WORKWISE SOLUTIONS

WHITE PAPER

# Helping Portfolio Companies Deploy AI

*A General Partner's Operating Partner Playbook*

**Dr. Leigh Coney**

Founder & Principal Consultant, WorkWise Solutions

Q2 2026

workwisesolutions.org

Contact@workwisesolutions.org

# Abstract

Private equity firms have largely answered the question of whether to invest in artificial intelligence at the firm level. The harder, less discussed problem is how to help portfolio companies deploy AI in ways that create durable enterprise value before exit. The latest industry data tells two stories at once. Portfolio companies are widely experimenting with generative AI, with the majority engaged in some form of testing or pilot work. Yet only a small minority have operationalised use cases that move EBITDA in a measurable way. The Massachusetts Institute of Technology NANDA initiative found that 95 percent of enterprise generative AI pilots fail to deliver any measurable profit and loss impact, even after $30 to $40 billion of corporate investment in 2025.

This paper introduces the Portfolio AI Deployment Framework (PADF), a structured playbook for the operating partners and value creation teams responsible for translating firm level AI conviction into portfolio level results. The framework rests on three pillars: selection discipline (choosing the right portfolio companies and use cases), implementation architecture (deploying AI in ways that survive contact with operating reality), and cross portfolio compounding (turning a single successful deployment into a repeatable playbook).

The paper is intended for general partners, operating partners, AI operating partners, value creation team leaders, and the chief executives and chief operating officers of portfolio companies who are accountable for AI outcomes. It addresses one central question: given that two thirds of operating partners are already stretched across five or more portfolio companies, how should a sponsor concentrate AI effort to maximise value creation across the portfolio?



## Keywords

AI operating partner, value creation, portfolio company AI deployment, EBITDA expansion, change management, generative AI implementation, private equity AI playbook, AI vendor selection, hold period strategy, exit preparation, AI governance, portfolio operations, KPMG PE Value Creation Survey, MIT GenAI Divide, Bain Global Private Equity Report

# Table of Contents

## 1. Introduction: Why Portfolio AI Is the Hardest Conversation in PE Right Now

The conversation about artificial intelligence in private equity has changed quickly over the past eighteen months. Two years ago, the question at most general partner meetings was whether to invest in AI at all. Today, the question is no longer whether, but where. Deloitte's 2025 GenAI in M&A Study, based on 1,000 senior corporate and private equity leaders, found that 86 percent of respondents had integrated generative AI into their M&A workflows. Among the private equity adopters specifically, 88 percent had invested at least $1 million in the technology for their deal teams, and 81 percent expected measurable return on that investment within one to three years (Deloitte, 2025).

The firm level conversation has matured. The portfolio level conversation has not. According to Bain & Company's Global Private Equity Report 2025, the majority of portfolio companies are in some phase of generative AI testing and development, but only roughly 20 percent have operationalised use cases that produce concrete results (Bain & Company, 2025). The Boston Consulting Group reaches a similar conclusion in its 2026 analysis of AI in private capital: most private equity firms cannot show meaningful returns from AI in many of their portfolio companies (BCG, 2026).

This is not a technology problem. The same large language models, vector databases, and orchestration platforms are available to a portfolio company in Cleveland and a large technology firm in Mountain View. The difference is operational: how the project is selected, who owns it, how it is staffed, how it is integrated with the workflow, and how the gains are measured. The Massachusetts Institute of Technology NANDA research finds that 95 percent of enterprise generative AI pilots fail to deliver any measurable profit and loss impact, despite an estimated $30 to $40 billion in corporate investment in 2025 (MIT NANDA, 2025). Most of those failures are not failures of model quality. They are failures of deployment.

For private equity sponsors, the cost of these failures compounds. A failed AI project at a portfolio company is not merely a wasted investment. It consumes scarce operating partner capacity, hardens portfolio company management against future initiatives, and risks producing the worst possible exit narrative: a company that tried AI and could not make it work. With holding periods near seven years and global buyout dry powder at roughly $1.3 trillion (Bain & Company, 2026), sponsors cannot afford a portfolio where AI initiatives fail at typical enterprise rates.

This paper sets out a framework, the Portfolio AI Deployment Framework (PADF), designed for the operating partners and value creation teams who carry day to day responsibility for portfolio AI outcomes. It does not claim that AI value creation is easy. It claims that the failures are systematic, the patterns are visible in the recent industry data, and the operating partner is uniquely positioned to interrupt them.

## 2. The State of Portfolio Company AI in 2026

Three patterns characterise the portfolio AI landscape in 2026.

First, almost every portfolio company is doing something. Industry surveys put adoption at near saturation. Bain's survey of investors representing $3.2 trillion in assets under management found that the majority of portfolio companies are in some phase of generative AI testing or pilot work (Bain & Company, 2025). KPMG's Q3 2025 AI Quarterly Pulse Survey, which covers organisations across banking, asset management and private equity, and technology, found that the share of organisations having deployed at least some AI agents nearly quadrupled in 2025, reaching 42 percent, up from 11 percent two quarters earlier (KPMG, 2025).

Second, almost none of these efforts are reaching the profit and loss statement. The MIT NANDA research, based on 52 executive interviews, surveys of 153 leaders, and analysis of 300 public AI deployments, concluded that only 5 percent of integrated AI pilots are extracting millions in value, while the vast majority remain stuck without measurable P&L impact. The MIT team named the resulting gap the GenAI Divide (MIT NANDA, 2025).

Third, the operating partner role is being asked to absorb this complexity without proportionate growth in capacity. KPMG's 2025 Global PE Value Creation Survey, based on 500 PE leader interviews, found that just 18 percent of respondents were operating partners, and two thirds of those were stretched across five or more portfolio companies (KPMG, 2025). The same report concludes that operational expertise headcount will need to roughly triple if firms are to capture today's value creation opportunities.

The combined picture is clear: high activity, low yield, and thin operating bandwidth. Into this gap, a new role has emerged. Korn Ferry's 2025 research documents the rise of the AI operating partner, a specialised version of the technology operating partner role focused specifically on generative AI value creation across the portfolio (Dube & Nejdawi, 2025). Some firms have made these full time hires; others use part time advisors. AlixPartners, in its 2025 guidance for operating partners, notes that operating partners are uniquely positioned to spot, vet, and scale AI use cases, but cautions that AI is a new discipline for many of them (AlixPartners, 2025).

This paper accepts the premise that the operating partner sits at the centre of portfolio AI value creation. It then asks the harder question: what should that partner actually do?

# 3. The Operating Partner's Unique Vantage Point

Three structural advantages position the operating partner well for portfolio AI work, and one structural disadvantage frames the entire role.

## 3.1 Cross Portfolio Visibility

A portfolio company chief executive sees one company. The operating partner sees ten or twenty. When a single asset succeeds with a particular vendor or use case, the operating partner is the only person in the system who can recognise that success in time to replicate it elsewhere. Bain's report describes this dynamic in the example of Vista Equity Partners, which runs annual hackathons across its portfolio in the United States and India, with successful prototypes scaling into revenue generating products at multiple companies (Bain & Company, 2025).

## 3.2 Selection Authority

Portfolio company management teams will instinctively pursue the AI projects that excite them, which are not always the projects with the strongest EBITDA case. The operating partner, sitting outside the operational fray, can apply portfolio level capital allocation discipline. FTI Consulting's analysis finds that EBITDA gains of 5 to 25 percent are realistic for well targeted AI deployments across industries, but warns that returns vary sharply with use case selection (FTI Consulting, 2025).

## 3.3 Vendor Leverage

A single portfolio company negotiating with an AI vendor is one customer. A sponsor negotiating on behalf of fifteen portfolio companies is a strategic account. This leverage matters more than it used to, because the MIT data shows that AI tools built by external vendors succeed roughly twice as often as internal builds (67 percent versus 33 percent) (MIT NANDA, 2025). If buying is the higher yield path, then the firm level vendor relationship becomes a real source of value.

## 3.4 The Bandwidth Constraint

The structural disadvantage is bandwidth. With two thirds of operating partners stretched across five or more portfolio companies (KPMG, 2025), the operating partner cannot personally drive AI projects in every asset simultaneously. The role is necessarily one of selection, sponsorship, and oversight, not direct execution. The framework that follows is built around this constraint.

# 4. Why Most Portfolio AI Deployments Fail

The MIT NANDA research identifies four structural factors behind the 95 percent failure rate of enterprise generative AI pilots (MIT NANDA, 2025). All four are visible in portfolio company deployments, and three are amplified by the specifics of the private equity operating context.

## 4.1 Brittle Workflow Integration

Pilots are typically designed as technology demonstrations rather than production systems. They work in clean test conditions and break under the noise, exceptions, and edge cases of real operations. In a portfolio company, this problem is intensified by lean technology functions and a limited tolerance for disruption.

## 4.2 Misalignment with Day to Day Operations

AI projects are often initiated by the chief executive or chief technology officer based on a perceived strategic priority, then handed to functional teams who do not see how it changes their daily work. BCG's analysis emphasises that mature use of horizontal AI tools requires investment in change management across the end to end workflow to ensure that AI tools are properly leveraged (BCG, 2025). Without that investment, adoption stalls.

## 4.3 Investment Bias Toward Sales and Marketing

The MIT research finds that AI budgets overwhelmingly favour sales and marketing applications, despite back office and operations producing higher returns (MIT NANDA, 2025). In portfolio companies, this same bias appears: revenue side AI projects feel more strategic and exciting, while finance, procurement, and operations projects produce more measurable EBITDA impact in shorter timeframes.

## 4.4 The Build Versus Buy Gap

The fourth factor, and the one that bites hardest in portfolio companies, is the build versus buy distinction. The MIT research finds that vendor built and partnership led AI projects succeed roughly 67 percent of the time, while internal builds succeed only 33 percent of the time (MIT NANDA, 2025). Yet many portfolio companies, particularly those with proud technology cultures, instinctively want to build their own. This instinct is sometimes correct. More often, it is the path to the failure column of the MIT statistics.

The Portfolio AI Deployment Framework set out in the next section is designed to interrupt these failure patterns systematically.

## 5. The Portfolio AI Deployment Framework (PADF)

The PADF is a three pillar framework for general partners and operating partners responsible for portfolio AI outcomes. It rests on a single premise: portfolio AI value creation is fundamentally a portfolio management problem, not a technology problem. Capital, attention, and operating bandwidth must be allocated as deliberately as investment capital.

| Pillar | Core Problem It Addresses | Design Principle |
|---|---|---|
| **I. Selection Discipline** | Portfolios are over invested in many shallow pilots and under invested in a few deep deployments. | Concentrate operating bandwidth on use cases and assets where AI fitness, reversibility, and hold period align. |
| **II. Implementation Architecture** | Most pilots are technology demonstrations that fail to survive contact with real operations. | Default to vendor or partner solutions, anchor in EBITDA, name a non CEO owner, and treat change management and data debt as first class workstreams. |
| **III. Cross Portfolio Compounding** | Each portfolio company effectively starts from zero, repeating the same mistakes. | Treat the first deployment as a portfolio asset. Negotiate vendors at the firm level. Convene a portfolio AI council and maintain an honest failure registry. |

Pillar I (Selection Discipline) addresses the question of where to deploy AI investment. Most portfolios are over invested in many shallow pilots and under invested in a few deep deployments. The pillar provides a method for narrowing.

Pillar II (Implementation Architecture) addresses how a deployment should be structured to survive contact with operating reality. This pillar draws heavily on the MIT NANDA findings on vendor selection, change management, and integration.

Pillar III (Cross Portfolio Compounding) addresses how a single success becomes a portfolio wide playbook. Without deliberate compounding, every portfolio company effectively starts from zero. With it, the marginal cost of the second, fifth, and tenth deployment falls sharply.

The framework consists of thirteen design principles distributed across the three pillars. They are intended to be used as a checklist by operating partners during portfolio AI planning, not as a rigid sequence.

## 6. Pillar I: Selection Discipline

Selection discipline answers a question that most portfolios avoid: which AI initiatives will the operating team actually back, and which will it deprioritise? Without an honest answer, scarce operating bandwidth gets spread thinly across many initiatives, none of which receives the support it needs to succeed.

### Principle 1: Use Case Fitness Over Technology Fitness

The first decision is not which AI tool to deploy. It is which business problem deserves AI attention. The strongest signals of use case fitness are: a workflow that is already heavily structured (so AI can be measured against a clear baseline), a repeatable task volume (so productivity gains compound), and a pre existing data substrate (so the project does not need to begin with a six month data cleanup). Use cases meeting all three criteria should be prioritised over use cases that pattern match to a fashionable AI demo but lack underlying fitness.

### Principle 2: Reversibility Weighted Prioritisation

Not every AI failure is equally costly. A failed AI experiment in marketing copy generation is recoverable in days. A failed AI deployment in pricing or credit decisioning can damage customer trust for quarters. Operating partners should weight expected EBITDA impact against the reversibility of failure. The ideal first deployment is high in expected impact and high in reversibility. Low impact and low reversibility projects should be deprioritised regardless of how interesting they sound.

### Principle 3: Hold Period Gating

A two year transformation does not work in a portfolio company that is eighteen months from exit. A three month productivity pilot may not be worth the change management cost in a company with a seven year remaining hold. Operating partners should explicitly map AI initiatives against remaining hold period and prioritise accordingly. Short hold companies should focus on pre built AI features in existing systems and on exit narrative initiatives. Long hold companies are appropriate for deeper transformations.

### Principle 4: Portfolio Company Readiness Scoring

Not every company in the portfolio is ready for the same kind of AI investment. Readiness has at least four dimensions: data infrastructure (is there reliable, accessible operational data?), leadership posture (does the chief executive view AI as opportunity or threat?), technology function maturity (is there anyone competent to own a deployment?), and frontline workforce capacity (are employees positioned to absorb workflow change?). Operating partners should score portfolio companies on these dimensions before committing capital, treating low readiness assets as future deployments rather than current ones.

## 7. Pillar II: Implementation Architecture

Selection discipline determines what to deploy. Implementation architecture determines whether the deployment survives. The principles in this pillar address the most common operational reasons portfolio AI projects fail to reach production.

### Principle 5: Buy Before Build, Partner Before Buy

The MIT NANDA research is unambiguous: vendor built and partnership solutions succeed roughly twice as often as internal builds (MIT NANDA, 2025). BCG's analysis reaches the same conclusion, recommending that portfolio companies buy rather than build, because horizontal AI tools are generally more economical and feature rich than those developed in house (BCG, 2025). The operating partner should treat internal build projects as the exception that requires explicit justification, not the default. Where building is justified, the project should be partnered with a specialist vendor rather than attempted entirely in house.

This principle has a corollary specific to portfolio companies. Portfolio company technology leaders sometimes resist the buy default because building creates an internal team and a perceived strategic asset. Operating partners should be aware of this incentive and counter it explicitly during use case approval.

### Principle 6: Anchor in EBITDA, Not Adoption

Adoption metrics (seats licensed, prompts submitted, users active) are easy to measure and almost meaningless as proxies for value. A common pattern is for a portfolio company to celebrate high adoption of a generative AI tool that has not produced any measurable change in cost, revenue, or working capital. Every funded AI initiative should have a stated EBITDA hypothesis and a measurement protocol agreed before deployment, not constructed retroactively.

The 100 day plan tradition in private equity offers a useful template. AI initiatives should be scoped with a specific EBITDA target and a measurement window appropriate to the use case. If the hypothesis cannot survive the measurement window, the initiative should be paused, redirected, or wound down.

### Principle 7: Designate a Portco AI Owner Who Is Not the CEO

A common failure mode is that the portfolio company chief executive is named as the executive sponsor for AI and is then drawn back into operating crises that crowd AI out. The result is initiatives without an active executive owner. The operating partner should require that a named executive at the portfolio company, ideally a chief operating officer, chief financial officer, or chief technology officer with specific authority, owns the AI workstream with measurable accountability. The chief executive remains the sponsor; the named owner remains the driver.

### Principle 8: Change Management as a First Class Workstream

The Bain report describes the resistance of portfolio company employees to AI as organ rejection, noting that successful firms tackle change management challenges directly rather than treating them as residual (Bain & Company, 2025). BCG's analysis is similar: the Reshape category of AI deployment requires change management of operating model and talent alongside new AI platforms (BCG, 2025). PwC's value creation guidance recommends decisive talent action early in the deployment cycle (PwC, 2026).

In practice, this means staffing change management as an explicit workstream with its own budget, its own owner, and its own milestones. It is not adequate to add change management to the technology project manager's responsibilities. The principal failure mode of well built AI tools is that they are not used.

### Principle 9: Treat Data Debt Before AI

Many portfolio companies, particularly those built through roll up strategies, carry substantial data debt: inconsistent definitions of basic operational metrics, fragmented systems, and missing or inaccurate master data. Layering AI on top of this debt produces unreliable outputs that erode trust quickly. The operating partner should require an honest data audit as a precondition to any AI initiative whose outputs depend on operational data quality. In some cases, the right first investment is not AI at all; it is the data infrastructure that will eventually make AI possible.

## 8. Pillar III: Cross Portfolio Compounding

The first two pillars optimise a single deployment. The third pillar is the difference between a sponsor that runs ten independent AI projects and a sponsor that builds a portfolio wide AI capability. Without deliberate compounding, the portfolio is doing the same work ten times. With it, each new deployment costs less and lands faster than the one before.

### Principle 10: The First Deployment Funds the Second

The cost of a successful AI deployment includes the cost of all the failures that produced it. Operating partners should treat the first successful deployment in a use case category as a portfolio asset, not a single company asset. Documentation, vendor relationships, training materials, change management playbooks, and measurement protocols should be captured in a transferable form during the first deployment, even when this slows the deployment slightly. This is the single most important habit for capturing portfolio level returns on AI investment.

### Principle 11: Vendor Master Agreements at the Firm Level

When an AI vendor proves successful in one portfolio company, the firm should negotiate master service agreements at the firm level that allow other portfolio companies to deploy on preferred commercial terms without renegotiating from scratch. This shortens deployment timelines materially. It also produces a partnership relationship with vendors who are then willing to invest in cross portfolio integration work. KPMG's 2025 value creation research notes that leading firms are explicitly centralising portfolio data and standardising vendor relationships to capture cross portfolio value (KPMG, 2025).

### Principle 12: A Portfolio Wide AI Council, Not Just an AI Operating Partner

The AI operating partner role is a useful innovation, but a single individual cannot carry the cross portfolio coordination load alone. Leading firms are beginning to convene quarterly councils of portfolio company AI owners (the named executives from Principle 7) to share what is working, what is failing, and what is being learned. These councils are most effective when they have explicit permission to discuss failures, not only successes. Bain's discussion of Apollo's structured playbook approach illustrates this, in which portfolio companies are asked to identify use cases and develop technology roadmaps that are then reviewed at the firm level (Bain & Company, 2025).

### Principle 13: Honest Failure Registries

The 95 percent failure rate documented by MIT is a portfolio wide phenomenon, not a unique problem at any single company. Yet within most portfolios, failures are minimised in reporting because admitting them feels like admitting underperformance. This creates a system where every portfolio company repeats the same mistakes. Operating partners should maintain a confidential, cross portfolio registry of AI initiatives that did not produce expected value, with documented reasons. This registry is one of the highest value artifacts an operating team can build, because it dramatically shortens the learning curve for the next deployment.

## 9. Implementation: A Phased Operating Partner Playbook

The PADF is most useful when applied as a phased operating sequence. The following timeline is calibrated to the typical operating partner's capacity constraint.

### Phase 1: Portfolio Diagnosis (Weeks 1 to 6)

Assess each portfolio company against the readiness criteria in Principle 4. Score and rank assets by combination of remaining hold period, AI readiness, and EBITDA opportunity. The output of this phase is a one page portfolio map showing where to invest operating bandwidth in the next twelve months and where to defer.

### Phase 2: Lead Asset Deployment (Months 2 to 6)

Select one or two lead assets where conditions for success are strongest. Deploy the full PADF: select use cases via Pillar I, structure the implementation via Pillar II, and document everything in transferable form per Principle 10. Resist the temptation to start everywhere at once. The MIT research notes that mid market organisations move faster from pilot to full implementation, with reported timelines around 90 days, while large enterprises often need nine months or more (MIT NANDA, 2025). Operating partners can use this difference deliberately, beginning with the assets where speed to production is realistic.

### Phase 3: Replication and Scaling (Months 6 to 18)

Use the documentation, vendor relationships, and playbooks from the lead deployments to accelerate the second and third waves. Convene the cross portfolio AI council quarterly. Begin building the failure registry. Negotiate firm level master agreements with the vendors who are emerging as partners.

### Phase 4: Exit Positioning (Twelve Months Pre Exit Through Close)

For assets approaching exit, shift focus from new AI deployments to documenting the value created by existing ones. Buyers are increasingly performing AI due diligence, examining a target's AI maturity, data integrity, and compliance readiness as material to enterprise value (Cuesta Partners, 2025). The operating partner's job at this stage is to ensure that the AI story is defensible: which initiatives, what measured impact, what governance, and what remains to be captured by the next owner.

## 10. Measuring Success: Beyond Adoption Metrics

The dominant failure mode in portfolio AI measurement is the substitution of activity metrics for outcome metrics. A serious measurement framework distinguishes three layers.

### Activity Metrics (necessary but insufficient)

Seats deployed, users active, prompts submitted, agents invoked. These confirm that something is happening but say nothing about value. They should be tracked, but never reported as the headline result of an AI initiative.

### Workflow Metrics (the bridge layer)

Tasks completed per hour, error rates, cycle times, customer touchpoints handled. These show that AI is changing how work gets done. They are leading indicators that an initiative is on track to deliver outcome impact.

### Outcome Metrics (the ones that matter at the firm level)

EBITDA impact, working capital release, revenue contribution, cost reduction. These should be measured against pre deployment baselines and reported to the investment committee on the same cadence as other value creation metrics.

The operating partner's role in measurement is not to design the dashboards. It is to insist that no AI initiative is funded without a stated outcome hypothesis, and that the hypothesis is tested honestly at the agreed point. A portfolio that learns this discipline will diverge sharply from one that does not, regardless of which specific technologies they deploy.

## 11. Conclusion and Recommendations

The portfolio AI conversation in 2026 is dominated by activity rather than outcomes. The vast majority of private equity firms have integrated generative AI into their M&A workflows, and the PE adopters are making meaningful financial commitments to do so. Most portfolio companies are running pilots. And yet the MIT NANDA research suggests that 95 percent of those efforts will not produce measurable P&L impact. The gap is not closing on its own.

The operating partner is uniquely positioned to close it. Cross portfolio visibility, selection authority, and vendor leverage are advantages that no portfolio company chief executive enjoys. The Portfolio AI Deployment Framework provides a structured way to use those advantages: select with discipline, deploy with architecture, and compound with intent.

The recommendations for sponsors deploying AI across their portfolios are as follows:

1. Treat portfolio AI as a portfolio management problem. Capital, bandwidth, and attention are scarce. Allocate them with the same discipline as investment capital.
2. Refuse the build default. The data on vendor versus internal build success rates is decisive. Build only with explicit justification and specialist partnership.
3. Anchor every initiative in EBITDA. Adoption metrics are not value. Require an outcome hypothesis at funding and test it honestly at the agreed point.
4. Stage the operating partner's involvement. One or two lead assets done well will produce more portfolio value than ten initiated and abandoned.
5. Capture the playbook the first time. The replicable artifacts produced by the lead deployment are worth more than the deployment itself.
6. Convene the cross portfolio council. The AI operating partner is necessary but not sufficient. The named owners across the portfolio need a shared forum.
7. Build a failure registry. The fastest way to compress the learning curve for the next deployment is to be honest about the last one.
8. Position for exit. Assets approaching exit should focus on documenting and defending what has been built, not initiating new deployments. AI due diligence by buyers will reward this discipline.

The firms that translate AI conviction into portfolio EBITDA will not be the ones that move fastest. **They will be the ones that select most carefully, deploy most deliberately, and compound most patiently.**

# References

AlixPartners. (2025). Practical AI for Private Equity Operating Partners. AlixPartners Insights.

Bain & Company. (2025). Field Notes from the Generative AI Insurgency in Private Equity. In Global Private Equity Report 2025. Bain & Company.

Bain & Company. (2026). Welcome to a New Era: Global Private Equity Report 2026. Bain & Company.

Boston Consulting Group. (2025). Private Equity: Value Creation in Portfolio Companies. Executive Perspectives on AI First Companies. BCG.

Boston Consulting Group. (2026). Inside the AI First Private Equity Firm. BCG Publications.

Cuesta Partners. (2025). The New Value Creation Playbook: How Mid Market Firms Are Moving First on AI. Cuesta Partners Insights.

Deloitte. (2025). 2025 GenAI in M&A Study. Deloitte Insights.

Dube, M., & Nejdawi, Y. (2025). The AI Operating Partner: The Latest PE Portfolio Value Creation Role? Korn Ferry Institute.

European Commission. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (AI Act).

FTI Consulting. (2025). Three Plays for Driving Value Creation in 2025. FTI Consulting.

KPMG International. (2025). 2025 Global PE Value Creation Survey. KPMG International.

KPMG International. (2025). AI Quarterly Pulse Survey, Q3 2025. KPMG.

MIT NANDA Initiative. (2025). The GenAI Divide: State of AI in Business 2025. Challapally, A., et al. MIT Media Lab Project NANDA.

PwC. (2026). What Is the Future of Portfolio Company Value Creation. PwC United States.

PwC. (2026). Global M&A Trends in Private Equity and Principal Investors: 2026 Outlook. PwC.

## 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 frameworks 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.

Dr. Coney's work is distinguished by its integration of behavioural science with practical technology deployment. He advises private equity firms, family offices, private credit funds, independent sponsors, and investment banks on AI strategy, governance design, and psychology informed change management.

This paper is part of an ongoing research series on responsible AI adoption in financial services. Previous publications include Agentic AI Governance in Private Equity: A Behavioural Framework for Autonomous Decision Systems (Q1 2026), Measuring AI ROI in Private Equity: A Framework for Decision Velocity vs Decision Quality (Q1 2026), and The Skill Erosion Paradox: Preserving Analytical Capability in AI Augmented Teams (Q1 2026).

## About WorkWise Solutions

WorkWise Solutions builds secure, purpose built AI systems for private equity, venture capital, family offices, private credit funds, and investment banking firms. The firm specialises in zero retention AI architecture that ensures proprietary deal flow and portfolio data never train public models.

WorkWise's approach is grounded in a core insight: most AI implementations fail not because of technology but because of broken workflows and poor adoption strategies. Every engagement integrates behavioural science and organisational psychology into the technical design, ensuring that AI systems become invisible, indispensable parts of how investment teams actually work.

**Contact:** Contact@workwisesolutions.org | workwisesolutions.org

**Schedule a consultation:** calendly.com/contact-atqf/30min