of differentiated capability rises; the firm calls it "12 is the new 5" (Bain & Company, 2026). That rule of thumb is, among other things, a labor statement. A fund cannot pay for double-digit operational value creation with a cost structure built around manual information processing, and it cannot generate that value with a team whose judgment pipeline has eroded. Firms that reallocate first will compound the advantage, because team redesign has a learning curve and the learning is proprietary.

7. The Transition: Three Moves in Twenty-Four Months

Move one: freeze the shape and instrument it (months 1 to 6)

Hire only at attrition. Baseline every professional's unassisted capability using the assessment structure from *The Skill Erosion Paradox*, and instrument the deal funnel so the firm knows its actual hours by stage. Most firms discover they are already running a ragged, unmanaged version of the obelisk, with agent tools doing analyst work informally and nobody owning verification.

Move two: run one obelisk squad (months 6 to 12)

Convert a single deal team to the target design: full agent stack, one deal engineer, three verification-owner associates with gated development plans. Measure it against the legacy squads on cycle time, investment committee kill rate, error escapes, and judgment hours per deal. The pilot exists to surface firm-specific failure modes before they scale, and to give skeptical partners a comparison they cannot argue with.

Move three: reprice and re-gate (months 12 to 24)

Roll the structure across squads as attrition permits, and avoid layoffs; firing juniors to buy software damages LP confidence and campus credibility for years. Publish the new career arc to candidates: smaller intake, higher pay, carry participation, promotion on demonstrated capability. Report the design to LPs proactively, with capability metrics as evidence that the judgment they are underwriting still exists in the building.

Twenty-four months is deliberate. Moving faster breaks the apprenticeship machinery this design exists to protect. Moving slower concedes the learning curve to competitors who started earlier.

8. Conclusion: The Fund That Trains on Purpose

The analyst pyramid was a technology for converting junior labor into processed information, with judgment as its byproduct. AI has made the conversion uneconomic, and the industry is responding the way industries respond, by cutting the input it no longer needs. The error is in believing that cut is the whole job.

The fund that emerges well from this transition is smaller at the base, senior in its center of gravity, engineered at its core, and deliberate about the one thing the pyramid produced that machines do not: people who know when the machine is wrong. Its associates are fewer, better paid, carried in the fund, and measurably developing. Its partners spend their hours on judgment rather than on supervising production. Its capacity rises while its headcount falls, and its succession risk is funded rather than deferred.

The earlier papers in this series dealt with verification, complacency, and capability. This one has dealt with structure, because structure is where those battles are won before they are fought. Firms that design the obelisk now, while the pyramid still staffs their competitors' deals, will own both the cost curve and the judgment supply of the next decade.

The pyramid trained people by accident. The AI-native fund will train them on purpose, or not at all.

References

- Axios. (2025, May 28). *Behind the curtain: Top AI CEO foresees white-collar bloodbath*. Axios.
- Bain & Company. (2026). *Global private equity report 2026: Welcome to a new era*. Bain & Company.
- Beane, M. (2024). *The skill code: How to save human ability in an age of intelligent machines*. Harper Business.
- Bloomberg News. (2025, June 5). *Vista CEO says AI to force 60% of SuperReturn crowd to seek work*. Bloomberg.
- Brynjolfsson, E., Chandar, B., & Chen, R. (2025). *Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence*. Stanford Digital Economy Lab working paper.
- Deloitte. (2025). *2025 M&A generative AI study*. Deloitte Development LLC.
- Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100(3), 363–406.
- Federal Reserve Bank of New York. (2026). *The labor market for recent college graduates*. Federal Reserve Bank of New York.
- Fortune. (2025, January 17). *Goldman Sachs CEO says that AI can draft 95% of an IPO prospectus in minutes*. Fortune Media.
- Heidrick & Struggles. (2025). *2025 North America private equity investment professional compensation survey*. Heidrick & Struggles International.
- Jensen Partners. (2025, July 23). *Private equity's early recruiting machine just broke, and everyone's watching what happens next*. Jensen Partners.
- Kim, A. G., Muhn, M., & Nikolaev, V. V. (2024). *Financial statement analysis with large language models*. Working paper, University of Chicago Booth School of Business.
- Koopman, S. (2025, June 23). *Big Four slash graduate jobs as AI takes on entry level work*. City AM.
- Lake, S. (2025, June 2). *Junior analysts, beware: Your coveted and cushy entry-level Wall Street jobs may soon be eliminated by AI*. Fortune.
- Maister, D. H. (1993). *Managing the professional service firm*. Free Press.
- OpenAI. (2025a). *GDPval: Evaluating AI model performance on real-world economically valuable tasks*. OpenAI.
- OpenAI. (2025b, December 11). *Introducing GPT-5.2*. OpenAI.
- Quilty-Harper, C. (2025, October 24). *AI promises to make private equity faster as competition heats up*. ThinQ by EQT.
- Thomson Reuters Institute. (2025). *2025 future of professionals report*. Thomson Reuters.

About the Author

Dr. Leigh Coney is the Founder and Principal Consultant of WorkWise Solutions. With a PhD in Organisational Psychology, Dr. Coney has spent over a decade at the intersection of AI, behavioural science, and organisational design. His research focuses on decision making in high stakes environments, with particular attention to why sophisticated AI systems fail to achieve adoption and how governance and deployment systems can be designed to account for human cognitive and social dynamics.

This paper is part of an ongoing research series on responsible AI adoption in financial services. Previous publications include *The Skill Erosion Paradox*, *Combating Automation Complacency*, and *Closing the Accountability Gap*.

About WorkWise Solutions

WorkWise Solutions specializes in psychology-driven AI implementations for Private Equity, Venture Capital, and Strategic Consulting firms. Our organizational design and deliberate development programs ensure that AI adoption strengthens, rather than erodes, the human judgment that defines investment excellence.

For advisory inquiries: leigh@workwisesolutions.org