

Buy, Build, or Wait

AI Adoption Resolutions in European Private Equity

Abstract

This thesis examines how organizational configurations within European private equity (PE) funds produce different artificial intelligence (AI) adoption resolutions under competitive pressure and limited partner (LP) scrutiny. Existing research on PE and digital adoption operates at the portfolio-company level and treats the fund as a homogeneous driver, leaving the question of how funds themselves organize AI capabilities in investment analysis largely unexamined. The thesis shifts the unit of analysis from the portfolio company to the fund itself and from quantitative outcome measurement to substantive organizational mechanisms. Following a stringent data-driven mandate, semi-structured expert interviews were conducted with eleven European PE Funds and analyzed through grounded-theory-inspired cross-case coding, with practical saturation in code emergence confirmed within the sample. Wide variation in operational AI adoption maturity is documented across the eleven Funds. The variation in AI maturity does not track ticket size, geography, sector focus, LP profile, or interviewee seniority. What explains it is configuration variance: operationally-mandated AI personnel as the primary feature, with bottom-up AI initiative and active-governance signature as amplifier features. Funds' AI-adoption decisions map onto the configuration variance and resolve along three characteristic lines: build-resolution, articulating-without-operationalizing, and foundational-deferral. These are resolutions Funds occupy rather than exclusive types, navigated dimension-by-dimension across functional areas in Funds' investment workflows. The configuration-variance framework and the resulting AI-adoption resolutions are read through the resource-based view, real options under Knightian uncertainty, institutional theory, and the relational view, examining how the resolutions Funds occupy relate to differential exposure across competitive capability, customer relationship, and human capital, with value chain taken up as a fourth dimension through forward projection.

Keywords

Private Equity, Artificial Intelligence, AI Adoption, Technology Investment, Europe

Authors

Valentin Karl Ganter, Nico Eichhorn

Institution / Program / Specialization

Stockholm School of Economics / MSc Finance / Corporate Finance

Supervisor / Examiner

David T. Robinson, Professor, Duke University; Visiting Professor, Stockholm School of Economics / Gualtiero Azzalini, Assistant Professor, Stockholm School of Economics

Submission Date

18 May 2026

Buy, Build, or Wait: AI Adoption Resolutions in European Private Equity
Master's Thesis
MSc Finance
Stockholm School of Economics
© Valentin Karl Ganter & Nico Eichhorn, 2026

Contents

CONTENTS.....	3
LIST OF TABLES	5
LIST OF FIGURES	6
LIST OF ABBREVIATIONS	7
1. INTRODUCTION	8
2. LITERATURE REVIEW	9
2.1. Private Equity Operations and the AI Adoption Question.....	9
2.1.1. Private Equity Value Creation and the Operational Era	9
2.1.2. AI as Value Driver and Existing Research on AI in Private Equity	9
2.1.3. Research Gap, Question, and Contribution	10
2.2. Theoretical Frameworks for Capability Investment Under Uncertainty .	11
2.3. Institutional Theory on Capital Allocation Dynamics.....	13
2.4. Labor Markets, Technological Change, and the Relational View	15
3. METHODOLOGY / EMPIRICAL FRAMEWORK	16
3.1. Research Design	16
3.2. Data Collection and Sample.....	17
3.3. Data Analysis.....	19
3.4. Quality Criteria.....	22
3.5. Limitations.....	23
4. FINDINGS.....	24
4.1. Wide Variation in AI Adoption Maturity	24
4.1.1. Shared Sample Characteristics.....	24
4.1.2. The Operational AI Maturity Matrix	24
4.1.3. Structural Factors	25
4.2. Organizational Configuration and AI Maturity Outcomes.....	26
4.2.1. Dedicated Technology Personnel	26
4.2.2. Bottom-Up AI Initiative	27
4.2.3. Active-Governance Signature.....	27
4.2.4. The Joint-Pattern Configuration	27
4.3. Wait-and-See Posture and Configuration-Dependent Deferral	28

4.3.1.	The Deferring Subset of the Sample.....	28
4.3.2.	Two Content Clusters within the Deferring Subset	28
4.3.3.	The Content Tracks Personnel Configuration	30
4.4.	Anticipated Competitive Dynamics.....	30
4.4.1.	Interviewee Articulations.....	30
4.4.2.	Three Sources of Competitive Pressure.....	31
4.4.3.	A Counter-Current and Market Structure Shifts.....	32
4.5.	Buy or Build? Substantive AI Investment in the Sample	33
4.5.1.	Strategic and Administrative AI Investment	33
4.5.2.	The Buy-vs-Build Navigation as Analytical Lens.....	33
4.5.3.	Build-Side Operationalization	34
4.5.4.	The Buy-vs-Build Navigation Across the Sample.....	36
5.	DISCUSSION.....	38
5.1.	Build Where the Edge Is, Buy (Time) Where It Isn't	38
5.2.	Customer Relationships: AI and the Brand-Name Threshold	40
5.3.	Human Capital: Where Edge Survives Commoditization	42
6.	CONCLUSION	45
	REFERENCES.....	48
	APPENDIX.....	52

List of Tables

Table 1: Operational AI Maturity Matrix.....	24
Table 2: Initial questionnaire catalog and thematic interview structure.	52
Table 3: Sample Overview.	54
Table 4: Expert Sample Overview.....	55
Table 5: Detailed Analytical Journey per Phase.	56
Table 6: Post-Audit Codebook.....	58
Table 7: Sourcing Dimension Score Levels.	64
Table 8: Due Diligence Dimension Score Levels.	65
Table 9: Knowledge & Data Infrastructure Dimension Score Levels.	66
Table 10: Reporting & Portfolio Monitoring Dimension Score Levels.	67
Table 11: Codes by Themes and Subcategories.	68
Table 12: Cited NCORs in Sections 4 and 5.....	74
Table 13: Additional Standard Quality Criteria Provisions.....	75
Table 14: Configuration Feature Distribution Across Sample Funds.	77

List of Figures

Figure 1: The Analytical Journey.....	19
Figure 2: Code Saturation Curve.	76

List of Abbreviations

AI	Artificial Intelligence
AUM	Assets Under Management
CRM	Customer Relationship Management
DD	Due Diligence
ESG	Environmental, Social, and Governance
EXP	Expert Interviewee
GP	General Partner
IC	Investment Committee
INT	Interview Transcript Identifier
IT	Information Technology
KPI	Key Performance Indicator
LLM	Large Language Model
LP	Limited Partner
MAXQDA	Qualitative data analysis software
NCOR	Normalized Co-Occurrence Ratio
NPV	Net Present Value
PE	Private Equity
PMI	Post-Merger Integration
ROI	Return on Investment
SaaS	Software-as-a-Service
SMCF	Small/Mid-Cap Fund
SWOT	Strengths, Weaknesses, Opportunities, Threats
T1-T5	Theme 1 through Theme 5 (codebook themes)
VC	Venture Capital
VRIN	Valuable, Rare, Inimitable, Non-substitutable

1. Introduction

Artificial intelligence (AI) is reshaping how firms across industries operate, with productivity and innovation effects increasingly documented in the empirical literature (Babina et al., 2024; Brynjolfsson et al., 2021). Private equity (PE), where competitive edge rests on the speed and quality of investment analysis, sits at the frontier of these dynamics. Yet PE funds vary widely, not only in how they leverage AI in their workflows but in how they reason about whether to engage it at all. Intuitively, structural fund characteristics seem promising separators of adoption variation. They are not. The scholarly literature offers no direct explanation: Existing empirical work studies how PE ownership affects technology adoption at the portfolio-company level (Agrawal and Tambe, 2016; Baik et al., 2024); the PE fund itself, as the unit that reasons about and operationalizes AI in its own workflows, remains largely unexamined.

This thesis examines how organizational configurations within European private equity funds produce different AI adoption resolutions under conditions of competitive pressure and limited partner scrutiny. The study follows a stringent data-driven mandate. Semi-structured expert interviews with eleven European private equity Funds were conducted, and the transcripts analyzed through grounded-theory-inspired cross-case coding (Gioia et al., 2013; Strauss and Corbin, 1998), with practical saturation in code emergence confirmed within the eleven-interview sample.

Wide variation in operational AI-adoption maturity is documented across the eleven structurally comparable Funds. The variation tracks organizational configuration rather than structural fund characteristics, and a configuration-variance framework is developed to explain the AI-adoption resolutions Funds occupy. Organizational configurations and the resulting AI-adoption resolutions are read through the resource-based view, real options under Knightian uncertainty, institutional theory, and the relational view to project differential exposure across competitive capability, customer relationship, and human capital dimensions.

The remainder of the thesis is structured as follows: Section 2 lays the foundation for theoretical engagement. Section 3 illustrates the methodological design, process, grounding, and limitations. Section 4 presents findings across five themes emerging from cross-case coding. Section 5 reads the established framework through theoretical literature across three dimensions of differential exposure. Section 6 assesses study implications and concludes.

2. Literature Review

2.1. Private Equity Operations and the AI Adoption Question

2.1.1. Private Equity Value Creation and the Operational Era

Private equity value creation has historically operated through three substantive levers: financial engineering through leverage and capital structure optimization, governance engineering through concentrated ownership and board oversight, and operational engineering at portfolio companies (Kaplan and Strömberg, 2009). The contemporary paradigm increasingly weights operational engineering as the primary value-creation mechanism, with financial-engineering returns compressed by competition on entry valuations. Acharya et al. (2013) document that operational improvements drive outsized returns to PE investors beyond what financial restructuring alone explains. Davis et al. (2014) extend this pattern to the productivity dimension using U.S. data: PE ownership produces measurable increases in portfolio-company productivity, with reallocation of resources and operational restructuring as the substantive mechanisms.

The shift toward operational value creation increases the weight of fund-level analytical capability and operational expertise – not only capital structure optimization – in differentiating funds. Bernstein and Sheen (2016) document that investor-level expertise shapes portfolio company outcomes substantively. Lerner et al. (2011) counter older framings of PE as inherently short-horizon: PE ownership does not systematically reduce long-run investment, including investment in innovation-related capabilities.

2.1.2. AI as Value Driver and Existing Research on AI in Private Equity

Artificial intelligence operates as a general-purpose technology rather than as another iteration of digital infrastructure. Machine learning systems learn patterns from data rather than executing pre-specified rules; large language models (LLM) process unstructured text at scales humans cannot directly engage; pattern-recognition capability operates across data volumes that conventional analytical tools cannot accommodate. The productivity effects of AI materialize over longer horizons than typical information technology (IT) investments because AI requires complementary investments in intangible capital – organizational restructuring, workforce reskilling, process redesign – to translate technological capability into measurable productivity (Brynjolfsson et al.,

2021). This productivity-J-curve framing builds on the foundational observation that IT value depends substantively on complementary organizational capabilities rather than on technology investment alone (Brynjolfsson and Hitt, 2000).

Existing research on PE-driven digital and AI adoption operates at the portfolio-company level, examining how PE ownership affects portfolio firms' technology adoption decisions. Agrawal and Tambe (2016) document that PE-backed firms enhance employees' IT-related human capital following buyouts. Subsequent work extends this lineage to the digital and AI dimension specifically: PE-backed firms increase digital technology investment post-acquisition, with effects strengthening when investors carry prior board experience at IT companies and intensifying after the AlexNet shock that marked AI's commercial-viability turning point (Baik et al., 2024).

This portfolio-company-level pattern operates within a broader empirical context. Babina et al. (2024) find that firms investing in AI outperform their peers on sales, headcount, and market value, with the gains traced primarily to new-product development rather than to efficiency improvements in existing processes. The largest and most analytically sophisticated firms capture these returns disproportionately, a skew that, on Babina et al. (2024) evidence, pushes measured industry concentration upward.

2.1.3. Research Gap, Question, and Contribution

The research gap this thesis addresses formulates through two observations. First, existing PE-and-digital-adoption literature operates at the portfolio-company level and treats the PE fund as a homogeneous driver, leaving the question of how funds themselves organize AI capabilities substantively unexamined. Second, existing research is predominantly quantitative, examining outcomes – adoption rates, productivity effects, firm growth – without engaging the substantive organizational mechanisms by which funds resolve AI-adoption decisions.

The research question this thesis takes up follows directly from these observations: how do organizational configurations within European PE funds produce different AI-adoption resolutions, and how do these configurations operate under conditions of competitive pressure and limited partner scrutiny? The thesis shifts the unit of analysis from the portfolio company to the fund itself and from quantitative outcome measurement to substantive organizational mechanisms. Configuration variance – cross-fund variation in three organizational features (Section 4.2) – is the analytical lens this thesis develops. The contribution is twofold. First, the thesis demonstrates

empirically how configuration variance maps onto three characteristic AI-adoption resolutions: build-resolution, articulating-without-operationalizing, and foundational-deferral, navigated dimension-by-dimension. Second, it engages theoretically with how those resolutions relate to differential exposure across competitive capability, customer relationship, and human capital.

2.2. Theoretical Frameworks for Capability Investment Under Uncertainty

Resource-Based and Knowledge-Based View: The resource-based view treats firms as bundles of productive resources. Penrose (1959) established the framework's foundational claim: firms differ systematically in the resources they accumulate over time, and these resource differences produce differentiated competitive position. Barney (1991) operationalized the framework at firm level through the VRIN criteria – resources that are valuable, rare, inimitable, and non-substitutable produce sustained competitive advantage. The framework establishes that internal capability development is rational where the resources at stake satisfy VRIN conditions, because external acquisition itself dissolves the rarity-and-inimitability conditions that produce the resource's strategic value. The knowledge-based view extends this framework by identifying knowledge as the resource type most strongly associated with sustained competitive position. Grant (1996) developed the substantive mechanism that underpins this claim: knowledge resides in individuals, and the firm's distinctive role is to integrate the specialist knowledge distributed across them. What is firm-specific is this integration – the coordination mechanisms and common knowledge that make it possible – rather than the knowledge itself, and it does not transfer cleanly across firm boundaries, because codification loses tacit content while the integrating context cannot be transmitted through market contracting. The implication for capability investment operates on two grounds. First, where capability is tied to firm-specific knowledge, external acquisition fails because that knowledge does not transfer through market contracting. Second, a capability that depends on integrating knowledge distributed across the firm's own specialists cannot be redeployed outside that integrating context without losing the integration that makes it productive – making internal development rational.

Strategic Capability Development and Sustained Commitment: Strategic Capability does not emerge from point-in-time investment decisions; it emerges through sustained organizational commitment over time. Teece (1986) established the foundation for this through the framework of complementary assets: firms profit from innovation only when the

complementary capabilities required to operationalize it are in place – absent that, returns accrue to imitators and asset holders instead; sustained capability emerges through the firm's capacity to reconfigure internal competencies over time (Teece et al., 1997). Innovation alone does not produce sustained competitive advantage. Applied to this thesis, strategic capability operates through a sequenced commitment dynamic – prior strategic commitment to deploying capital toward capability, organizational architecture aligned with that commitment, then operational scaffolding sustained over time. Each link in the sequence is necessary but insufficient; capability emerges from integration rather than from any single component. Applied methodologically, observable resource deployment is read as the visible portion of an upstream commitment dynamic – personnel, infrastructure, and operational tools as markers of the commitment that produced them rather than substitutes for it. Point-in-time analysis cannot directly observe the underlying organizational commitment; the observable features serve as its visible traces.

Option Value and Investment Under Uncertainty: Investment under uncertainty operates in a decision environment that standard expected-value frameworks do not capture. Knight (1921) established the foundational distinction: risk operates with knowable probability distributions, where agents can calculate expected values and apply standard decision-theoretic frameworks; uncertainty operates without knowable probability distributions, where the standard calculation does not apply. This distinction is the foundation real-options theory later builds on: under uncertainty, premature commitment forecloses option value the agent cannot currently evaluate, which is what makes disciplined deferral a substantively rational response. Dixit and Pindyck (1994) operationalized Knightian uncertainty (Knight, 1921) as investment-deferral logic through real options theory. Investments under uncertainty carry option value that premature commitment destroys; the value of deferring depends on how much new information arrives during the wait, on whether the wait can be exercised when conditions clarify, and on the irreversibility of the commitment. Disciplined deferral is rational when option value of new information exceeds cost of delay. Applied to capability investment under technological uncertainty: where the technological substrate evolves substantively during the waiting period, premature commitment to infrastructure that may be obsolete destroys value that disciplined deferral preserves. Applied to this thesis, a further distinction is drawn within deferral: informed waiting, where option value is assessed and the conditions to exercise are in place, versus passive deferral, where the option cannot be exercised because the organizational capability to act on new information is absent.

Novy-Marx (2007) developed an equilibrium model of investment timing for heterogeneous firms in competitive uncertain environments. The substantive findings extend the partial-equilibrium real options framework: real option premia remain significant in equilibrium and do not get arbitrated away by competition, with firms delaying irreversible investment and choosing not to undertake some positive-Net present value projects even under competitive conditions. Equilibrium delay remains significant relative to partial-equilibrium predictions because each firm's deferral decision is informed by anticipation of other firms' eventual investments. In the model, firm heterogeneity yields a stable cross-sectional distribution of investment-timing decisions rather than convergence. Applied to this thesis, that result is read at the organizational level: funds differing in operational substrate, capability position, and commitment configuration are not competing identically on the same investment opportunity, so their divergent wait-versus-build decisions are a rational equilibrium outcome of that heterogeneity rather than analytical noise.

2.3. Institutional Theory on Capital Allocation Dynamics

Institutional Pressure and Capital Allocation: DiMaggio and Powell (1983) developed institutional isomorphism as the framework for explaining convergence among organizations operating in shared institutional environments. The substantive mechanism operates through three pressure types: coercive isomorphism arising from external dependence on resource providers and regulatory authorities; mimetic isomorphism arising from imitation of successful peer organizations under uncertainty; and normative isomorphism arising from professionalization and shared educational and professional networks. Organizational practices converge within institutional fields through the systematic operation of these three pressure mechanisms, with relative weight depending on the institutional environment. Aggarwal et al. (2011) provided empirical demonstration of the institutional-investor pressure mechanism specifically: institutional investors monitoring the firms they invest in produce measurable convergence in those firms' adoption of governance practices. Institutional investors operate as carriers of institutional pressure, with the pressure becoming legible at the firm-structure level when the investor base develops capacity to evaluate substantive practice rather than relying on formal-compliance proxies. The framework grounds the analytical move from institutional pressure as general convergence mechanism (DiMaggio and Powell, 1983) to capital-allocation-specific mechanism.

Applied to this thesis, capital allocators that develop the capacity to evaluate substantive practice rather than formal-compliance proxies exert convergence pressure unevenly across the firm distribution, concentrated on firms exposed to substantive evaluation rather than insulated by reputation.

Symbolic versus Substantive Adoption: Westphal and Zajac (2001) demonstrated the substantive distinction between formal-policy adoption and operational practice. Firms responding to institutional pressure can adopt symbolically – formal compliance with policy commitments without substantive operational change – or substantively – operational change consistent with the formal commitment. The decoupling between policy and practice is methodologically important because it distinguishes firms that respond to institutional pressure through visible signaling from firms that respond through substantive operational restructuring. Symbolic adoption operates as a strategic response when the cost of substantive operational change exceeds the cost of formal compliance. Applied to this thesis, the distinction between symbolic and substantive adoption becomes legible to an evaluating audience only when that audience can identify decoupling between policy and practice. This extends to capital-allocation evaluation: as capital allocators develop analytical capacity to distinguish substantive operational capability from formal-compliance signaling, the strategic value of symbolic adoption diminishes, and substantive operational practice becomes the dimension on which capital allocation turns.

Information Asymmetry and Agency in Capital Allocation: Akerlof (1970) developed the foundational framework for analyzing markets under asymmetric information through the market-for-lemons analysis. When sellers know quality information that buyers do not, and buyers cannot distinguish high-quality goods from low-quality at the point of transaction, high-quality goods exit the market because their price cannot rise above the average buyers are willing to pay. The substantive mechanism operates through the buyer's signal-evaluation problem: when reliable signals of quality are unavailable or costly to verify, buyers default to proxies that may or may not correlate with actual quality. Application to capital allocation: capital allocators face information asymmetry between funds with substantive operational capability and funds with formal-compliance signaling, and the allocator's signal-evaluation problem is which signal type to weight more heavily under imperfect information. Jensen and Meckling (1976) developed agency theory as the framework for analyzing principal-agent relationships under information asymmetry and incentive misalignment. Agency cost operates through three categories: monitoring costs incurred

by principals to verify agent behavior, bonding costs incurred by agents to demonstrate alignment with principal interests, and residual loss representing the divergence between agent behavior and principal-optimal outcomes that monitoring and bonding cannot fully eliminate. Application to capital allocation: limited partners operate as principals delegating capital deployment to fund managers as agents, with monitoring costs constraining due diligence (DD) depth on each allocation decision. When monitoring is constrained, principals rely on signal-based evaluation, and the strategic question for both principals and agents is which signals carry the most information about substantive operational capability. The three frameworks combined establish that capital allocation under information asymmetry depends on which signals allocators weight most heavily, and that the relative weighting between substantive infrastructure and reputational proxies depends on the allocator's monitoring capacity, signal information costs, and strategic position within the institutional environment (Aggarwal et al., 2011; DiMaggio and Powell, 1983).

2.4. Labor Markets, Technological Change, and the Relational View

Skill-Biased Technological Change and AI Workforce Effects: Bresnahan et al. (2002) developed the foundational framework for how information technology, organizational restructuring, and workforce skill-mix interact to shape firm productivity. The central claim is that productivity gains from technological investment depend on whether the firm restructures organizational practices and skill composition alongside the technology itself. The task-based framework of Autor et al. (2003) distinguishes two registers: technology substitutes for routine cognitive tasks where it replaces labor inputs at lower per-unit cost; technology complements non-routine judgment-intensive work where it amplifies rather than displaces the human contribution. Both registers operate at firm level simultaneously and produce differential exposure across the workforce by task composition. Acemoglu et al. (2022) extended this framework to occupation-level granularity through vacancy-level evidence on AI exposure: AI-exposed firms reshape hiring and skill-mix patterns concentrated in occupations with high task-substitution exposure rather than reducing workforce uniformly. Jevons (1865) observed the efficiency paradox in the case of coal: efficiency gains in a productive activity can expand demand for the underlying input rather than contracting it, as lower per-unit cost makes more activity economically viable.

The Relational View and Sources of Edge Beyond Firm Boundaries: Dyer and Singh (1998) developed the relational view, which complements the resource-based and knowledge-

based foundations Section 2.2 establishes. Where the resource-based view (Barney, 1991; Penrose, 1959) and the knowledge-based view (Grant, 1996) examines resources within firm boundaries as sources of sustained competitive advantage, the relational view examines competitive advantage arising between firms through inter-firm relationships. The framework complements rather than competes with the within-firm view, identifying competitive advantage at a different organizational level. Relational rents are supernormal profits that neither firm could generate alone but emerge from the relationship between them. Dyer and Singh (1998) identify four sources: relation-specific assets, knowledge-sharing routines, complementary resources and capabilities, and effective governance. The methodological implication is that firms facing commoditization of within-firm analytical capability retain access to competitive advantage through the relational dimension – sources of edge that within-firm commoditization does not erode.

3. Methodology / Empirical Framework

3.1. Research Design

This thesis adopts a qualitative, exploratory multiple-case study design. The exploratory rather than hypothesis-testing design follows from the state of the literature on AI adoption in private equity: available evidence is fragmented, no consolidated theoretical framework specifies the organizational determinants of adoption, and accounts of adoption drivers conflict across studies and practitioner sources. The hypothesis space was therefore underdetermined; the analytical objective was instead to develop an explanatory understanding of variation in an underexplored organizational phenomenon by building explanation from rich empirical material rather than imposing *ex ante* measurement structures. Case study research is particularly suitable in this configuration of theoretical maturity (Edmondson and McManus, 2007; Eisenhardt, 1989).

The choice of a multiple- over a single-case design follows from the analytical purpose of the thesis. Rather than producing an in-depth description of a single bounded organizational setting, the study seeks to identify patterned similarities and differences across Funds that vary in size segment, geography, and organizational configuration; a multiple-case design therefore offers stronger leverage for cross-case comparison and theory development, particularly where replication logic rather than sampling logic guides case selection and interpretation (Eisenhardt, 1989; Yin, 2018). This design choice also reflects the principle of methodological fit, according to

which research design should be aligned with the maturity of prior theory and the type of contribution sought (Edmondson and McManus, 2007). Case-based research should not rest on description alone, but should use empirical variation to sharpen a broader conceptual contribution, which is the role the multiple-case design serves in this thesis (Siggelkow, 2007).

3.2. Data Collection and Sample

Data was collected through semi-structured expert interviews. This method was selected because it combines the thematic comparability required for cross-case analysis with the flexibility needed to probe respondent-specific experiences, interpretations, and organizational processes (Kvale and Brinkmann, 2009; Rubin and Rubin, 2005). Consistent with Kvale and Brinkmann's understanding of the qualitative interview as a knowledge-producing conversation, the aim was not merely to extract factual information, but to elicit informed accounts of how organizational practices, constraints, and adoption logics are perceived and experienced by participants. This was particularly important because AI adoption in private equity is shaped by multiple organizational factors rather than by a single formal decision point.

The interview guide (Table 2) is organized into thematic blocks covering workflow practices, strategy and governance, future outlook, stakeholder relationships, and peer-divergence anomalies. It was designed to capture broad coverage of AI readiness and organizational sentiment rather than to test specific propositions; the semi-structured format allowed emergent topics to become substantive material. The guide was applied stably across all interviews and was not iteratively refined on the basis of early interviews. Interviews were recorded and transcribed in first draft using Read.ai, where the interviewee and the interviewers had given prior verbal or written consent. The first-draft transcripts were transcribed verbatim in the original interview language and then manually reviewed against the recorded video by the interviewers and adjusted to capture actual wording exactly. Each transcript was categorized with a header section documenting Fund characteristics, the interview identifier (INT label), and the Fund identifier (Fund label) used throughout the analysis.

The European private equity market in which the sample Funds operate is sufficiently concentrated that combinations of structural Fund characteristics would render individual Funds re-identifiable. Several interviewees explicitly requested anonymization as a precondition for participation, and the interview material engages strategically sensitive content that warrants strict

confidentiality protections. Fund-level characteristics are therefore disclosed at two levels: per-Fund disclosure (Table 3) covers four characteristics – interview identifier (INT label), Fund identifier (Fund label), ticket size category (Growth, Venture Capital, Large-Cap, Mid-Cap, Small-Cap), and interviewee seniority (Junior, Senior) – while three further characteristics (office geography, sector focus, limited partner (LP) profile) are reported only as aggregate sample-level distributions wherever analytically relevant in the Findings chapter, allowing verification of the structural-elimination claim in Section 4.1.3 while preserving per-Fund identification only for analytically required characteristics. Transcripts were pseudonymized to operationalize this disclosure: Fund names were replaced with the Fund identifiers used throughout the thesis (Fund 1 through Fund 11), proprietary system names with descriptive placeholders (e.g., ‘[Fund vehicle name]’, ‘[Fund 10 sourcing tool]’), and any geographic or contextual detail that would render a respondent identifiable was generalized.

Fieldwork included five expert interviews beyond the eleven core cases, conducted to capture multi-stakeholder perspectives on AI adoption in PE. Four involved stakeholders outside PE funds and are retained as fieldwork context only, having proved non-substantive for the analytical chapters: the LP interviewee was a junior at a fund-of-funds rather than a direct LP allocator; the consultant did not work on PE deals; the banker described an advisor-side pitch document unavailable for examination; and the technology professional did not work on PE implementation. The fifth interview, with a senior dedicated AI professional at a Large-Cap PE Fund (EXP01, INT01), offered direct experience with PE-internal AI implementation and is engaged as an expert-informant in the Discussion and Conclusion. INT01 was not recorded (only structured field notes are available) and is therefore held outside the cross-case coding sample (Bogner et al., 2009; Meuser and Nagel, 1991). Table 4 provides the full expert sample overview.


The final analytical sample consists of eleven verbatim-coded interviews with private equity professionals, each yielding a Fund-level case for cross-case analysis. The eleven Funds span the major European private equity strategy classes – Venture Capital (VC) (n=1), Growth (n=2), Small-Cap (n=2), Mid-Cap (n=3), and Large-Cap (n=3) – across three geographic contexts (Geography 1, 2, 3). The ticket size class approximates Fund size segment as a deliberate anonymization choice; assets-under-management (AUM) is not reported. Interviewee seniority is reported in two bands: Junior (n=8), and Senior (n=3). Additional factors are sector focus: Technology-focused, diversified, non-technology-focused; and LP profile: Primarily Institutional,

Primarily Private. Cases were selected purposively rather than representatively, following a logic of theoretical sampling aimed at capturing meaningful variation across relevant organizational contexts rather than statistical representativeness (Eisenhardt, 1989). Table 3 provides the sample overview. The analytical chapters scope to pre-investment operations, the workflow axis comparable across all eleven cases; portfolio company AI adoption is studied in the adjacent quantitative literature (Baik et al., 2024).

3.3. Data Analysis

The analytical journey from transcripts to the cross-case structure used in the Findings chapter unfolded across three coding phases, a structured cross-interview audit, the emergence of a five-theme codebook, and the inductive construction of the Operational AI Maturity Matrix (hereafter AI Matrix). Each step was data-driven: codes were generated from interview language rather than imposed from prior frameworks, the theme structure was derived from the codebook rather than being pre-specified, and the AI Matrix dimensions were derived from the workflow areas covered comparably across the interviews. The analytical approach is therefore grounded theory-inspired and abductive – codes and themes emerged from the interview material, with theoretical literature engaged at the interpretation stage to connect these patterns to a broader analytical framework. Figure 1 summarizes the analytical journey; Table 5 provides the detailed procedure per phase.

Figure 1: The Analytical Journey. From data collection (Section 3.2) through three coding phases, the cross-interview audit, and thematic categorization to the analytical instruments — the AI Matrix and the Normalized Co-Occurrence Ratio (NCOR) (Section 3.3). Single arrows indicate procedural sequence; two-sided arrows indicate iterative processes.



Throughout the analysis, the distinction between artificial intelligence and broader technology was operationalized at the code level. Technology in the broader sense – digital tools,

data infrastructure, workflow systems, productivity software – provides the substrate; artificial intelligence is the specific class of capability that may be layered on top of that substrate to produce inference, retrieval, generation, or scoring functions, and the chapter’s analytical interest is in this AI layer specifically rather than in the underlying technology substrate. Codes capturing AI integration into investment workflows (large language model use, AI-driven retrieval, AI-screened search, machine-learning-based scoring) were kept analytically separate from codes capturing technology infrastructure more broadly (CRM systems, data warehouses, manual reporting tools, productivity software).

Pseudonymized interview transcripts were imported into MAXQDA and analyzed through three coding phases, following a grounded theory-inspired approach (Strauss and Corbin, 1998). Open coding generated approximately 200 first-order codes closely grounded in interview language; MAXQDA’s AI assistance was only used in this phase to suggest code names and definitions, with all suggestions reviewed by the researchers and accepted, modified, or rejected based on substantive fit with the empirical material. Axial coding consolidated these into higher-order categories; selective coding identified core categories and the narrative logic linking emerging themes across cases. The reduction to a final codebook of 162 leaf-level codes operated through code merging (concept-level consolidation, traceable via the ‘+’ marker) and Descriptive Archive separation of interview-specific content not connecting to cross-case themes (~20 codes). Codes are treated as interpretive analytical categories: cross-case analysis operates at binary code coverage rather than code frequency, with carrier counts describing patterns of articulation across the sample rather than metric observations. Throughout the phases, a structured cross-interview audit ensured codes were applied consistently across transcripts, consistent with constant comparative analysis (Glaser and Strauss, 1967) and second-order theme construction (Gioia et al., 2013). The final post-audit codebook (Table 6) is the basis for all cross-case analyses reported in the Findings chapter.

All codes in the final post-audit codebook were clustered into higher-level thematic categories. Codes engaging the operational status quo of AI use were consolidated into the AI Matrix (Table 1), structured around the interview questionnaire’s coverage of the PE investment workflows and Fund operating systems. The matrix scores each of the eleven Funds on four operational dimensions – Sourcing, Due Diligence, Knowledge and Data Infrastructure, and Reporting and Portfolio Monitoring – using a 0–3 scale anchored to observed AI integration. Each

dimension yields a per-Fund score; the four together produce a composite operational mean. Score anchors were calibrated from the most operationally articulated examples in the transcripts rather than from a pre-specified rubric. The four dimensions capture the output of organizational engagement with AI; configuration features (personnel, governance, culture) are intentionally excluded from the matrix and analyzed as inputs in Section 4.2 to avoid circular construction. The score descriptions and codes consistent with scores are documented in Tables 7–10. Composite means are partitioned into three tiers: High ≥ 2.0 ; Middle 1.0–1.99; Low < 1.0 .

The AI Matrix construction is part of the analytical sequence the chapter follows. Once the matrix had operationalized the cross-case maturity distribution (Theme 1), the variation visible in the matrix was tested against structural Fund characteristics and found not to track them (Section 4.1.3). The unanswered question – what features distinguish Funds along the maturity dimension if not structural ones – directed the analysis to the organizational configuration features under Theme 2 (Section 4.2). What remained was the question of what Funds at lower and middle tier articulate about their non-adoption; this directed the analysis to the deferral content under Theme 3 (Section 4.3). The deferral analysis surfaced articulations of the broader competitive environment, directing the analysis to the codes on peer pressure, LP pressure, and industry-level commoditization under Theme 4 (Section 4.4). The structural conditions Funds anticipate then framed the remaining question what substantive AI investment looks like within this environment, taken up under Theme 5 (Section 4.5), which engages knowledge-institutionalization. Each theme transition was driven by the unanswered question from the prior theme’s empirical material; the five-section structure of the Findings chapter was earned by this analytical sequence.

Five themes emerged from clustering: Theme 1 ‘AI Maturity Variation & Context’ (Section 4.1), Theme 2 ‘Organizational Configuration Features’ (Section 4.2), Theme 3 ‘Wait-and-See Deferral Patterns’ (Section 4.3), Theme 4 ‘Anticipated Competitive Dynamics’ (Section 4.4), and Theme 5 ‘Knowledge Institutionalization Gap’ (Section 4.5). Theme assignment is not exclusive – 21 codes touching multiple analytical domains appear under two themes each, producing 183 theme-level rows from 162 unique codes (Table 11); cross-theme analysis is conducted at the unique-code level to avoid double-counting.

Cross-case patterns are analyzed through the Normalized Co-Occurrence Ratio (NCOR), which measures how strongly two codes (A, B) co-occur on the same Funds as a proportion of the maximum possible co-occurrence:

$$NCOR(A, B) = \frac{|C_A \cap C_B|}{\min(|C_A|, |C_B|)} \quad (1)$$

where C_A and C_B denote the sets of Funds carrying codes A and B. A 100% NCOR indicates the rarer code's carrier set is fully contained in the more common code's carrier set; lower values indicate partial containment. Table 12 provides the NCORs engaged in the Findings section.

Cross-case analysis at the code carriage level is complemented by within-case analysis where small-N subsets warrant it (see Section 4.5.3), with the analytical contribution being the configuration each case occupies rather than a population-level distributional claim (Eisenhardt, 1989; Yin, 2018).

3.4. Quality Criteria

Several measures were employed to strengthen the methodological quality of the study, following Lincoln and Guba's (1985) framework for trustworthiness in qualitative research (credibility, transferability, dependability, confirmability). The provisions most analytically consequential for the cross-case design are discussed below; additional standard provisions (triangulation across Fund characteristics, negative case analysis, researcher reflexivity) are documented in Table 13.

Theoretical saturation was assessed through a saturation curve plotting cumulative code emergence across the chronological order of data collection (Figure 2); new-code emergence declined steeply after the fourth interview and flattened by the seventh, confirming practical saturation within the eleven-interview sample (Glaser and Strauss, 1967; Guest et al., 2006).

Each Fund is represented by interview material from one respondent; articulations are individual rather than institutional positions, and the chapter operates on articulation-pattern level rather than at the level of institutional fact-claim (Eisenhardt and Graebner, 2007; Yin, 2018).

The chapter engages a forward-looking topic – strategic AI adoption rests on respondents' views about how competitive conditions will evolve. Current state and forward-looking articulation cannot be cleanly separated; the Findings chapter documents current strategic positions alongside the forward-looking views that ground them, with standalone forward-looking projections receiving substantive analytical engagement in the Discussion.

The study's claims should be understood as analytic rather than statistical generalization (Yin, 2018) – the objective is to develop concepts and theoretical relationships that may be transferable beyond the immediate cases under study.

3.5. Limitations

This study is subject to several limitations.

First, the cross-case sample is small ($N=11$), consistent with qualitative case-study research where analytic depth is prioritized over statistical breadth (Eisenhardt, 1989; Yin, 2018). The eleven Funds support the analytical contribution – empirical demonstration of configuration variance – but cannot ground population-level claims about the prevalence of any configuration in the European PE industry. Saturation analysis (Section 3.4) indicates the analytical codebook stabilized within the eleven-interview sample, but additional cases would extend the empirical range of configuration types observed.

Second, the material is based on self-reported accounts, meaning that adoption maturity and organizational capability were not independently verified through direct observation or internal documents.

Third, the sample is skewed toward junior-level perspectives (Junior $n=8$, Senior $n=3$), meaning that the direct voice of General Partner (GP)- or Partner-level decision-makers is limited. This matters because strategic technology decisions may be framed differently at the level where budgets and final governance authority reside.

Fourth, the sample is geographically concentrated to European coverage. The findings should therefore be understood as analytically grounded in a specific European context rather than as exhaustive of the worldwide private equity landscape.

Fifth, some of the interviews were conducted in languages other than English, while verbatim quotes from interviews are presented in English translation in the chapter. Although the translation was carried out with care, some nuance, tone, or context-specific meaning may have been attenuated, particularly where respondents relied on idiomatic formulations.

Sixth, the semi-structured interview format combined with the time constraints typical of practitioner research required interviewer judgment about which guide questions to prioritize per interview. Coverage of specific topics consequently varies across transcripts. The analytical claims of the chapter rely on patterns visible across substantively engaged segments rather than on uniform coverage of every topic across every Fund, consistent with the logic of qualitative case-study research where claims are grounded in segments where the topic was substantively discussed (Kvale and Brinkmann, 2009; Rubin and Rubin, 2005)

Finally, the cross-sectional design and qualitative co-occurrence basis of the cross-case analysis preclude directional causal inference. NCORs report joint-articulation patterns rather than measurement of causal effect; theoretical interpretation in the Discussion draws on the empirical material as patterned association consistent with prior theory.

4. Findings

4.1. Wide Variation in AI Adoption Maturity

4.1.1. Shared Sample Characteristics

The eleven Funds in the sample share key structural features. They are all European private equity firms with established institutional positioning and comparable structural conditions. They invest in similar geographies, target overlapping industry segments, and face the same broad market conditions. The question this section examines is whether the structural similarity of the sample is associated with similar AI adoption outcomes, or whether observed variation operates below the level of structural similarity.

4.1.2. The Operational AI Maturity Matrix

The AI Matrix (Section 3.3; rubric and per-Fund justifications in Tables 7-10) scores each of the eleven Funds across the four operational AI capability dimensions. The composite mean anchors cross-Fund comparisons in this and subsequent sections.

Table 1: Operational AI Maturity Matrix. Per-Fund Scoring Documentation. Eleven Funds scored on four operational dimensions (0–3 scale; rubric in Tables 7-10); Mean is the arithmetic average across dimensions; tiers at High ≥ 2.0 , Middle 1.0-1.99, Low < 1.0 .

Fund	Sourcing	Due Diligence	Knowledge & Data	Reporting & Portfolio	Mean	Tier
Fund 5	3	2	3	2	2.50	High
Fund 2	3	3	2	2	2.50	High
Fund 10	3	1	3	2	2.25	High
Fund 8	2	1	3	1	1.75	Middle
Fund 7	1	2	3	0	1.50	Middle
Fund 4	1	1	2	1	1.25	Middle
Fund 3	2	1	2	0	1.25	Middle
Fund 1	1	1	1	0	0.75	Low
Fund 11	1	1	0	0	0.50	Low
Fund 9	0	1	0	0	0.25	Low
Fund 6	0	0	0	1	0.25	Low

The distribution of mean operational scores is asymmetric and structurally differentiated. Three Funds occupy the upper tier with means at or above 2.0: Fund 5 (mean 2.50), Fund 2 (2.50), and Fund 10 (2.25). Four Funds occupy the middle tier with means between 1.0 and 1.99: Fund 8 (1.75), Fund 7 (1.50), Fund 4 (1.25), and Fund 3 (1.25). Four Funds occupy the lower tier with means below 1.0: Fund 1 (0.75), Fund 11 (0.50), Fund 9 (0.25), and Fund 6 (0.25). The wide spread of operational AI capability across an otherwise comparable sample of European PE Funds, and the absence of any Fund operationalizing AI uniformly across all four dimensions, is the central empirical observation.

4.1.3. Structural Factors

The variation documented in the AI Matrix raises the question of whether structural factors that distinguish PE funds at the vehicle level or interviewee seniority account for the observed differences in AI adoption maturity.

Ticket size: The upper tier (Funds 5, 2, 10) contains a Mid-Cap Fund, a VC Fund, and a Growth Fund. The middle tier contains Mid-Cap and Large-Cap Funds. The lower tier contains Growth, Small-Cap, and Large-Cap Funds. Ticket size is not associated with consistent maturity outcomes.

Three further structural characteristics – office geography, sector focus, and LP profile – are reported only as aggregate sample-level distributions to preserve confidentiality (per Section 3.2). Across all three characteristics, within-category variance is wide: every examined category contains Funds spanning at least 1.0 point of operational maturity, and most categories span at least 2.0 points. For the three geographic cohorts, two cohorts span at least 2.0 maturity points and one cohort spans at least 1.0. For the three sector focus categories, two span at least 2.0 maturity points and one spans at least 1.5. For the two LP profile categories, one spans at least 2.0 and the other at least 1.5. Within-category variance of this magnitude indicates that no examined characteristic cleanly tracks the maturity distribution.

Interviewee seniority: The sample combines Junior and Senior interviewees, with both groups represented across all three maturity tiers. Interviewee seniority is not associated with consistent maturity outcomes.

Across the examined characteristics, the variation documented in the AI Matrix is not related to structural fund features: Funds with similar ticket size, seniority, geography, sector

focus, and LP profile occupy different tiers of operational AI capability. Interviewee Seniority does not explain variation either.

4.2. Organizational Configuration and AI Maturity Outcomes

4.2.1. Dedicated Technology Personnel

The ‘Dedicated technology personnel with ownership’ code appears in 7 of 11 Funds (Funds 1, 2, 3, 5, 7, 8, 10). ‘Absence of dedicated technology personnel’ is coded in the remaining 4 Funds (Funds 4, 6, 9, 11). The binary partition is exhaustive across the sample.

The personnel feature is associated with the strongest gradient observed in the section. Funds with the code have a mean operational AI capability score of 1.79; Funds without it have a mean of 0.56. The mandate scope of the personnel role refines this gradient further.

The first group consists of five Funds whose dedicated technology personnel role is operationally mandated for AI integration into deal-making workflows. These are Funds 2, 5, 7, 8, and 10, with operational means of 2.50, 2.50, 1.50, 1.75, and 2.25 (group mean 2.10). The second group consists of two Funds carrying the personnel code without operational mandate – the role is scoped to cyber security, IT, portfolio-company digitization, or general process improvement rather than AI integration into the investment workflow. These are Fund 1 (operational mean 0.75) and Fund 3 (operational mean 1.25; group mean 1.0). The third group consists of four Funds with no dedicated technology personnel: Funds 4, 6, 9, and 11, with operational means of 1.25, 0.25, 0.25, and 0.50 (group mean 0.56).

Fund 4 (INT06) articulates the operationally-mandated role as the structural feature its own configuration lacks: "their work is more like cyber security [...] there has to be this kind of person [responsible for AI], and that person does not exist at [the Fund]" (Fund 4, INT06). Fund 9 (INT13) articulates personnel absence as a consequence of the firm-wide top-management orientation toward "AUM per head", with no operational AI personnel allocation across the firm (Fund 9, INT13). The second and third groups overlap at the boundary – Fund 3 (group 2, mean 1.25) and Fund 4 (group 3, mean 1.25) tie on the operational mean – but every group-1 Fund scores at or above 1.50, no group-3 Fund scores above 1.25, and group 2 sits between.

4.2.2. Bottom-Up AI Initiative

The ‘Bottom-up AI initiatives’ code is carried by 5 of 11 Funds (Funds 2, 4, 5, 6, 11). Funds with the code have a mean operational score of 1.40; Funds without it have a mean of 1.29. Bottom-up AI initiative alone does not differentiate operational AI capability outcomes in this sample.

The analysis examines the joint occurrence of ‘Bottom-up AI initiatives’ with either operationally-mandated AI personnel or the absence of it. Fund 5 (bottom-up code, mandated personnel, operational mean 2.50) and Fund 2 (bottom-up code, mandated personnel, operational mean 2.50) sit at the upper end of the AI Matrix. Fund 4 (bottom-up code, absence of personnel, operational mean 1.25), Fund 11 (bottom-up code, absence of personnel, 0.50), and Fund 6 (bottom-up code, absence of personnel, 0.25) sit lower. The bottom-up feature is associated with higher operational AI scoring only in combination with operationally-mandated AI personnel.

4.2.3. Active-Governance Signature

The C-suite governance pattern is captured by the co-occurrence of two code pairs. ‘Centralized technology decision making’ (n=7) or ‘C-suite decision for major technology adoption’ (n=7) paired with ‘Concrete discovery steps to understand AI opportunities’ (n=4). The presence of this pattern in one Fund is defined as active-governance signature in this section: centralized technology decision authority — a structural proxy for where AI-adoption decisions are made — paired with documented AI discovery activity. This pattern is observed in four Funds: Fund 5 (operational mean 2.50), Fund 8 (1.75), Fund 11 (0.50), and Fund 6 (0.25), spanning the full maturity distribution — indicating that the active-governance signature is not associated with maturity outcomes on its own.

These four Funds split along the operationally-mandated AI personnel partition: Funds 5 and 8 carry it (group mean 2.13); Funds 11 and 6 do not (group mean 0.38). The active-governance signature, parallel to the bottom-up AI initiative feature, is associated with higher operational outcomes in the sample only when paired with operationally-mandated AI personnel.

4.2.4. The Joint-Pattern Configuration

Three configuration features have been examined for their association with operational AI capability outcomes in the sample. Operationally-mandated AI personnel is observed in every Fund in the upper tier (Funds 5, 2, 10) and in two of the four middle-tier Funds (Funds 7, 8); no

Fund scores above 1.25 without it. The personnel feature alone is associated with the strongest gradient observed in the section (Section 4.2.1). Bottom-up AI initiative and active-governance signature are each observed across the maturity distribution but are associated with higher operational outcomes only when paired with operationally-mandated AI personnel (Sections 4.2.2 and 4.2.3).

Within the operationally-mandated AI personnel group, amplifier features are associated with higher operational outcomes, but the relationship is not monotonic – Fund 10 (personnel alone, 2.25)¹ sits above Fund 8 (personnel plus both amplifiers, 1.75). Table 14 presents the full feature distribution across the sample.

4.3. Wait-and-See Posture and Configuration-Dependent Deferral

4.3.1. The Deferring Subset of the Sample

The ‘Wait-and-see AI adoption strategy’ code is present in 6 of 11 Funds: Funds 1, 3, 6, 7, 9, and 11. These Funds articulate a deliberate posture of holding off on substantive AI investment; they are referred to as the deferring Funds throughout this section. The remaining five Funds (2, 4, 5, 8, 10) do not carry this code; the codes characterizing their engagement with AI investment and the operational infrastructure several have built are examined in Section 4.5. The postures are not mutually exclusive: some Funds in the deferring subset also carry codes engaged in Section 4.5, indicating dimension-by-dimension navigation of AI adoption rather than uniform Fund-level posture. The deferring subset includes the four lowest maturity Funds (1, 11, 9, 6) and two middle-tier Funds (7 and 3).

4.3.2. Two Content Clusters within the Deferring Subset

The codes capturing deferral content within these six Funds cluster into two analytically distinct groups: granular and foundational concerns. The granular cluster comprises ‘Skepticism: Output reliability’ (n=2), ‘Skepticism: Technical feasibility limits’ (n=4), ‘Data validity as bottleneck for AI adoption’ (n=3), and ‘Trust in human work over AI tools’ (n=4); these are tool-level concerns about specific AI implementations. The foundational cluster comprises ‘Skepticism: ROI Uncertainty’ (n=4), ‘Lacking clean data storage processes’ (n=4), and ‘Cost concerns for AI’

¹INT14’s coverage emphasized operationalization over internal mechanisms; Fund 10’s apparent absence of amplifier features may reflect this rather than feature absence (Section 3.5).

(n=3); these are configuration-level concerns about whether AI evaluation is even possible, with Return on Investment (ROI) doubts articulated in the abstract, absence of clean data infrastructure, and cost concerns articulated as structural rather than tool-specific.

Within the deferring subset, Funds 1 and 7 carry only granular cluster codes, while Funds 3, 6, 9, and 11 carry foundational cluster codes (some alongside granular ones). The clusters are not mutually exclusive at the level of code carriage. Because foundational concerns operate upstream of granular ones – doubts about whether AI evaluation is possible at all precede doubts about specific tool performance – cluster placement follows the deepest concern articulated: a Fund carrying any foundational concerns is placed in the foundational cluster, deferring at the level of whether AI evaluation is possible; a Fund carrying only granular concerns is placed in the granular cluster, deferring at the level of specific tools.

Fund 7 (INT09) illustrates the granular cluster. The interviewee names current tool-level limitations as the deferral content: "the AI tools simply do not yet function well. Financial outputs are never correct. Always adjustments, always adjustments" (Fund 7, INT09). The deferral is conditioned on observable improvement in tool outputs rather than on configuration prerequisites. Fund 1 (INT02) articulates 'Overreliance on AI could harm investment decisions' as a specific decision risk tied to specific tool categories, with deferral content naming particular AI applications and the trade-offs involved. In both cases the deferral is granular: tool-specific constraints, revisable concerns, articulation at the level of specific AI tools rather than at the level of whether AI evaluation is possible at all.

Fund 9 (INT13) illustrates the foundational cluster. The interviewee describes the absence of basic data infrastructure: "there are 8000 different Excel templates for the same analysis [...] everyone builds their own and uses that" (Fund 9, INT13). The absence of basic data infrastructure operates upstream of specific-tool evaluation. Funds 6 and 11 carry the foundational cluster on similar grounds: ROI uncertainty articulated as structural rather than tool-specific, data infrastructure deficits, and cost concerns articulated as configuration constraints. Fund 3 (INT05) carries the heaviest foundational profile in the deferring subset, with all three foundational codes – 'Skepticism: ROI Uncertainty', 'Lacking clean data storage processes', and 'Cost concerns for AI' – alongside structural-skeptical codes on PE conservatism and predictable competitive advantage. Fund 3's interviewee articulates a structural barrier at the responsibility-transfer level, naming reluctance to delegate decision responsibility to AI as a firm-wide cultural condition: "I

am reluctant to give up responsibility and reluctant to relinquish it" (Fund 3, INT05). The cultural posture toward responsibility delegation sits upstream of any specific tool evaluation.

4.3.3. The Content Tracks Personnel Configuration

Among the six deferring Funds, three carry the ‘Dedicated technology personnel with ownership’ code (Funds 1, 3, 7) and three do not (Funds 6, 9, 11). Within the personnel-present group, Funds 1 and 3 carry the scope-only personnel mandate covering portfolio-company digitization rather than Fund workflow AI implementation, while Fund 7 carries the operationally-mandated AI personnel role – consistent with the partition documented in Section 4.2.1.

The cluster placements established in 4.3.2 – granular (Funds 1, 7) and foundational (Funds 3, 6, 9, 11) – map onto the personnel partition. Granular cluster Funds are both personnel-present. Foundational cluster Funds 6, 9, and 11 are personnel-absent; Fund 3 is the one boundary case, carrying a scope-only personnel mandate while its substantive deferral content sits in the foundational cluster.

The foundational cluster codes co-occur with ‘Absence of dedicated technology personnel’ at 75% NCOR for the two largest contributing codes: ‘Skepticism: ROI Uncertainty’ (3 of 4 carriers personnel-absent) and ‘Lacking clean data storage processes’ (3 of 4 personnel-absent). In each case the personnel-present carrier is Fund 3, the boundary case whose scope-only mandate does not provide the operational AI integration that the operational-mandate group carries.

The cluster placement maps onto the AI Matrix: the foundational cluster sits at a group mean of 0.56 (Funds 3, 6, 9, 11 at 1.25, 0.25, 0.25, 0.50) while the granular cluster sits at 1.13 (Funds 1 and 7 at 0.75 and 1.50).

4.4. Anticipated Competitive Dynamics

4.4.1. Interviewee Articulations

The ‘Competitive pressure from peers for AI adoption’ code appears in 8 of 11 Funds, spanning the maturity distribution: Funds at the upper end (2, 5), middle band (3, 4, 8), and lower end (1, 9, 11). Fund 5 (INT07) captures the dynamic from the upper end: "I think this happens completely automatically. Because if you have conservative funds that have not engaged with this at all, they will simply see – the competition is winning deals, is more efficient, has exclusive situations, based on tools, based on AI. And then it is somehow inevitable, you have your pioneers and the others

simply have to follow, otherwise they cannot keep up" (Fund 5, INT07). Fund 9 (INT13) articulates from the lower end with a specific mechanism: "What helps most, I think, is somehow competitive pressure, because there things always happen. When the senior fund manager suddenly fears that [other fund] is already doing this, then we have it too" (Fund 9, INT13).

4.4.2. Three Sources of Competitive Pressure

The first source of competitive pressure is peer pressure operating through deal sourcing. 'AI value concentrated in deal sourcing' is articulated by 6 of 11 Funds and 'AI potential in deal sourcing' by 8 of 11. Fund 4 (INT06) anchors the value: "I would say most value can be found in sourcing. Because if you can find a company that is not on the radar and then you can do a good deal based on that, that really outdoes any kind of administrative time saved that you would get. I would much rather have one more hour of administrative work each day, but do a great deal at the end of the year that no other player could have found" (Fund 4, INT06). Fund 1 (INT02) articulates competitive pressure operating through this channel: "in the early stages where it is super standardized, [the Fund] would want to use something for competitive advantage. So, you want to have the best methods for finding these companies. And [the Fund] could get pressure from other players, who do it more effectively" (Fund 1, INT02).

The second is LP pressure. Five Funds in the sample carry 'LP pressure for technological adoption' (Funds 3, 4, 5, 8, 11). The carriers articulate LP pressure from three substantive perspectives. Fund 4 articulates where the pressure operates and where it does not: LP pressure operates strongly for reporting, transparency, and data availability – Environmental, Social, Governance (ESG) reporting and key performance indicator (KPI) tracking are described as having been pushed by LPs – but the interviewee draws an explicit limit on the commercial side, observing that "commercial excellence and sourcing edge would really need to go from inside the organization" (Fund 4, INT06). Fund 5 articulates what sharpens the pressure: LPs themselves are using AI to evaluate fund returns against risk-adjusted benchmarks, raising the analytical sophistication of the scrutiny. Fund 8 articulates the response: a proprietary relationship tracking system that "comes across very well" in LP scrutiny (Fund 8, INT10). Funds 3 and 11 carry the code without articulating a distinct pressuring mechanism. Read together, these articulations engage the LP relationship at the level of where pressure operates, what sharpens it, and what holds up under it.

The third is industry-level commoditization. ‘AI intensifying industry competition’ (n=5) co-occurs with ‘Commoditization of expertise through AI tools’ (n=5) at 100% NCOR. Fund 4 makes the connection explicit: "Information will become more symmetric. [...] easier data analysis and more like digestment of public data with AI tools just makes your edge less sharp [...] competition will become more fierce" (Fund 4, INT06). This source is articulated as operating at the industry level rather than the Fund level: interviewees describe expertise commoditization and edge erosion as competitive dynamics that operate beyond any single Fund’s adoption posture.

4.4.3. A Counter-Current and Market Structure Shifts

The competitive pressure analysis above assesses AI adoption as pressure operating on Funds. A counter-position is articulated by Fund 3 (INT05): "The expectation is rather the other way around, that as a Fund you eventually exert pressure on your consultants and say, we provide you with this data room and we expect that by tomorrow morning we have the whole thing fully analyzed. And that only works with AI. Period. And if that does not work then there is a problem. Then they cannot deliver. And then that was their last assignment" (Fund 3, INT05). Fund 3 transmits pressure to consultants rather than receiving it from peer funds.

The ‘AI-driven shift in consultant dependency’ code (n=4) and the ‘Consultants leverage proprietary datasets’ code (n=5) together engage AI dynamics operating in consultant relationships in multiple directions. Fund 4’s articulation that consultants are losing scope to ChatGPT-deep-research-equivalent tools (Fund 4, INT06) engages the disintermediation direction; Fund 5’s note that consultants who already use AI are bringing that capability to their PE clients (Fund 5, INT07) engages the intermediation direction. The counter-position Fund 3 articulates therefore does not strictly contradict the analysis above; it locates the competitive pressure in a different layer of the value chain.

Segment-level competitive dynamics are also articulated, with Mid-Cap consolidation as the specific case Funds engage. Fund 7 (INT09): "There is simply no right to win for a classic generalist mid-cap fund. Large-cap consolidates, runs mid-cap strategies to grow, and absorbs the expertise that way [...] unless they have a strong sector specialization" (Fund 7, INT09). Fund 4 articulates the lateral dynamic across segments: "larger caps start to look at mid-to-large cap, mid-to-large cap starts to look at mid-cap, mid-cap starts to look at smaller cap to get that advantage in a smaller investment" (Fund 4, INT06). Fund 8 articulates the counter-register: Small-Cap deal-

making rests on relationships and negotiation rather than data depth, structurally limiting AI-driven encroachment (Fund 8, INT10). The ‘Competitive spillover across PE market segments’ code is carried by 6 of 11 Funds, with five carriers (Funds 4, 5, 7, 8, 9) also articulating ‘AI as competitive differentiation’ (n=7), yielding 83% NCOR.

4.5. Buy or Build? Substantive AI Investment in the Sample

4.5.1. Strategic and Administrative AI Investment

The ‘Strategic vs. administrative AI value distinction’ code is articulated in 6 of 11 Funds (Funds 2, 4, 5, 8, 9, 11), spanning the maturity distribution. ‘Efficiency gains from AI’ is articulated in all 11 Funds: engagement with AI’s efficiency potential is the shared floor across the sample, and the strategic-administrative distinction is what sits above it. The distinction is between AI investments producing administrative efficiency and AI investments producing strategic differentiation through proprietary knowledge and defensible competitive positioning. Fund 4’s articulation in 4.4.2 anchors this boundary at the LP-pressure level: administrative AI is what external pressure can drive; strategic AI must come from inside the organization.

The two categories map onto different organizational drivers. Administrative AI investment is associated with external triggers – LP-driven reporting, centralized procurement, broadly available third-party tools – while strategic AI investment is associated with internal initiative. ‘Internally driven AI adoption’ (n=4) co-occurs with ‘AI potential for knowledge institutionalization’ at 100% NCOR: Funds 2, 5, 8, and 10 articulate the institutionalization potential. ‘Externally driven AI adoption’ (n=3) co-occurs with ‘Strategic vs. administrative AI value distinction’ at 100% NCOR: Funds 4, 9, and 11 frame the distinction as a limit on what external pressure can drive. The two driver codes do not overlap (0% NCOR).

4.5.2. The Buy-vs-Build Navigation as Analytical Lens

Theme 5 names the knowledge institutionalization gap as the substantive phenomenon; this section engages it through the buy-vs-build navigation. The buy-vs-build codes are housed under Theme 2 as dispositions toward acquiring AI capability (per Section 3.3) and are engaged here as indicators of how each Fund’s buy-vs-build resolution positions it relative to the institutionalization frontier. Three codes capture the buy-vs-build navigation directly: ‘Buy vs build tension’ (n=4, Funds 3, 5, 7, 10), ‘Build-over-buy approach’ (n=5, Funds 3, 5, 6, 7, 10), and

‘Buy-over-build approach’ (n=6, Funds 2, 4, 5, 8, 9, 10). Nine of eleven Funds articulate at least one of the three. The four Funds that explicitly carry ‘Buy vs build tension’ are also build-over-buy carriers (100% NCOR), associating the explicitly articulated tension with build-side preference.

Some Funds carry both build-over-buy and buy-over-build simultaneously. Funds 5 and 10 articulate the navigation as dimension-by-dimension across functional areas rather than as an absolute commitment to either one.

Fund 7 (INT09) articulates the substantive logic for build-side preference at the level of edge protection, framing AI as model-training risk on proprietary data: uploading past investment committee (IC) papers and analytical work to commercial AI tools makes that data available for training models that other PE firms may use, eroding the firm’s analytical edge (Fund 7, INT09).

4.5.3. Build-Side Operationalization

The build resolution that the buy-vs-build navigation engages is documented through ‘AI potential for knowledge institutionalization’, articulated by 7 of 11 Funds (Funds 2, 4, 5, 7, 8, 10, 11). These Funds articulate that AI’s strategic potential lies in capturing Fund-internal knowledge into infrastructure that operates at the firm level. ‘Weak institutional knowledge retention’ (n=4, Funds 4, 8, 9, 11) and ‘Knowledge tied to individuals’ (n=3, Funds 4, 8, 9) name the underlying problem; Fund 9 (INT13) articulates concretely: Fund knowledge sits “in the inbox or brain of a senior”, accessible only through that individual and lost when the individual leaves.

Of the seven Funds articulating ‘AI potential for knowledge institutionalization’, Funds 5, 7, 8, and 10 have built proprietary fund-internal data infrastructure at the operational level. All four carry ‘AI as competitive differentiation’ (4 of 7 sample-wide carriers). Three of four carry ‘Internal data as proprietary moat’ (Funds 5, 7, 8); Fund 10 carries ‘Tech as selling point to LPs’ and articulates the same differentiation logic at the LP-positioning level.

Fund 5 (INT07) has built proprietary AI into its deal tool, automating IC-paper generation, slide construction, and chart creation from fund-internal historical data and templates. The interviewee articulates the within-fund navigation between build and buy directly: "In-house, and the CRM with third-party" (Fund 5, INT07). Proprietary infrastructure was built for the document and IC-paper workflow; commercial CRM tooling covers the relationship-management layer.

Fund 7 (INT09) has built an internal [AI-Retrieval-Tool] that searches across past IC papers, deal histories, and analyst notes within the firm's data perimeter. The interviewee articulates the build-decision rationale at the level of competitive risk: "You don't want to upload all your past IC papers somewhere so that others have them too – so that the LLMs of other PEs are trained on your information" (Fund 7, INT09). The substantive logic engages privacy and Chinese-wall constraints that external commercial tools do not resolve. Fund 7 also articulates a wait-and-see posture on adjacent dimensions where current AI tools do not yet meet performance thresholds (per Section 4.3.2); the dual position documents that build resolution operates dimension-by-dimension within the Fund.

Fund 8 (INT10) has integrated a proprietary LLM layer over the firm's database, allowing investors to query past deal documentation, planning materials, and historical investment context across the firm's accumulated work in natural language. The interviewee articulates what the infrastructure substantively produces: "We had a company recently that the larger fund had looked at eight years ago, and without [tool] or the search across our systems we would not have known or come across it. We could speak to the people who worked on it then [...] and we could see the documents from back then, the planning from back then, and understand what happened in between" (Fund 8, INT10). The substantive logic engages institutional retention as a concrete operational capability – the infrastructure holds firm-specific knowledge that would otherwise exit with personnel turnover.

Fund 10 (INT14) has built a proprietary algorithmic sourcing platform that scores companies across the European target universe against fund-internal investment criteria, plus a proprietary intranet and data integration layer. Three operationally-mandated AI personnel are tasked exclusively with maintaining the proprietary sourcing tool and the intranet. The interviewee articulates what the build operationalization produces at the LP-positioning level: "look, you're gonna own our businesses at some point, whether it's with [competitor] or [competitor], or with us in the stage before – we find all businesses that become interesting, and there we say [our sourcing tool] systematically covers the space [...] to ninety percent" (Fund 10, INT14). The proprietary infrastructure operates as both sourcing edge and fund-positioning currency.

While all four operationalizing Funds score 3 on the Knowledge & Data dimension – consistent with the build operationalization itself engaging the institutionalization frontier – they also distribute substantively across the AI Matrix more broadly. Three of the four sit at upper or

upper-middle tier overall: Fund 5 at 2.50, Fund 10 at 2.25, Fund 8 at 1.75. Fund 7 sits at 1.50 with a 0 on the Reporting & Portfolio dimension – operational mandate concentrated on knowledge-retrieval while adjacent dimensions sit lower, consistent with the within-fund navigation pattern documented in Section 4.3.2.

The four operationalizing Funds all carry operationally-mandated AI personnel (per Section 4.2.1) and built their tools internally.

4.5.4. The Buy-vs-Build Navigation Across the Sample

Engagement with AI through low-barrier commercial tooling – third-party Software-as-a-Service (SaaS), general-purpose AI assistants adopted on non-edge dimensions – is near-universal across the sample and does not by itself constitute a resolution; ‘Efficiency gains from AI’ is articulated by all eleven Funds (Section 4.5.1). The buy-vs-build navigation that distinguishes Funds concerns substantive commitment on edge dimensions: whether a Fund builds proprietary infrastructure, resolves toward a substantive buy, or defers that commitment. The four Funds documented in 4.5.3 operationalize the strategic build-side on edge dimensions. The remaining seven engage the navigation in two distinct sub-patterns: articulating the institutionalization potential without operationalizing it, and remaining outside the institutionalization frontier entirely.

Three Funds – 2, 4, and 11 – articulate ‘AI potential for knowledge institutionalization’ without operationalizing it. All three carry ‘Low-barrier experimentation with AI tools’ alongside the institutionalization-potential articulation; each engages buy-side resolution or wait-and-see as current-state AI investment. Fund 2 sits at upper-tier maturity (2.50 on the matrix) through buy-side breadth – scoring 3 on Sourcing and Due Diligence through commercial tooling but 2 on Knowledge & Data. Its buy-side resolution is deliberate: Fund 2 (INT04) carries ‘Buy-over-build approach’ alongside ‘Team size constraints’, framing a clear SaaS buy-commitment as the rational choice for a VC fund of its size. The institutionalization potential is nonetheless articulated, and located specifically in larger VC funds’ proprietary infrastructure – the segment carrying ‘AI potential for knowledge institutionalization’ also carries ‘Internal data as proprietary moat’ and ‘Scale advantage in larger funds’: "from all this data they have, that can already give you an indication [...] they have these internal databases that create quite a lot of value" (Fund 2, INT04). As a fully buy-committed fund identifying an edge its own approach does not reach, Fund 2 corroborates from the buy-side the ceiling that Section 4.5.3 documents from the build-side. Fund

4 sits at middle-tier (1.25) and articulates the buy-build boundary explicitly: administrative AI is what external pressure drives, the commercial side must come from inside the organization (per Section 4.4.2). Fund 11 sits at low-tier (0.50), carries the wait-and-see posture (per Section 4.3.2), and articulates the gap: "Too resource-constrained to take their head out of the water [...] there's a belief that it will require a lot of tedious work to get everything up and running and the output or the value of it might not be super clear" (Fund 11, INT16). Of the three Funds articulating the potential (Funds 2, 4, and 11), Fund 11 is the sole carrier of foundational deferral, inviting the reading that the potential is noted but, in contrast to the operationally-engaged buy-side positions of Funds 2 and 4, the articulation is made from a deferring posture upstream of resolution. All three articulate the institutionalization potential, but each articulates the non-operationalization from a different underlying posture.

Four Funds sit outside the institutionalization frontier (Funds 1, 3, 6, 9 do not carry 'AI potential for knowledge institutionalization'). Funds 3 and 6 carry 'Build-over-buy approach' as articulated preference without operationalizing; both sit in the foundational deferral cluster (per Section 4.3.2), with build-side preference articulated as aspiration rather than current-state commitment. Fund 9 carries 'Buy-over-build approach' alongside externally-driven adoption, engaging buy-side resolution through commercial tooling within the foundational deferral posture. Fund 1 sits in the granular deferral cluster, articulating tool-level concerns; it carries neither the buy-vs-build codes nor the institutionalization-potential articulation. These four Funds position at low-tier or low-middle-tier in the matrix (0.25 to 1.25), consistent with the absence of build-side operationalization and of the configuration features documented in Section 4.2.

Outcomes do not map cleanly to discrete categories. Funds 5 and 10 carry both build-over-buy and buy-over-build, navigating dimension-by-dimension across functional areas; Fund 8 carries buy-over-build at code level while operationalizing build-side infrastructure on the institutional-knowledge dimension; Fund 7 carries build-over-buy on knowledge retrieval while articulating wait-and-see on adjacent dimensions. Individual Funds navigate the buy-vs-build question across multiple functional dimensions simultaneously, with operational outcomes that do not always align with articulated preferences and with configuration scaffolding that determines whether the navigation can be substantively engaged.

5. Discussion

5.1. Build Where the Edge Is, Buy (Time) Where It Isn't

The sample documents a genuine disagreement on whether building proprietary AI infrastructure is rational in private equity. Fund 4 articulates a buy-side reading directly: no fund will build because it is not an efficient use of resources (Fund 4, INT06). Four Funds in the sample (Funds 5, 7, 8, 10) have built proprietary AI infrastructure on fund-internal data. The disagreement reflects different fund-level resolutions of the buy-vs-build question on AI as competitive capability. The resolutions are not random – they track configuration variance the chapter documents.

Build commitment surfaces on dimensions where edge is concentrated. Fund 10 articulates the constraint against procurement directly: “you can rarely use out-of-the-box solutions because every fund operates differently” (Fund 10, INT14). High-value edge tools require integration with proprietary fund-internal data, organizational workflow specificity, and analytical logic that does not transfer cleanly from commercial offerings. Fund 7 articulates a different constraint on the institutional knowledge dimension: proprietary IC papers cannot be uploaded to external commercial AI tools without exposing the Fund’s competitive content to systems that train competitors’ models. Tool-fit and competitive risk are two different reasons against buy on edge dimensions, both grounded in the substantive content that proprietary fund-internal data engages. This aligns with the knowledge-based view of strategic resources (Grant, 1996), building on the resource-based foundation (Barney, 1991): a capability that integrates firm-specific knowledge with proprietary data and workflow logic is causally ambiguous and difficult to imitate, and cannot be procured externally without losing the integration that makes it valuable. Where edge is not concentrated – CRM tooling, document collaboration, workflow infrastructure – funds buy. Thus, the buy-vs-build resolution operates dimension-by-dimension within the fund: build where edge is concentrated, buy where it is not.

What separates funds that build on edge from funds that do not is organizational commitment to AI as analytical territory. The chapter’s strongest empirical signal of fund-level engagement is operationally-mandated AI personnel (Section 4.2.1): all four operationalizing Funds carry the personnel feature, and the personnel-present versus personnel-absent partition produces the strongest gradient observed on the AI Matrix. The resource-based and capability-

development literature frames productive resources as something that does not deploy itself (Penrose, 1959; Teece, 1986). Read through the sequenced commitment dynamic set out above, this empirical pattern marks the visible trace of an upstream process: prior strategic commitment to deploying capital toward capability, then organizational architecture aligned with that commitment, then operational scaffolding sustained over time. The personnel correlation observable in the cross-case data is the visible portion of that process – a commitment the methodology cannot directly observe: commitment to AI as analytical territory shaped the hiring decision, the hiring decision materialized as the personnel feature, and the hired personnel built the infrastructure the chapter documents. Fund 10's three operationally-mandated AI personnel maintaining proprietary algorithmic infrastructure, and the upper-tier maturity score that follows, illustrate that dynamic.

Funds that defer build commitment on edge dimensions are not making a lesser decision; they continue to engage AI through commercial tooling at low-barrier level on non-edge dimensions while deferring substantive build investment where uncertainty is highest. Under Knightian uncertainty (Knight, 1921) – uncertainty whose probability distribution cannot be calculated – committing capital to build infrastructure that may be obsolete in two years is itself irrational. Waiting holds option value (Dixit and Pindyck, 1994): the value of deferring depends on how much new information arrives during the wait, and on whether the wait can be exercised when conditions clarify. Disciplined deferral operates within an equilibrium: when funds heterogeneous in their operational setup face common technological uncertainty over AI capability evolution, equilibrium delay can exceed what partial-equilibrium analysis predicts because each fund's deferral decision is informed by anticipation of other funds' eventual investments (Novy-Marx, 2007). The substantive conditions reinforce this dynamic – where ROI from premature AI commitment is unclear and operational tools are not yet at the level required, disciplined deferral allows the option value of waiting to be preserved while fund-level heterogeneity, applied to this thesis, makes that deferral a stable equilibrium rather than a competitively forced position. Whether the option is actually exercisable is what divides the deferring Funds in the sample (Section 4.3). Granular-cluster Funds (Funds 1, 7) defer on implementation where current tool-level performance does not yet justify build commitment, while the organizational conditions to exercise the option when conditions clarify are in place. Fund 7 illustrates this dimension-by-dimension navigation: build operationalization on knowledge retrieval where edge is clear,

deferral on adjacent dimensions where current tool performance has not yet clarified. This is informed waiting – option value assessed and preserved. Foundational cluster Funds (Funds 3, 6, 9, 11) defer upstream of evaluation: the organizational commitment that produces substantive engagement is not yet in place, though the wait-and-see articulation indicates these Funds recognize AI as carrying value worth waiting for rather than dismissing it as without value. Fund 11 articulates this from inside: the Fund is too resource-constrained to engage AI substantively, anticipates the implementation work as tedious and the value as unclear (Fund 11, INT16). This is structural rather than strategic deferral – Knightian uncertainty without the conditions to exercise the option when it becomes in-the-money. Both clusters defer build commitment on edge dimensions, but only the granular cluster could exercise the option when conditions clarify.

The disagreement in the sample can be resolved. Fund 4's articulation that no fund will build is one rational reading among several that the sample documents. The configuration variance the chapter examines operates within a broader empirical pattern: at the firm-economy level, AI-driven growth concentrates among more analytically capable firms (Babina et al., 2024). The chapter's findings operationalize this pattern at fund level: where commitment is in place and edge is concentrated, funds build; where it is not, they buy or defer. Under the option-value and resource-based-view, all three resolutions are rational under their respective conditions; the apparent contradiction reflects configuration variance, not different views of what is efficient.

5.2. Customer Relationships: AI and the Brand-Name Threshold

The competitive-capability reading developed in Section 5.1 extends to the customer-relationship dimension that determines capital flow. Fund 11 articulates the LP relationship as a customer relationship – "you also need to think about, like, who's the customer. And it's essentially LPs. Like, that's the ones you're working for" (Fund 11, INT16). Fund 3 articulates the consequence: LP capital flight produces the same end state as competitive obsolescence – the result is the same whether competitors advance or LPs walk away. The question this section takes up follows from Section 4.4.2's substantive perspectives: if LP pressure is shifting in the direction the carriers articulate, what threshold approaches going forward, and what does crossing it mean for capital allocation?

The codes describe a shift in transit. Fund 4 marks the present-day equilibrium: pressure operates strongly on reporting, transparency, and ESG, but commercial-edge investment "would

really need to go from inside the organization" (Fund 4, INT06). Fund 5 names what makes the equilibrium unstable – LPs themselves layering AI onto fund evaluation, screening returns against risk and against deal flow. Fund 8 illustrates the response side: proprietary infrastructure registering positively in LP scrutiny without LPs explicitly demanding it. Fund 10 articulates the substantive economic argument the response operationalizes: "you're gonna own our businesses at some point, whether it's with [competitor] or [competitor], or with us in the stage before – we find all businesses that become interesting, and there we say [our sourcing tool] systematically covers the space [...] to ninety percent" (Fund 10, INT14) – earlier-stage capture of the same eventual portfolio companies through systematic algorithmic coverage of the European target universe, pitched to LPs as the differentiation that justifies the allocation. The four vantage points span the institutional-pressure mechanisms DiMaggio and Powell (1983) identify as coercive, mimetic, and normative, with Fund 10's articulation extending the spectrum into substantive economic positioning. Westphal and Zajac (2001) provide the framework for reading these vantage points together: as LP analytical capacity rises, the symbolic-versus-substantive distinction becomes legible at the customer-relationship level rather than only at the firm-structure level.

Where the shift could converge follows from the strategic-versus-administrative distinction documented in Section 4.5.1, which co-occurs with externally-driven adoption at 100% NCOR (Funds 4, 9, 11). Administrative AI infrastructure becomes table stakes – necessary for fundraising but no longer differentiating – while strategic AI infrastructure could become the dimension on which evaluative judgment turns. The symbolic-substantive frame (Westphal and Zajac, 2001) surfaces a failure mode: Fund 3 names "greenwash hires" – funds hiring AI personnel because LP perception rewards the signal without substantive AI investment behind it. As LPs develop the analytical capacity to distinguish substantive infrastructure from compliance hires, the symbolic-compliance response becomes harder to sustain at the LP-evaluation level.

The shift, however, runs against an opposing reading of LP behavior that the data supports equally well. Four Funds articulate fundraising scarcity as the present-day market reality (Funds 3, 5, 8, 9) – the structural backdrop that makes the conflict consequential. Fund 8 names the opposing reading directly: "especially when capital is scarce, LPs default to top-10 names because nobody is going to tell you 'hey, you invested in [top-10 fund], that's a [expletive] idea'" (Fund 8, INT10). Under capital scarcity, LP behavior is structurally conservative – capital concentrates with brand-name funds because the allocation is reputationally defensible regardless of substantive

differentiation. The same Fund holds both readings in adjacent passages of the same interview: analytical-capacity rise destabilizing the equilibrium, and brand-name protection insulating top-tier funds. Akerlof's (1970) information-asymmetry framework and Jensen and Meckling's (1976) principal-agent dynamics surface what the conflict is substantively about – when both substantive infrastructure and reputational proxies signal capability to capital allocators, which signal LP evaluation weights more heavily, and under what conditions the weighting shifts. The question this section engages is which reading dominates as the dynamic plays out.

The conflict could resolve asymmetrically across the configuration distribution. Brand-name protection insulates top-tier funds from the LP-evaluation threshold shift – their reputational defensibility holds regardless of tech infrastructure. Non-brand-name funds, where the foundational-deferral cluster (Funds 3, 6, 9, 11) sits, face the threshold from the more exposed side, because tech-native differentiation is one mechanism that could break the brand-name allocation default under capital scarcity. Within-sample evidence grounds the asymmetric resolution: two of the build-resolution Funds – Funds 8 and 10 – surfaced tech-nativeness as the mechanism by which they pitch differentiation to LPs. They are not waiting for the LP-evaluation threshold to shift; they are constructing it. Foundational-deferral Funds in the non-brand-name segment would face increasingly costly catch-up dynamics if LP analytical capability rises. As LPs develop the capacity to evaluate fund AI capability substantively, the institutional-investor pressure Aggarwal et al. (2011) document operates here at the LP-PE-fund level – convergence pressure concentrates below the brand-name threshold, where reputational defensibility does not insulate funds from substantive evaluation.

LP evolution operates as the enforcement mechanism that converts 5.1's configuration variance into capital-allocation consequence, but the enforcement is segment-conditional, operating below the top-tier threshold rather than industry-wide. The disagreement on whether LPs reorient toward tech-native funds or default to brand names does not resolve into a single industry-wide answer; it resolves into asymmetric pressure that operates differently across the configuration distribution the chapter has documented.

5.3. Human Capital: Where Edge Survives Commoditization

The readings developed in Sections 5.1 and 5.2 extend to the human-capital dimension that determines what PE work substantively becomes as AI restructures the analytical layer. The cross-

case sample articulates this dimension across two registers: what AI commoditizes, and what AI does not. The question this section takes up follows from the substantive perspectives Funds articulate in Section 4.5: if AI productivity gains compress the analytical work historically performed by juniors, what survives commoditization, and how does the configuration variance the chapter has documented map onto the differential exposure that follows?

The substantive shift Funds articulate concentrates on the analytical layer. Five Funds across the sample engage the displacement reading. Fund 11 articulates the long-horizon position from the Small-Cap side: "I don't see analysts and definitely not interns as very useful in a couple of years. What an analyst does or what an intern does is very easy to automate" (Fund 11, INT16). Fund 7 articulates a stronger displacement reading at shorter horizon, with the substantive mechanism that AI commoditizes the screening, business-description, financial-extraction, and modelling work that constitutes junior-analyst output (Fund 7, INT09). Fund 5 articulates the productivity-expansion variant: rather than displacement, the analyst of the future concentrates on the central, important workflows while the routine analytical layer is taken over by AI tools (Fund 5, INT07). The displacement and productivity-expansion readings are not mutually exclusive – they engage two registers of skill-biased technological change, substitution for routine cognitive tasks and complementarity with non-routine judgment-intensive work (Autor et al., 2003; Bresnahan et al., 2002), and the cross-case sample articulates both registers simultaneously rather than converging on a single trajectory.

What survives commoditization operates on two distinct dimensions. The firm-internal capability dimension Section 5.1 documents – proprietary infrastructure, fund-specific knowledge integration, asset-specific operational substrate – is one source of edge that AI commoditization does not erode. The second dimension operates across organizational boundaries: as analytical expertise commoditizes, what remains substantively valuable is relationships and commercial sense. Fund 4 articulates this directly: "expertise gets commoditized [...] What becomes even more important is relationships [...] through relationships, you can also build an edge in the process" (Fund 4, INT06). Fund 7 names the same dimension from the operational side: "your network is your net worth in this industry", with "commercial sense even more important" as the driver of returns once analytical work commoditizes (Fund 7, INT09). The relational view extends the resource-based foundation Section 2.2 establishes: relational rents arise across organizational boundaries through inter-organizational relationships rather than through firm-internal resources

alone (Dyer and Singh, 1998). The framework reads relationships as a theoretically distinct source of competitive advantage that AI commoditization does not reduce, complementary to the firm-internal capability dimension Section 5.1 establishes. Beyond the cross-case sample, expert input from a senior dedicated AI professional at a Large-Cap PE Fund (EXP 1, INT01) articulates the structural consequence directly through a pyramid-to-diamond reading of the PE workforce: the analytical layer compresses while judgment-carrying senior capacity expands proportionally. AI exposure varies systematically across occupations and tasks at firm level (Acemoglu et al., 2022); the pyramid-to-diamond framing operationalizes this variance for PE workforce structure.

The same expert source (EXP 1, INT01) articulates the framework that captures what happens to deal volume under AI productivity gains. Asked about the effect on PE workforce, EXP01 invoked the Jevons paradox – the observation originating in Jevons' (1865) analysis of coal that efficiency gains expand rather than contract demand for the underlying input because lower per-unit cost makes more activity economically viable. Applied to PE, the substantive mechanism is that AI productivity per analyst-hour expands the volume of opportunities funds can substantively evaluate; higher value-creation capabilities let funds participate in more processes, hold larger active opportunity sets, and engage targets earlier in their commercial trajectory. Fund 8 articulates this from inside an operationalizing Fund: with AI augmentation analysts will generate significantly more output, becoming productive at scale rather than displaced (Fund 8, INT10).

The resolutions configuration variance produces map onto the human-capital dimension as differential exposure. Build-resolution Funds technologically scale relationship management at Fund-level: Fund 10's systematic touchpoint coverage, Fund 8's two-year cumulative tracking of every asset, advisor, and industry expert, and Fund 7's institutional databases all institutionalize relational value at scale rather than leaving it dependent on individual personnel. The relationship dimension that commoditization does not erode is captured by build-resolution Funds through the same configurational logic Section 5.1 establishes. Foundational-deferral Funds face the opposite condition: the talent scarcity Fund 3 articulates from inside the cluster – implementation talent is rare, the people who can build from scratch are hard to find, and the cost is substantial (Fund 3, INT05). Build-resolution Funds capture the productivity expansion and access AI implementation talent earlier; foundational-deferral Funds face the labor-market dynamics from the more vulnerable side. The human-capital disadvantage compounds the catch-up costs the deferral cluster

already faces on the capability (Section 5.1) and LP-evaluation (Section 5.2) dimensions – the same resolutions, read across three dimensions, produce one reinforcing rather than independent disadvantage.

6. Conclusion

The eleven Funds share key structural features yet vary widely in how far they have built AI capability. Where each Fund sits on the maturity distribution does not track ticket size, geography, sector focus, LP profile, or interviewee seniority. What it does track is configuration variance: operationally-mandated AI personnel as the primary feature (Section 4.2.1), with bottom-up AI initiative and active-governance signature as amplifiers (Sections 4.2.2–4.2.3). These observable features follow downstream from a Fund’s commitment to AI as analytical territory – the upstream disposition that shapes whether it engages substantive AI investment at all. Across the sample, the AI-adoption decision resolves along three characteristic lines. Four Funds (5, 7, 8, 10) operationalized build-side infrastructure on the dimensions where edge is concentrated; two (2, 4) articulate the institutionalization potential without building; four (3, 6, 9, 11) defer upstream of evaluation, lacking the infrastructure to assess AI’s potential value; Fund 1 defers at the granular tool level (per Section 4.3.2). These are resolutions Funds occupy, not types Funds belong to: the buy-vs-build question is navigated dimension-by-dimension, so a Fund’s placement names its primary resolution while other dimensions might resolve differently. The Discussion reads these resolutions across three dimensions of differential exposure – competitive capability (5.1), LP relationship (5.2), and human capital (5.3). What follows extends the reading to a fourth dimension, the value chain, and projects what these resolutions imply for segment dynamics, capital allocation, and labor-market trajectories.

The advisor layer is where the resolutions Funds occupy produce their sharpest forward projection. A senior dedicated AI professional at a Large-Cap PE Fund (EXP 1, INT01) describes a directive that captures the trajectory: one of the Fund’s offices asked the internal AI team how to replace its commercial due diligence provider entirely. What Fund 4 articulated as a future possibility, consultancy work shrinking as AI tools internalize the analytical scope previously bought from outside, EXP01 reports as an already-issued operational mandate. Build-resolution Funds are executing this displacement; the foundational-deferral cluster remains dependent on the advisor layer those Funds are exiting. Fund 3 articulates a position one step earlier: pressure

transmission to its consultants for AI-driven deliverables as a precondition for engagement (Fund 3, INT05). A counter-reading runs in the opposite direction: Funds 5, 7, and 9 describe consultants moving up the AI capability curve faster than their PE clients, with proprietary datasets and AI-augmented deliverables Mid-Cap funds cannot easily replicate.

The same logic projects forward to segment structure. Fund 4 captures the substantive driver of Large-Cap encroachment into smaller deal sizes: "having a commercial understanding that is edgy [...] is harder [in Large-Cap, with public data], which is why larger caps start to look at mid-to-large cap, mid-to-large cap starts to look at mid-cap, mid-cap starts to look at smaller cap to get that advantage in a smaller investment" (Fund 4, INT06). Two mechanisms enable the trajectory: AI efficiency rises with data volume (Fund 3, INT05), and the management-fee economics of larger funds make the AI infrastructure affordable in the first place (Fund 8, INT10). On top of these, the Jevons paradox operates – efficiency gains do not reduce headcount but expand the opportunities each analyst can substantively evaluate (EXP 1, INT01) – making smaller-ticket deals previously unprofitable to evaluate at scale economically viable at the AI-mature Large-Cap frontier (Jevons, 1865). The structural-elimination finding from Section 4.1.3 remains intact because the trajectory operates on the substrate of deal-making rather than ticket size. A counter-reading is articulated by Funds closer to that substrate: Fund 8 holds that Small-Cap deal-making runs on relationships, negotiation, and contacts rather than data depth, and that "Tech here cannot completely eliminate the small ones" (Fund 8, INT10); sector specialization, according to Funds 5 and 7, is the surviving niche where generalist Mid-Cap consolidates upward. The forward dynamic is conditional on both configuration and deal substrate: Funds with both AI infrastructure and a data-rich substrate engage smaller segments through AI-enabled efficiency, while Funds in relationship-driven segments are structurally insulated from this kind of encroachment.

The capital-allocation dimension extends Section 5.2's argument forward. As 5.2 showed, brand-name protection insulates top-tier funds from rising LP analytical capacity. The tech-native Funds in the sample (Funds 8 and 10 most explicitly) are not waiting for the LP-evaluation threshold to shift; they are constructing the criteria by which it would operate, pitching tech-nativeness as the differentiation that justifies allocation under capital scarcity. The labor-market dimension Section 5.3 develops projects forward through the same logic. Jevons productivity expansion lets analysts evaluate more opportunities at constant headcount, but capturing that expansion requires the build infrastructure that absorbs the additional throughput. Build-resolution

Funds capture both the productivity expansion and the institutionalized relational infrastructure Section 5.3 documents; the foundational-deferral cluster faces the implementation-talent scarcity Fund 3 articulates from inside, compounding catch-up costs across every other dimension.

Read across the four dimensions – capability, customer relationship, human capital, and value chain – the build-resolution Funds capture efficiency gains, access new segments, absorb labor-market restructuring as productivity expansion, and displace advisor scope; the foundational-deferral Funds face commoditization, manufactured pressure, advisor dependency, segment encroachment, and labor-market displacement as reinforcing rather than independent dynamics. Configuration variance is not one organizational dimension among several. It is the dimension the chapter’s forward projection finds consequential across each register of differential exposure.

Of the limitations detailed in Section 3.5, two suggest direct directions for future research. The sample’s European concentration invites geographic comparative work extending the analysis to US or Asian samples. Junior-skewed interviewee seniority means strategic decision-rationale is predominantly observed one step below the GP and Partner level; for Junior-represented Funds, replication at GP-level would test whether the configuration-variance reading holds when surfaced from where the upstream commitment is made. A third direction follows from the analytical scope of the thesis: pre-investment dealmaking was the workflow axis comparable across all eleven cases (per Section 3.2), and extending the analysis to portfolio-company AI adoption would connect fund-level configuration variance to portfolio outcomes (Baik et al., 2024) – the channel through which Kaplan and Strömberg’s (2009) operational-engineering paradigm of contemporary PE value creation would operate under AI dynamics.

References

- Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1), S293–S340. <https://doi.org/10.1086/718327>
- Acharya, V. V., Gottschalg, O. F., Hahn, M., & Kehoe, C. (2013). Corporate governance and value creation: Evidence from private equity. *Review of Financial Studies*, 26(2), 368–402. <https://doi.org/10.1093/rfs/hhs117>
- Aggarwal, R., Erel, I., Ferreira, M., & Matos, P. (2011). Does governance travel around the world? Evidence from institutional investors. *Journal of Financial Economics*, 100(1), 154–181. <https://doi.org/10.1016/j.jfineco.2010.10.018>
- Agrawal, A., & Tambe, P. (2016). Private equity and workers' career paths: The role of technological change. *Review of Financial Studies*, 29(9), 2455–2489. <https://doi.org/10.1093/rfs/hhw025>
- Akerlof, G. A. (1970). The market for “lemons”: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488–500. <https://doi.org/10.2307/1879431>
- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4), 1279–1333. <https://doi.org/10.1162/003355303322552801>
- Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151, 103745. <https://doi.org/10.1016/j.jfineco.2023.103745>
- Baik, B. K., Chen, W., & Srinivasan, S. (2024). *Private equity and the adoption of digital technologies* (Harvard Business School Accounting & Management Unit Working Paper Nos. 24–070). Harvard Business School. <https://doi.org/10.2139/ssrn.4885664>
- Barney, J. B. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1), 99–120. <https://doi.org/10.1177/014920639101700108>
- Bernstein, S., & Sheen, A. (2016). The operational consequences of private equity buyouts: Evidence from the restaurant industry. *Review of Financial Studies*, 29(9), 2387–2418. <https://doi.org/10.1093/rfs/hhw037>
- Bogner, A., Littig, B., & Menz, W. (Eds.). (2009). *Interviewing experts*. Palgrave Macmillan.
- Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *Quarterly Journal of Economics*, 117(1), 339–376. <https://doi.org/10.1162/003355302753399526>

- Brynjolfsson, E., & Hitt, L. M. (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4), 23–48. <https://doi.org/10.1257/jep.14.4.23>
- Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1), 333–372. <https://doi.org/10.1257/mac.20180386>
- Davis, S. J., Haltiwanger, J., Handley, K., Jarmin, R., Lerner, J., & Miranda, J. (2014). Private equity, jobs, and productivity. *American Economic Review*, 104(12), 3956–3990. <https://doi.org/10.1257/aer.104.12.3956>
- DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147–160. <https://doi.org/10.2307/2095101>
- Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press.
- Dyer, J. H., & Singh, H. (1998). The relational view: Cooperative strategy and sources of interorganizational competitive advantage. *Academy of Management Review*, 23(4), 660–679. <https://doi.org/10.2307/259056>
- Edmondson, A. C., & McManus, S. E. (2007). Methodological fit in management field research. *Academy of Management Review*, 32(4), 1246–1264. <https://doi.org/10.5465/amr.2007.26586086>
- Eisenhardt, K. M. (1989). Building theories from case study research. *Academy of Management Review*, 14(4), 532–550. <https://doi.org/10.5465/amr.1989.4308385>
- Eisenhardt, K. M., & Graebner, M. E. (2007). Theory building from cases: Opportunities and challenges. *Academy of Management Journal*, 50(1), 25–32. <https://doi.org/10.5465/amj.2007.24160888>
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking qualitative rigor in inductive research: Notes on the Gioia methodology. *Organizational Research Methods*, 16(1), 15–31. <https://doi.org/10.1177/1094428112452151>
- Glaser, B. G., & Strauss, A. L. (1967). *The discovery of grounded theory: Strategies for qualitative research*. Aldine Publishing Company.
- Grant, R. M. (1996). Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17(S2), 109–122. <https://doi.org/10.1002/smj.4250171110>
- Guest, G., Bunce, A., & Johnson, L. (2006). How many interviews are enough?: An experiment with data saturation and variability. *Field Methods*, 18(1), 59–82. <https://doi.org/10.1177/1525822X05279903>

- Jensen, M. C., & Meckling, W. H. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4), 305–360. [https://doi.org/10.1016/0304-405X\(76\)90026-X](https://doi.org/10.1016/0304-405X(76)90026-X)
- Jevons, W. S. (1865). *The coal question: An inquiry concerning the progress of the nation, and the probable exhaustion of our coal-mines* (1st ed.). Macmillan and Co.
- Kaplan, S. N., & Strömberg, P. (2009). Leveraged buyouts and private equity. *Journal of Economic Perspectives*, 23(1), 121–146. <https://doi.org/10.1257/jep.23.1.121>
- Knight, F. H. (1921). *Risk, uncertainty and profit*. Houghton Mifflin.
- Kvale, S., & Brinkmann, S. (2009). *InterViews: Learning the craft of qualitative research interviewing* (2nd ed.). Sage Publications.
- Lerner, J., Sorensen, M., & Strömberg, P. (2011). Private equity and long-run investment: The case of innovation. *Journal of Finance*, 66(2), 445–477. <https://doi.org/10.1111/j.1540-6261.2010.01639.x>
- Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. Sage Publications.
- Meuser, M., & Nagel, U. (1991). ExpertInneninterviews—Vielfach erprobt, wenig bedacht. In D. Garz & K. Kraimer (Eds.), *Qualitativ-empirische Sozialforschung: Konzepte, Methoden, Analysen* (pp. 441–471). VS Verlag für Sozialwissenschaften. https://doi.org/10.1007/978-3-322-97024-4_14
- Novy-Marx, R. (2007). An equilibrium model of investment under uncertainty. *Review of Financial Studies*, 20(5), 1461–1502. <https://doi.org/10.1093/rfs/hhm013>
- Penrose, E. T. (1959). *The theory of the growth of the firm*. Oxford University Press.
- Rubin, H. J., & Rubin, I. S. (2005). *Qualitative interviewing: The art of hearing data* (2nd ed.). Sage Publications.
- Siggelkow, N. (2007). Persuasion with case studies. *Academy of Management Journal*, 50(1), 20–24. <https://doi.org/10.5465/amj.2007.24160882>
- Strauss, A. L., & Corbin, J. M. (1998). *Basics of qualitative research: Techniques and procedures for developing grounded theory* (2nd ed.). Sage Publications.
- Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy*, 15(6), 285–305. [https://doi.org/10.1016/0048-7333\(86\)90027-2](https://doi.org/10.1016/0048-7333(86)90027-2)
- Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509–533. [https://doi.org/10.1002/\(SICI\)1097-0266\(199708\)18:7%3C509::AID-SMJ882%3E3.0.CO;2-Z](https://doi.org/10.1002/(SICI)1097-0266(199708)18:7%3C509::AID-SMJ882%3E3.0.CO;2-Z)

Westphal, J. D., & Zajac, E. J. (2001). Decoupling policy from practice: The case of stock repurchase programs. *Administrative Science Quarterly*, 46(2), 202–228.
<https://doi.org/10.2307/2667086>

Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). Sage Publications.

Appendix

Table 2: Initial questionnaire catalog and thematic interview structure. Thematic blocks structuring semi-structured interviews; questions applied flexibly per Section 3.2.

Personal	Name	
	Company	
	Position	
Characteristics	Deal Flow: type, quantity, industry, size (LC, MC, SC)	
	LP Profile	
	Where are you on the tech adoption curve vs peers? (1–10 self-rating + “why that score?”)	
	Do you use technology in these workflows? If yes, what and how well does it work? Concrete KPIs? If not, why? What’s next?	Sourcing
		Screening
		Due Diligence
		Value Creation / Post-Merger Integration
		Monitoring / Reporting
		Exit
	What is your sophistication of technology integration? Discuss Tools, Data, AI, Operating Systems, Inhouse vs. Outsourced, On premise vs. Cloud integrations	
	Who owns tech decisions? (Partner, Ops, Deal team, external providers)?	
	What roles did you add/change because of technology requirements?	
	What is the biggest internal blocker for technology?: Incentives, time, trust, security, data, cost, or governance?	
	What skills will a junior/senior need in 5–10 years? What skills become obsolete?	
SWOT	Strengths	What technologies matter most for your firm/investments right now?
	Weaknesses in Data	What data do you have access to in your Fund?

		What's your biggest data pain: quality, permissions, integration, homogeneity?
	Opportunities & Importance	Biggest Value Potential across Lifecycle?
		Do you think PE Funds will build vs buy vs partner and see tech as assistance or automatization? What's the decision rules?
		How important is technology adoption to you? Priority? Willingness to pay? Operating vs. Deal expenses?
		Walk me through the most meaningful/recent technology investment, why did it matter, what was the result?
	Threats	Your view on the future of technology, how it shapes market dynamics (competitive, consolidation)
		Is there a consolidation trend ongoing in the PE market? Is Tech adoption having an influence on that?
		Do you think Fund size matters in how well tech works?
		Are smaller Funds lagging behind? If yes, do you know why?
Stakeholder Dynamics	How do the following influence your technology strategy?	Other Funds
		LPs (explicit requirements?, Fundraising outcomes?)
		Consultants
		Other
Divergence Anomalies	You haven't adopted [technology], why did you not yet do so, given that others have?	
	Should we talk to someone else in your network?	

Table 3: Sample Overview. Per-Fund disclosure of interview ID, Fund pseudonym, ticket size category, and interviewee seniority for the eleven analytical-sample interviews. Geography, sector focus, and LP profile reported only as aggregate distributions for anonymization (Section 3.2).

Fund	Interview ID	Ticket Size	Interviewee seniority
Fund 1	INT02	Growth	Junior
Fund 2	INT04	VC	Junior
Fund 3	INT05	Large-Cap	Junior
Fund 4	INT06	Mid-Cap	Senior
Fund 5	INT07	Mid-Cap	Junior
Fund 6	INT08	Small-Cap	Senior
Fund 7	INT09	Large-Cap	Junior
Fund 8	INT10	Mid-Cap	Junior
Fund 9	INT13	Large-Cap	Junior
Fund 10	INT14	Growth	Senior
Fund 11	INT16	Small-Cap	Junior

Additional factors

Office geography	Geography 1, Geography 2, Geography 3
Sector focus	Technology-focused, Diversified, Non-technology-focused
LP profile coverage	Primarily institutional, Primarily private

Table 4: Expert Sample Overview. The five expert interviews conducted during fieldwork outside the eleven-interview cross-case analytical sample. INT01 (EXP 1) is engaged as expert-informant material in the Discussion and Conclusion; the remaining four are retained as fieldwork context only. See Section 3.2 for the sampling rationale and exclusion criteria.

Name	Interview ID	Profile	Interviewee seniority
EXP 1	INT01	AI Professional, Large-Cap PE	Senior
EXP 2	INT12	Consultant	Junior
EXP 3	INT11	Investment Banking Professional	Junior
EXP 4	INT15	Technology Professional	Junior
EXP 5	INT03	Fund-of-Funds Professional	Junior

Table 5: Detailed Analytical Journey per Phase. Documentation of the coding apparatus from open coding through theme construction; columns name the procedural step, methodological grounding, and output of each phase.

Phase	Procedural Step	Methodological Grounding	Output
1. Open Coding	Pseudonymized transcripts imported into MAXQDA and open-coded through a semi-manual hybrid workflow. MAXQDA AI assistance suggested candidate code names and definitions; every suggestion was reviewed and adjudicated by the researchers against the empirical material. AI assistance was confined to this phase.	Strauss and Corbin (1998) grounded-theory-inspired approach. Inductive, data-near coding preserving empirical nuance before imposing higher-order abstraction.	Approximately 200 first-order codes closely grounded in interview language.
2. Axial Coding	Initial codes consolidated and reorganized into a smaller set of higher-order categories. Codes capturing one underlying concept under different surface formulations were merged into consolidated codes; merged codes carry the ‘+’ marker in the final code list, preserving traceability to their first-order constituents.	Strauss and Corbin (1998). Constant comparative analysis (Glaser and Strauss, 1967) operating across all transcripts.	Reduced codebook with consolidated higher-order categories.
3. Selective Coding	Core categories and the narrative logic linking themes across cases were identified. Roughly 20 interview-specific codes not connecting to cross-case themes were moved to a Descriptive Archive bucket within the MAXQDA project, separating them from the analytical codebook while keeping them in the project’s audit trail.	Strauss and Corbin (1998). Progression from descriptive coding toward coherent analytical structure used in the Findings chapter.	Final analytical codebook of 162 leaf-level codes (Table 6).
Cross-Interview Audit	Structured per-transcript audit checking two things: whether codes that emerged in later interviews also applied to earlier transcripts, and whether codes already attached to a transcript sat on the correct passages. Categories were continuously compared across interviews and against one another.	Constant comparative analysis (Glaser and Strauss, 1967); the move from first-order informant concepts toward second-order themes and aggregate dimensions (Gioia et al., 2013).	Post-audit codebook (Table 6) – basis for all cross-case analyses reported in the Findings chapter.

Phase	Procedural Step	Methodological Grounding	Output
Theme Construction	All post-audit codebook codes were clustered into higher-level themes through inductive clustering, yielding five core themes. Theme assignment is non-exclusive: a code touching multiple analytical domains is listed under each relevant theme. Dual-themed codes are resolved to the unique-code level for cross-theme analysis.	Inductive theme emergence consistent with grounded theory-inspired approach. Methodological discipline against double-counting in cross-theme analysis.	Five-theme structure (Table 11).
Matrix Construction	Codes engaging the operational status quo of AI use were consolidated into the AI Matrix, structured on the interview questionnaire's coverage of the PE investment workflow and fund operating systems. Four operational dimensions were scored on a 0–3 scale, with anchors calibrated from the most operationally articulated transcript examples rather than a pre-specified rubric; the per-dimension scores yield a composite mean. AI integration is scored specifically – technology presence without AI integration scores no higher than basic workflow technology.	Input/output discipline: configuration features (personnel, governance, culture) are treated as inputs (Section 4.2), not scored on the matrix. Scoring them on the same instrument as workflow outcomes would make the construction circular – explanans and explanandum sharing one scale.	Operational AI Maturity Matrix (Table 1). Rubric and per-Fund scoring justifications (Tables 7–10).
Normalized Co-Occurrence Ratio (NCOR)	Co-occurrence instrument capturing how strongly two codes co-occur on the same Funds, scaled against the maximum possible given each code's carrier count. Definition and Equation 1 are given in Section 3.3.	Set-theoretic overlap measure. A descriptive proportion, not an inferential statistic: with N=11 it carries no conventional significance, and joint articulation does not establish causal direction.	NCOR values reported alongside carrier counts of both codes in the Findings chapter, allowing readers to assess the strength of the containment claim against the underlying sample sizes (Table 12).
Code Treatment Framework	Codes are treated as interpretive analytical categories constructed from the interview material, not as quantitative measurements. A code's presence on a transcript reflects whether the concept it names was articulated within the code's codebook-defined scope. Cross-case analysis operates on binary code coverage rather than code frequency.	Strauss and Corbin (1998); Glaser and Strauss (1967). Constant comparative logic refining codes across the coding phases into interpretive categories rather than quantitative measures.	Cross-case analytical claims rest on cross-case carrier patterns rather than within-interview intensity of discussion topics.

Table 6: Post-Audit Codebook. Binarized Fund-Coverage Matrix for the 162 leaf-level codes constituting the analytical codebook in Alphabetic Order. Cell value = 1 if code carried by Fund, 0 otherwise. Sum column reports each code's carrier count across the eleven Funds.

Code	Fund 11	Fund 6	Fund 10	Fund 1	Fund 5	Fund 8	Fund 4	Fund 9	Fund 7	Fund 3	Fund 2	Sum
Absence of dedicated technology personnel	1	1	0	0	0	0	1	1	0	0	0	4
AI as competitive differentiation	0	0	1	0	1	1	1	1	1	0	1	7
AI automation trajectory	1	0	1	0	0	1	0	1	1	0	1	6
AI call notetaker	1	0	0	1	1	1	0	0	0	0	1	5
AI decision guidance rather than full automation	0	1	0	0	1	0	0	1	1	0	0	4
AI edge requires empirical testing to validate	0	0	0	0	0	0	0	0	0	1	0	1
AI efficiency scales with data volume	1	1	1	0	0	1	0	0	1	1	0	6
AI hype skepticism	0	0	0	0	0	0	0	0	0	1	0	1
AI implementation in automotive sector	0	1	0	0	0	0	0	0	0	0	0	1
AI intensifying industry competition	1	0	0	0	1	1	1	0	1	0	0	5
AI potential for knowledge institutionalization	1	0	1	0	1	1	1	0	1	0	1	7
AI potential in deal sourcing	1	0	1	1	1	1	1	0	0	1	1	8
AI potential in specific workflow stages	1	1	0	0	1	0	0	1	1	1	0	6
AI value concentrated in deal sourcing	1	0	1	0	1	1	1	0	0	0	1	6
AI-connected knowledge retrieval	0	0	0	0	0	1	0	0	1	0	0	2
AI-driven shift in consultant dependency	1	0	0	0	0	0	1	0	0	1	1	4
Analysts will be fully replaced by AI in the future	1	0	0	0	0	0	0	1	1	0	0	3
Anti-ideological tech implementation in PortCos	0	1	0	0	0	0	1	1	1	0	0	4
Automated portfolio financial reporting	0	0	1	0	1	0	0	0	0	0	1	3
Automated sentiment analysis of reference calls	0	0	0	0	0	0	0	0	0	0	1	1
Automated Slide generation potential	1	0	0	0	0	0	0	1	0	1	0	3
Automated slides generation	0	0	0	1	1	0	0	0	0	0	0	2
Automatic Scoring algorithm in sourcing	0	0	1	0	1	0	0	0	0	0	0	2
Automation across investment workflow steps	0	0	1	0	1	0	0	0	0	0	1	3
Automation of CRM updates	0	0	1	0	1	1	0	0	0	0	1	4
Bank-driven deal flow	0	0	1	1	1	1	1	0	0	0	0	5
Bottom-up AI initiatives	1	1	0	0	1	0	1	0	0	0	1	5

Code	Fund 11	Fund 6	Fund 10	Fund 1	Fund 5	Fund 8	Fund 4	Fund 9	Fund 7	Fund 3	Fund 2	Sum
Build-over-buy approach	0	1	1	0	1	0	0	0	1	1	0	5
Buy vs build tension	0	0	1	0	1	0	0	0	1	1	0	4
Buy-over-build approach	0	0	1	0	1	1	1	1	0	0	1	6
Centralized technology decision making	1	0	0	1	1	1	1	1	0	1	0	7
Commercial sense is value driver for productivity in the future	0	0	0	0	1	0	1	0	1	0	0	3
Commoditization of expertise through AI tools	1	0	0	0	1	1	1	0	1	0	0	5
Competitive intensity in deal sourcing	0	0	1	1	1	0	1	0	0	1	1	6
Competitive loss from departing individuals	0	0	0	0	0	1	1	0	0	0	0	2
Competitive loss risk from AI	1	0	0	0	0	0	1	0	0	0	0	2
Competitive pressure from peers for AI adoption	1	0	0	1	1	1	1	1	0	1	1	8
Competitive spillover across PE market segments	0	0	0	0	1	1	1	1	1	1	0	6
Concrete discovery steps to understand AI opportunities	1	1	0	0	1	1	0	0	0	0	0	4
Consultant deliveries increase in complexity due to AI	0	0	0	0	0	0	0	0	1	0	0	1
Consultants leverage proprietary datasets (+)	1	1	0	0	1	0	0	1	0	1	0	5
Context-driven tech spending	1	1	0	0	0	0	0	0	0	0	0	2
Cost concerns for AI	1	0	0	0	0	1	0	0	0	1	0	3
Cost-driven approval processes	0	0	1	1	0	0	1	1	0	0	1	5
Costly specialized data access	0	0	0	0	0	0	0	0	0	1	0	1
Cross-functional tool collaboration	0	0	1	0	1	0	0	0	1	1	0	4
Cross-office deal collaboration tools	0	0	1	1	1	1	0	0	1	1	0	6
Cross-tool collaboration	0	0	0	0	0	1	0	0	0	0	1	2
C-suite decision for major technology adoption	1	1	1	0	1	0	1	1	0	1	0	7
Data format heterogeneity barrier	0	0	0	0	0	0	0	1	0	1	0	2
Data safety concerns (+)	0	0	0	0	0	1	0	0	1	1	0	3
Data validity as bottleneck for AI adoption	1	1	0	0	0	0	0	0	1	0	0	3
Data-driven Large caps, opinion-driven SMCs	1	0	0	0	1	1	0	0	0	1	0	4
DD AI advantage favors large cap	0	0	0	0	0	0	0	0	0	1	0	1
Deal management tool integrated into various deal stages	0	0	1	0	1	1	0	0	1	1	0	5

Code	Fund 11	Fund 6	Fund 10	Fund 1	Fund 5	Fund 8	Fund 4	Fund 9	Fund 7	Fund 3	Fund 2	Sum
Deal pipeline management tools	1	0	1	1	1	1	1	0	1	1	1	9
Dedicated technology personnel with ownership	0	0	1	1	1	1	0	0	1	1	1	7
Dedicated technology team	0	0	1	1	1	1	1	0	1	1	1	8
Desire to understand use cases of AI before implementation	0	1	0	0	0	0	0	1	0	0	0	2
Digitization as a lever in value creation (+)	0	0	0	0	1	0	1	0	1	1	1	5
Diverging financials across data providers	0	0	1	0	0	0	0	0	0	0	0	1
Efficiency gains from AI	1	1	1	1	1	1	1	1	1	1	1	11
Errors in manual reporting processes	0	0	0	0	1	0	1	0	1	0	0	3
Expert networks in due diligence	0	0	0	1	0	0	0	0	0	0	0	1
Externally driven AI Adoption	1	0	0	0	0	0	1	1	0	0	0	3
Firm-level AI tool deployment	1	1	0	1	1	1	0	1	1	1	0	8
Fundraising problems in the PE industry	0	0	0	0	1	1	0	1	0	1	0	4
Future skill requirements of PE professionals	1	0	0	1	1	1	0	1	1	0	0	6
Generic AI access without deep integration	1	0	0	1	0	0	0	1	0	1	0	4
Human process integrity as AI boundary	1	0	0	1	1	0	0	1	0	0	0	4
Information symmetry increases due to AI	0	0	0	1	1	0	1	0	1	0	0	4
Internal data as proprietary moat	1	0	0	0	1	1	0	0	1	0	1	5
Internal database AI integration	0	0	1	0	1	0	0	0	0	0	0	2
Internally driven AI adoption	0	0	1	0	1	1	0	0	0	0	1	4
Intern-as-substitute-for-automation	1	0	0	1	0	0	0	1	0	0	0	3
Investment team takes care of portfolio management	0	0	0	0	0	0	1	0	0	0	0	1
Investment vs. Operations team structure	0	1	1	0	1	0	0	1	0	0	0	4
Irreplaceability of critical judgement in PE	1	1	0	0	1	1	0	1	0	1	0	6
Issues with AI notetakers	0	0	0	0	0	1	0	1	0	0	0	2
Junior analyst workload	1	0	0	1	0	0	1	1	1	1	0	6
Junior demand implementation support	0	0	0	0	0	0	0	0	0	1	0	1
Knowledge tied to individuals	0	0	0	0	0	1	1	1	0	0	0	3
Lack of automated company monitoring	0	0	0	1	0	0	1	0	0	0	0	2
Lack of structured data provider integration	0	0	0	1	0	1	1	0	0	1	0	4

Code	Fund 11	Fund 6	Fund 10	Fund 1	Fund 5	Fund 8	Fund 4	Fund 9	Fund 7	Fund 3	Fund 2	Sum
Lack of technology use in value creation (post signing)	1	1	1	0	1	0	1	1	1	0	0	7
Lack of training for AI tools	0	1	0	1	0	0	0	0	0	0	0	2
Lacking clean data storage processes	1	0	0	0	0	0	1	1	0	1	0	4
Large cap deals are not found through active sourcing	0	0	0	1	0	0	0	1	0	0	0	2
Larger Funds can afford better tech systems	1	0	0	0	0	1	0	0	1	0	0	3
Limited adoption of available tools	1	0	0	1	1	1	0	1	0	0	0	5
Limited integration across investment workflows	0	0	0	1	0	0	1	1	0	1	1	5
Limited internal deal sourcing	0	0	0	1	0	0	1	0	0	0	0	2
Limited sell-side signals	0	0	0	1	0	1	0	0	0	0	0	2
Low organizational focus on process improvement	1	1	0	1	0	1	1	1	0	1	0	7
Low organizational salience of junior pain points	1	0	0	1	0	0	0	1	0	1	0	4
Low priority on inbound automation	0	0	0	0	0	0	0	1	0	0	1	2
Low-barrier experimentation with AI tools	1	1	0	0	0	1	1	1	0	0	1	6
LP demand for transparency and governance	0	0	0	0	0	1	1	0	0	0	0	2
LP non-interference in operational decisions	0	0	0	0	0	0	1	0	0	0	0	1
LP pressure for technological adoption (+)	1	0	0	0	1	1	1	0	0	1	0	5
LP-driven reporting requirements	0	0	1	0	0	1	1	0	0	0	0	3
Manual processes in Due Diligence	1	1	0	0	0	0	1	1	0	1	0	5
Manual reporting workflows	1	1	0	1	1	0	1	1	1	1	1	9
Manual workflows in sourcing and screening	1	1	0	1	1	0	1	1	1	0	0	7
Need for manual cross-checking in ops	0	1	0	0	1	0	0	0	0	0	0	2
No PortCo platform access	0	0	1	0	0	0	0	0	0	1	0	2
No universal AI tool for ops area	0	1	1	0	0	0	0	0	0	0	1	3
Old economy industry context	0	1	0	0	1	0	0	0	0	0	0	2
Only industry specialized Mid-Cap will survive	0	0	0	0	1	0	0	0	1	0	0	2
Open LLM usage	1	0	0	1	1	1	0	1	1	1	1	8
Openness to AI adoption	1	1	0	1	1	1	1	1	0	0	1	8
Operational scope of the tech team	0	0	1	0	0	0	1	0	0	0	1	3

Code	Fund 11	Fund 6	Fund 10	Fund 1	Fund 5	Fund 8	Fund 4	Fund 9	Fund 7	Fund 3	Fund 2	Sum
Organizational inertia in technology adoption	0	1	0	1	0	1	0	1	0	1	0	5
Outbound sourcing dominance	0	0	1	0	1	1	1	0	0	0	1	5
Overreliance on AI could harm investment decisions	0	0	1	1	0	1	0	0	0	0	0	3
Paid expert calls	0	0	0	1	1	0	0	1	0	0	0	3
PE mindset is structurally conservative	0	1	0	0	0	0	0	0	0	1	0	2
Peer pressure rejection for AI adoption	0	1	0	0	0	0	0	1	1	0	0	3
Personal reliability aversion (+)	0	1	0	0	0	0	0	0	1	1	0	3
Planned AI adoption	0	0	1	1	1	1	0	1	0	0	0	5
Portfolio team tech ownership	0	0	0	0	0	0	0	0	0	1	0	1
Potential for AI notetaker	0	0	0	0	0	1	0	1	1	1	0	4
Preset technology stack (+) (+)	0	0	1	1	0	1	0	0	0	0	1	4
Proactive company outreach rationale	1	0	1	1	1	1	1	1	0	1	1	9
Productivity is priority (+)	1	1	1	0	1	1	0	0	1	1	0	7
Public registry integration	0	0	1	0	0	0	0	0	0	0	1	2
Replicability of tech stack across Funds	0	0	0	0	1	1	1	0	0	0	1	4
Scale advantage in larger Funds	1	0	0	1	1	1	0	0	1	1	1	7
Search for informational advantage in smaller markets	0	0	0	0	0	1	1	0	1	0	1	4
Sector expertise as competitive edge	0	0	0	0	0	1	1	0	0	0	1	3
Senior gatekeeping	0	1	0	0	0	1	0	1	0	1	0	4
Separate PortCo tracking	0	0	0	0	0	0	0	0	1	1	0	2
Shared target tracking criteria	0	0	0	1	0	0	0	0	0	0	0	1
Signal-based top funnel sourcing tools	0	0	1	0	1	0	0	0	0	0	1	3
Skepticism about predictable competitive advantage	0	1	1	0	0	0	0	1	0	1	0	4
Skepticism: Output reliability	0	0	0	0	0	0	0	0	1	1	0	2
Skepticism: ROI Uncertainty	1	1	0	0	0	0	0	1	0	1	0	4
Skepticism: Technical feasibility limits	1	1	0	0	0	1	0	0	0	1	0	4
SMCF deals less data-driven	1	0	0	0	1	1	0	0	0	1	0	4
Sourcing AI advantage favors MidCap	0	0	0	0	0	1	0	0	0	1	0	2
Specialized DD data providers used	0	0	0	0	0	0	0	0	0	1	0	1
Standardized sourcing workflows	0	0	1	1	1	1	1	0	0	0	0	5
Strategic vs. administrative AI value distinction	1	0	0	0	1	1	1	1	0	0	1	6

Code	Fund 11	Fund 6	Fund 10	Fund 1	Fund 5	Fund 8	Fund 4	Fund 9	Fund 7	Fund 3	Fund 2	Sum
Talent scarcity for AI implementation	0	1	0	0	0	1	0	0	0	1	0	3
Team size constraints	1	1	0	1	0	1	0	1	0	0	1	6
Tech as selling point to LPs	0	0	1	0	0	1	0	0	0	0	0	2
Tech lag not necessarily disadvantageous for smaller Funds	0	0	0	0	0	1	0	0	0	1	0	2
Tech level in PEs is dependent on sector context	0	1	1	0	1	0	0	0	0	0	0	3
Tech native operator mindset	1	1	1	0	1	0	0	0	0	0	1	5
Technology adoption driven by operational pain points	1	1	0	0	0	0	1	0	0	1	1	5
Technology adoption in PortCos depends on context	0	1	1	0	0	0	1	0	0	1	1	5
Technology self assessment	1	0	1	1	1	1	1	1	1	1	1	10
Technology usage in Due Diligence	1	0	0	1	0	1	1	1	1	0	1	7
Technology use in sourcing and screening	1	0	1	1	1	1	0	0	1	0	1	7
Third-party CRM	1	0	0	1	1	1	1	0	1	1	1	8
Touchpoint and relationship tracking	0	0	1	0	1	1	0	0	1	1	0	5
Trust in human work over AI tools	1	0	0	1	0	1	0	0	0	1	0	4
Unclear AI policy	0	1	0	0	0	0	0	1	0	1	0	3
Value of relationships	1	0	0	1	0	1	1	1	1	0	1	7
Wait-and-see AI adoption strategy	1	1	0	1	0	0	0	1	1	1	0	6
Weak institutional knowledge retention	1	0	0	0	0	1	1	1	0	0	0	4
Workforce reduction pressure from AI	1	0	0	0	1	0	0	1	1	1	0	5
Total	67	46	50	52	75	77	62	62	54	73	51	669

Table 7: Sourcing Dimension Score Levels. Score anchors for the Sourcing dimension of the AI Matrix (Table 1). Scores calibrated from interview material; example codes consistent with each level listed in the rightmost column.

Score	Description	Codes consistent with this level
3	Proprietary AI-driven sourcing platform with internal-data scoring or active deal-flow generation; AI-supported workflow tightly integrated across multiple sourcing stages.	‘Signal-based top funnel sourcing tools’ (n=3); ‘Automatic Scoring algorithm in sourcing’ (n=2); ‘Public registry integration’ (n=2); supported by ‘Outbound sourcing dominance’, ‘Standardized sourcing workflows’, ‘Touchpoint and relationship tracking’.
2	Substantive third-party or in-house tooling integrated across at least one sourcing sub-process (signal scanning, screening, pipeline tracking) with operational use; AI-supported in some form but not proprietary platform.	‘Technology use in sourcing and screening’ (n=7) plus at least one of ‘Standardized sourcing workflows’ (n=5), ‘Touchpoint and relationship tracking’ (n=5); ‘Deal pipeline management tools’ (n=9) used in operationally substantive way.
1	Generic third-party tools (e.g., LinkedIn, PitchBook, basic CRM) without AI-driven scoring or signal generation; sourcing remains predominantly manual or referral-based.	‘Deal pipeline management tools’ (n=9) carried but used as generic CRM/list rather than as integrated platform; ‘Proactive company outreach rationale’ (n=9) without supporting tooling.
0	No AI integration in deal sourcing; manual or referral-based pipeline; no observed sourcing-side technology use.	Absence of any of the codes above, or only weak co-occurrence with explicit articulation that no sourcing technology is used (e.g., ‘Large cap deals are not found through active sourcing’ as standalone).

Table 8: Due Diligence Dimension Score Levels. Score anchors for the Due Diligence dimension of the AI Matrix (Table 1). Scores calibrated from interview material; example codes consistent with each level listed in the rightmost column.

Score	Description	Codes consistent with this level
3	AI-integrated DD across multiple workflow stages with proprietary or deeply customized tools; AI substantively integrated into document screening, summary generation, and IC document production.	‘Technology usage in Due Diligence’ (n=7) co-occurring with ‘Automated slides generation’ (n=2) and proprietary or near-proprietary integration; ‘Automated sentiment analysis of reference calls’ (n=1) at the leading edge.
2	AI-supported DD on specific sub-processes (note-taking, document summarization, reference-call transcripts, screening of specific document types) using established third-party or general-purpose AI tools.	‘Technology usage in Due Diligence’ (n=7) plus ‘AI call notetaker’ (n=5) used substantively; ‘Paid expert calls’ (n=3) with AI-supported transcript or summary generation; ChatGPT Enterprise or Copilot-style tools applied to DD documents.
1	Generic AI tools (e.g., ChatGPT) used for ad-hoc DD tasks without DD-workflow integration; no systematic AI use across DD stages.	‘Technology usage in Due Diligence’ (n=7) carried but described as ad-hoc / individual analyst use; ‘AI call notetaker’ (n=5) used inconsistently or in pilot stage; no integrated DD platform.
0	No AI use in DD workflows; manual analysis only; expert calls and document review entirely without AI support.	Absence of any DD-related AI codes; explicit articulation of manual DD process.

Table 9: Knowledge & Data Infrastructure Dimension Score Levels. Score anchors for the Knowledge & Data Infrastructure dimension of the AI Matrix (Table 1). Scores calibrated from interview material; example codes consistent with each level listed in the rightmost column.

Score	Description	Codes consistent with this level
3	Operational AI integration with proprietary Fund-internal data, in any form. Apex tier captures the operationalizing configurations identified in Section 4.5.4. Internal databases, IC paper repositories, or document libraries with AI-driven retrieval, integration, or generation.	‘Internal database AI integration’ (n=2: Funds 5, 10); ‘AI-connected knowledge retrieval’ (n=2: Funds 7, 8). The two codes do not overlap at any Fund. Both signal score 3.
2	Foundational data infrastructure with AI-supported tagging, search, or document organization; substantive operational AI integration on a bounded use case (e.g., document retrieval); articulated frontier without full operationalization.	‘AI potential for knowledge institutionalization’ (n=7) co-occurring with substantive infrastructure investment; ‘Cross-office deal collaboration tools’ (n=6) and ‘Deal management tool integrated into various deal stages’ (n=5) used as integrated platforms; ‘Third-party CRM’ (n=8) with substantive integration.
1	Generic data infrastructure (CRM, document management) without AI integration; data accessible but not AI-usable.	‘Third-party CRM’ (n=8) carried as generic system; ‘Limited integration across investment workflows’ (n=5) signals fragmented infrastructure; absence of either operationalization code.
0	No proprietary data infrastructure articulated; data fragmented across individual files, inboxes, or senior memory; no AI integration with internal data.	‘Lacking clean data storage processes’ (n=4) and ‘Weak institutional knowledge retention’ (n=4) co-occurring with absence of proprietary infrastructure; ‘Knowledge tied to individuals’ (n=3) as primary mode of knowledge capture.

Table 10: Reporting & Portfolio Monitoring Dimension Score Levels. Score anchors for the Reporting & Portfolio Monitoring dimension of the AI Matrix (Table 1). Scores calibrated from interview material; example codes consistent with each level listed in the rightmost column.

Score	Description	Codes consistent with this level
3	AI-driven portfolio monitoring or value-creation workflows; anomaly detection, automated KPI tracking, AI-driven reporting integrated across portfolio companies.	‘Automated portfolio financial reporting’ (n=3); ‘Automation across investment workflow steps’ (n=3); ‘Automation of CRM updates’ (n=4) used substantively in reporting; absence of ‘Manual reporting workflows’ or its strict bounding.
2	AI-supported reporting workflows (automated extraction, standardized templates with automated processing, draft generation); systematic but not fully AI-driven.	‘Automated portfolio financial reporting’ (n=3) used in standardized but partially-manual flow; ‘Automation of CRM updates’ (n=4); structured PortCo data flows.
1	Standard reporting workflows; AI used for ad-hoc tasks only; manual processes dominate but with some technology support.	‘Manual reporting workflows’ (n=9) carried alongside ad-hoc AI assistance; structured tooling without AI integration; ‘LP-driven reporting requirements’ (n=3) without AI fulfillment.
0	No AI in portfolio monitoring or reporting; entirely manual; no observable technology integration in the post-signing phase.	‘Manual reporting workflows’ (n=9) as standalone; ‘Errors in manual reporting processes’ (n=3); ‘Lack of technology use in value creation (post signing)’ (n=7).

Table 11: Codes by Themes and Subcategories. Theme assignment of the 162 leaf-level codes across five inductively constructed themes. Sum column reports each code's carrier count across the eleven Funds. 21 codes touching multiple analytical domains appear under two themes each and are shaded light blue, producing 183 theme-level rows from the 162 unique codes.

Theme 1: AI Maturity Variation & Context	SUM
1.1 Technology self assessment	
AI as competitive differentiation	7
Tech native operator mindset	5
Technology self assessment	10
1.2 Sourcing & Screening – Status Quo	
Automatic Scoring algorithm in sourcing	2
Deal pipeline management tools	9
Large cap deals are not found through active sourcing	2
Outbound sourcing dominance	5
Proactive company outreach rationale	9
Public registry integration	2
Shared target tracking criteria	1
Signal-based top funnel sourcing tools	3
Standardized sourcing workflows	5
Technology use in sourcing and screening	7
Touchpoint and relationship tracking	5
1.3 Due Diligence – Status Quo	
Automated sentiment analysis of reference calls	1
Automated slides generation	2
Costly specialized data access	1
Expert networks in due diligence	1
Paid expert calls	3
Specialized DD data providers used	1
Technology usage in Due Diligence	7
1.4 Portfolio / PMI – Status Quo	
Anti-ideological tech implementation in PortCos	4
Automated portfolio financial reporting	3
Digitization as a lever in value creation (+)	5
Investment team takes care of portfolio management	1
No PortCo platform access	2
Portfolio team tech ownership	1
Separate PortCo tracking	2
Technology adoption in PortCos depends on context	5
1.5 Reporting & LP Relations – Status Quo	
Errors in manual reporting processes	3
LP demand for transparency and governance	2
LP non-interference in operational decisions	1

LP-driven reporting requirements	3
Manual reporting workflows	9
1.6 Holistic Tool Infrastructure	
Cross-tool collaboration	2
AI call notetaker	5
Automation across investment workflow steps	3
Automation of CRM updates	4
Cross-functional tool collaboration	4
Cross-office deal collaboration tools	6
Deal management tool integrated into various deal stages	5
Firm-level AI tool deployment	8
Generic AI access without deep integration	4
Limited adoption of available tools	5
Open LLM usage	8
Preset technology stack (+) (+)	4
Third-party CRM	8
1.7 Fund Size, Sector & Context	
AI efficiency scales with data volume	6
AI implementation in automotive sector	1
Bank-driven deal flow	5
Data-driven Large caps, opinion-driven SMCFs	4
Larger Funds can afford better tech systems	3
Old economy industry context	2
SMCF deals less data-driven	4
Scale advantage in larger Funds	7
Tech level in PEs is dependent on sector context	3
Total	230

Theme 2: Organizational Configuration Features	SUM
2.1 Personnel & Ownership	
Absence of dedicated technology personnel	4
Dedicated technology personnel with ownership	7
Dedicated technology team	8
Junior analyst workload	6
Operational scope of the tech team	3
Talent scarcity for AI implementation	3
Team size constraints	6
2.2 Governance & Decision-Making	
C-suite decision for major technology adoption	7
Centralized technology decision making	7
Context-driven tech spending	2
Cost-driven approval processes	5

Investment vs. Operations team structure	4
Senior gatekeeping	4
Unclear AI policy	3
2.3 Build vs. Buy Strategy	
Build-over-buy approach	5
Buy vs build tension	4
Buy-over-build approach	6
2.4 Culture & Initiative	
Bottom-up AI initiatives	5
Lack of training for AI tools	2
Low organizational focus on process improvement	7
Low organizational salience of junior pain points	4
Low priority on inbound automation	2
Low-barrier experimentation with AI tools	6
Organizational inertia in technology adoption	5
Productivity is priority (+)	7
Technology adoption driven by operational pain points	5
2.5 Principal-Agent Gap	
Intern-as-substitute-for-automation	3
Junior demand implementation support	1
Need for manual cross-checking in ops	2
2.6 Sector Expertise & Relationships	
Sector expertise as competitive edge	3
Value of relationships	7
Total	143

Theme 3: Wait-and-See Deferral Patterns	SUM
3.1 The Wait-and-See Pattern	
Desire to understand use cases of AI before implementation	2
Openness to AI adoption	8
Planned AI adoption	5
Tech lag not necessarily disadvantageous for smaller Funds	2
Wait-and-see AI adoption strategy	6
3.2 Active Monitoring – Rational Delay Evidence	
Concrete discovery steps to understand AI opportunities	4
Internally driven AI adoption	4
Low-barrier experimentation with AI tools	6
3.3 Passive Inertia – Strategic Blindspot Evidence	
AI hype skepticism	1
PE mindset is structurally conservative	2
Low organizational focus on process improvement	7

Organizational inertia in technology adoption	5
3.4 Skepticism Cluster	
AI edge requires empirical testing to validate	1
Automated Slide generation potential	3
Human process integrity as AI boundary	4
Irreplaceability of critical judgement in PE	6
Overreliance on AI could harm investment decisions	3
Personal reliability aversion (+)	3
Skepticism about predictable competitive advantage	4
Skepticism: Output reliability	2
Skepticism: ROI Uncertainty	4
Skepticism: Technical feasibility limits	4
Trust in human work over AI tools	4
3.5 Structural Blockers	
Diverging financials across data providers	1
Cost concerns for AI	3
Data format heterogeneity barrier	2
Data safety concerns (+)	3
Data validity as bottleneck for AI adoption	3
Externally driven AI Adoption	3
Lacking clean data storage processes	4
Limited integration across investment workflows	5
No universal AI tool for ops area	3
3.6 Drivers of Adoption	
Potential for AI notetaker	4
Competitive pressure from peers for AI adoption	8
Efficiency gains from AI	11
LP pressure for technological adoption (+)	5
Peer pressure rejection for AI adoption	3
Tech as selling point to LPs	2
Productivity is priority (+)	7
Total	157

Theme 4: Anticipated Competitive Dynamics	SUM
4.1 Competitive Pressure Mechanisms	
AI intensifying industry competition	5
Competitive intensity in deal sourcing	6
Competitive loss from departing individuals	2
Competitive loss risk from AI	2
Competitive spillover across PE market segments	6
Information symmetry increase due to AI	4
Search for informational advantage in smaller markets	4

4.2 The Three Cascade Triggers	
Fundraising problems in the PE industry	4
Future skill requirements of PE professionals	6
Workforce reduction pressure from AI	5
LP pressure for technological adoption (+)	5
Externally driven AI Adoption	3
4.3 AI as Moat – Proprietary vs. Generic	
AI potential in deal sourcing	8
AI value concentrated in deal sourcing	6
Internal data as proprietary moat	5
Internal database AI integration	2
Replicability of tech stack across Funds	4
Strategic vs. administrative AI value distinction	6
AI-connected knowledge retrieval	2
4.4 Market Structure & Consolidation	
Consultant deliveries increase in complexity due to AI	1
AI decision guidance rather than full automation	4
AI-driven shift in consultant dependency	4
Commoditization of expertise through AI tools	5
Consultants leverage proprietary datasets (+)	5
DD AI advantage favors large cap	1
Only industry specialized Mid-Cap will survive	2
Sourcing AI advantage favors MidCap	2
4.5 Future of PE Professionals	
AI automation trajectory	6
Analysts will be fully replaced by AI in the future	3
Commercial sense is value driver for productivity in the future	3
Value of relationships	7
Future skill requirements of PE professionals	6
Total	134

Theme 5: Knowledge Institutionalization Gap	SUM
5.1 The Fragmentation Problem	
Knowledge tied to individuals	3
Lack of structured data provider integration	4
Limited internal deal sourcing	2
Limited sell-side signals	2
Weak institutional knowledge retention	4
Lacking clean data storage processes	4
Competitive loss from departing individuals	2
5.2 The Proof of Concept	
AI-connected knowledge retrieval	2

Lack of automated company monitoring	2
AI potential for knowledge institutionalization	7
AI potential in specific workflow stages	6
Internal database AI integration	2
5.3 Knowledge Infrastructure	
Third-party CRM	8
Touchpoint and relationship tracking	5
Automation of CRM updates	4
Cross-office deal collaboration tools	6
Deal management tool integrated into various deal stages	5
5.4 Consultant Dependency as Symptom	
Consultants leverage proprietary datasets (+)	5
AI-driven shift in consultant dependency	4
5.5 Pain Points Feeding the Gap	
Issues with AI notetakers	2
Lack of technology use in value creation (post signing)	7
Manual processes in Due Diligence	5
Manual workflows in sourcing and screening	7
Manual reporting workflows	9
Limited integration across investment workflows	5
Total	112


Table 12: Cited NCORs in Sections 4 and 5. Normalized Co-Occurrence Ratios cited in the Findings and Discussion chapters. $NCOR(A,B) = |CA \cap CB| / \min(|CA|, |CB|)$, where CA and CB denote the sets of Funds carrying codes A and B (Section 3.3, Eq. (1)).

Section	Code A	Code B	Co-occurrences	NCOR	Source cross-tab
4.3.3	Absence of dedicated technology personnel (n = 4)	Skepticism: ROI Uncertainty (n = 4)	3	75% NCOR (3/4)	T2 × T3
4.3.3	Absence of dedicated technology personnel (n = 4)	Lacking clean data storage processes (n = 4)	3	75% NCOR (3/4)	T2 × T3
4.4.2	AI intensifying industry competition (n = 5)	Commoditization of expertise through AI tools (n = 5)	5	100% NCOR (5/5)	T4 × T4
4.4.3	Competitive spillover across PE market segments (n = 6)	AI as competitive differentiation (n = 7)	5	83% NCOR (5/6)	T1 × T4
4.5.1	Internally driven AI adoption (n = 4)	AI potential for knowledge institutionalization (n = 7)	4	100% NCOR (4/4)	T3 × T5
4.5.1 & 5.2	Externally driven AI adoption (n = 3)	Strategic vs. administrative AI value distinction (n = 6)	3	100% NCOR (3/3)	T3 × T4
4.5.1	Internally driven AI adoption (n = 4)	Externally driven AI adoption (n = 3)	0	0% NCOR (0/3)	T3 × T3
4.5.2	Buy vs build tension (n = 4)	Build-over-buy approach (n = 5)	4	100% NCOR (4/4)	T2 × T2

Table 13: Additional Standard Quality Criteria Provisions. Methodological quality provisions following Lincoln and Guba’s (1985) trustworthiness framework, beyond those discussed in Section 3.4.

Provision	Description	Methodological Grounding
Triangulation	The study relied on triangulation not through combining multiple methods, but through systematic comparison across Fund type, geography, and seniority. This enabled the analysis to assess whether themes recurred across heterogeneous organizational settings rather than appearing only in isolated cases, strengthening credibility through structured cross-Fund-type, cross-geography, and cross-seniority comparison.	Lincoln and Guba (1985) credibility provision; triangulation as structured comparative discipline rather than method combination.
Negative Case Analysis	The study incorporated negative case analysis by explicitly documenting instances that did not fit the dominant pattern. Deviant cases were treated not as methodological noise but as analytically valuable evidence capable of refining, challenging, or qualifying emerging interpretations. This was particularly important in a thesis concerned with variation in adoption maturity, where outlier cases carry substantive analytical weight rather than reducing to measurement error.	Negative case analysis as discipline against confirmation bias in qualitative research; treatment of deviance as analytically informative.
Researcher Reflexivity	The study acknowledged the importance of researcher reflexivity. Given the contemporary salience of AI and the strong narratives surrounding technological disruption in finance, the researchers remained attentive to the risk of over-ascribing coherence, intentionality, or strategic sophistication to respondents’ accounts, using reflexivity as a discipline of interpretation to keep findings grounded in the empirical material rather than in prior assumptions.	Kvale and Brinkmann (2009): qualitative interviewing approaches recognizing the interviewer’s role in shaping both the production and interpretation of data.

Figure 2: Code Saturation Curve. New-code emergence per interview in chronological order of data collection (28 January – 7 April), with cumulative code count plotted alongside. New-code emergence declined steeply after the fourth interview and flattened by the seventh, confirming practical saturation within the eleven-interview sample.



Note: Bars indicate new codes generated per interview; line indicates cumulative code count. Percentages shown are cumulative codes as % of total (162).

Table 14: Configuration Feature Distribution Across Sample Funds. Binary indicators of three configuration features and the resulting joint-pattern grouping. AI Matrix scores for cross-reference.

Fund	Mean Score	Tier	Operationally-Mandated AI Personnel	Bottom-Up AI Initiatives	Active-Governance Signature
Fund 5	2.50	High	✓	✓	✓
Fund 2	2.50	High	✓	✓	–
Fund 10	2.25	High	✓	–	–
Fund 8	1.75	Middle	✓	✓	✓
Fund 7	1.50	Middle	✓	–	–
Fund 4	1.25	Middle	–	✓	–
Fund 3	1.25	Middle	(scope)	–	–
Fund 1	0.75	Low	(scope)	–	–
Fund 11	0.50	Low	–	✓	✓
Fund 9	0.25	Low	–	✓	–
Fund 6	0.25	Low	–	✓	✓

Note: ✓ = configuration feature is coded for the fund; – = feature is not coded; (scope) = personnel role coded but mandate excludes the investment workflow. Tier thresholds: High ≥ 2.0 ; Middle 1.0–1.99; Low < 1.0 . Configuration features defined in Sections 4.2.1–4.2.3.

AI Appendix

What AI tools have been used and how?

Anthropic Claude Opus 4.7 (adaptive), Perplexity Computer, Read.ai, and MAXQDA's AI Assist were used in this thesis. Claude served as a quarriable backend for our binarized code-coverage and theme-correlation matrices, and as a writing aid for clarity and redundancy checks. In a separate chat, Claude was also used to stress-test our sample anonymization: because curious readers are likely to query AI tools and the published thesis will be scraped by them, we prompted candidate per-fund disclosures and tightened our anonymization discipline based on what the model could plausibly infer. Perplexity Computer produced indicative dimension clusterings on the binarized code matrix. Read.ai recorded interviews and generated initial transcripts. MAXQDA's AI Assist suggested codes during open coding. No AI tool was used to generate data, produce empirical findings, or draft sections without human editing.

In what ways have these tools contributed to increasing the quality of the thesis?

Querying our own matrices through Claude let us examine code coverages and correlations more systematically than manual inspection allowed, which strengthened the empirical grounding of our analysis. Perplexity Computer's clusterings offered useful starting hypotheses for deriving higher-order dimensions. Read.ai reduced time investment for transcript generation. For the text itself, Claude helped tighten sentences and improve flow, making the thesis more accessible to readers outside our specific topic.

What potential risks were found using AI and what measures were taken to reduce these risks?

The two risks we encountered were hallucination and imprecision, and they often reinforced each other. The models confirmed what we proposed, occasionally fabricated supporting details, and were not accurate enough to be trusted without human challenging.

We addressed this by anchoring every AI output to verifiable material on our side. Coding and clustering outputs were validated against the original MAXQDA data and interview transcripts. Textual arguments were cross-checked against our references. Every code from MAXQDA's AI Assist passed through analyst review before being retained. We also actively prompted the models to argue against their own suggestions where this was useful, to offset their bias toward agreement.

What are the insights gained from using AI tools in the writing process?

AI proved most useful where we supplied the structure, and least useful when asked to interpret source material independently. When engaging with academic literature it routinely invented references or shifted the meaning of sources through paraphrasing, which forced enough manual verification to neutralize much of the time saved. Our overall takeaway is that AI raised the ceiling of data handling and how clearly we could write, but the substantive judgments, the coding decisions, and the argumentation remained ours to make.