

American University Kyiv

A Capstone Project

**AI in the Boardroom: Director Sensemaking in Ukrainian Corporate
Governance**

**Штучний інтелект у залі засідань: осмислення директорами феномену
ШІ в українському корпоративному управлінні.**

by Artem Skorupych

Presented in Partial Fulfillment of the Requirements
for the Degree

Master

APPROVED BY:

Artem Klyuchnikov, PhD

Adjunct Professor at American University Kyiv

2026

TABLE OF CONTENTS

ABSTRACT	4
CHAPTER 1: INTRODUCTION	5
CHAPTER 2: LITERATURE REVIEW AND THEORETICAL FRAMEWORK	6
2.1 Theoretical Framework: Understanding AI Through Multiple Lenses.....	7
2.2 Critical Challenges: Explainability and Performativity	13
2.3 The Gap: What We Don't Know About AI in Boardrooms.....	15
CHAPTER 3: RESEARCH METHODOLOGY	18
3.1 Research Design: Why Qualitative?	18
3.2 Data Collection Methods	19
3.3 Sampling Strategy	20
3.4 Data Analysis	21
3.5 Research Quality	23
3.6 Ethical Considerations	24
3.7 Limitations	25
CHAPTER 4: FINDINGS	26
4.1 Analytical Approach.....	26
4.2 Participant Characteristics	28
4.3 How Directors Make Sense of AI (RQ1).....	29
4.3.1 AI as Practical Solution to Board Challenges	29
4.3.2 AI as Source of Governance Problems	31
4.3.3 AI as Data and Security Vulnerability.....	33
4.3.4 The Dialectic of Promise and Risk	35
4.4 How Composition and Context Shape Perspectives (RQ2).....	36
4.4.1 Board Composition: Technical Expertise as Double-Edged Sword	36
4.4.2 Organizational Context: War, Regulation, and Sector	39
4.5 Governance Implications and Practices (RQ3).....	42
4.5.1 How AI Might Change Boards	42
4.5.2 What Practices Boards Need.....	43
CHAPTER 5: DISCUSSION	47
5.1 Introduction.....	47
5.1.1 Distinguishing Theory-Led and Data-Led Insights	48
5.2 Understanding AI's Dual Nature (RQ1).....	50

5.2.1 Bounded Rationality and Cognitive Extension.....	50
5.2.2 Agency Concerns and the Monitoring Function	51
5.2.3 Sensemaking Under Technological Disruption.....	52
5.3 Shaping Forces: Composition and Context (RQ2)	54
5.3.1 Upper Echelons Effects and Epistemic Power.....	54
5.3.2 Socio-Technical Embeddedness and Context Dependency	55
5.3.3 Institutional Pressures and Legitimacy Seeking	56
5.4 Governance Innovation and Adaptation (RQ3)	58
5.4.1 Enduring Principles, Evolving Practices.....	58
5.4.2 Human-in-the-Loop as Foundational but Insufficient	59
5.4.3 Governance as Collective Learning Process.....	60
5.5 Integrating Insights Across Research Questions.....	61
5.5.1 How Understanding Shapes and Is Shaped by Context.....	61
5.5.2 How Context and Composition Interact	62
5.5.3 How Practices Reflect and Reinforce Understanding.....	63
5.5.4 Theoretical Integration: Toward Contingent AI Governance.....	64
CHAPTER 6: CONCLUSION.....	65
6.1 Summary of Key Insights	65
6.2 Theoretical Contributions	66
6.3 Practical Implications.....	68
6.4 Limitations and Boundary Conditions	70
6.5 Future Research Directions	71
6.6 Concluding Reflections.....	73
REFERENCES.....	76
APPENDIX A: Data Collection Instruments	81
APPENDIX B: Data Structure Tables	85

ABSTRACT

While artificial intelligence increasingly influences organizational decision-making, how corporate board directors make sense of AI for governance purposes remains underexplored, particularly in non-Western and high-stress contexts. This study investigates three questions: how boards currently engage with AI, how board composition and organizational context shape directors' perspectives, and what governance practices directors identify as necessary for responsible adoption. Using a qualitative design, the research collected data through semi-structured interviews (n=5) and written qualitative surveys (n=17) with Ukrainian board directors across banking, technology, energy, healthcare, and other sectors. Analysis followed Gioia-inspired methodology, progressing from first-order concepts through interpretive themes to aggregate theoretical dimensions. Findings reveal directors hold dialectical understanding of AI—simultaneously recognizing its potential to address cognitive constraints (information overload, backward focus, data fragmentation) while creating governance risks (explanation difficulties, accountability ambiguity, judgment erosion). Board composition, particularly the mix of technical and traditional expertise, systematically shapes these perspectives, while Ukrainian wartime conditions create paradoxical pressures making AI both more urgent and more risky. Directors converge on governance practices emphasizing human-in-the-loop principles, formal frameworks, transparency, and director capability-building. The study contributes to bounded rationality, upper echelons, and socio-technical systems theories while demonstrating how extreme contexts function as theoretical microscopes revealing dynamics relevant to boards globally.

Keywords: artificial intelligence, corporate governance, board directors, sensemaking, bounded rationality, Ukraine

CHAPTER 1. Introduction

Imagine a boardroom discussion where directors debate a critical strategic decision—perhaps a major acquisition or market entry. Traditionally, this deliberation draws on directors' experience, management's analysis, and perhaps consultant reports. Now imagine the same discussion enhanced by artificial intelligence: algorithms processing thousands of data points, modeling multiple scenarios, identifying patterns invisible to human analysis. This is no longer hypothetical. AI is entering the boardroom, and it raises fundamental questions about the nature of governance itself.

While AI's operational benefits are well-documented (Brynjolfsson & McAfee, 2017; Davenport & Ronanki, 2018), its implications for board governance remain surprisingly underexplored. Research has documented AI's transformative effects in operational contexts (Shrestha et al., 2019), yet questions remain about how it transforms the dynamics of power, accountability, and judgment at the very top of the organization. When algorithms inform strategic choices, who is truly accountable? When directors cannot fully understand AI's logic, can they fulfill their fiduciary duties? These are not merely technical questions—they strike at the heart of corporate governance.

This study addresses these challenges through three research questions:

RQ1: How do corporate boards currently engage with and make sense of AI for strategic decision-making? This question explores whether boards are discussing, considering, piloting, or using AI, what prompts these discussions, and how directors reason about AI's potential benefits and risks. Understanding both current engagement patterns and director sensemaking establishes the empirical baseline and reveals how directors weigh competing values like accuracy versus interpretability.

RQ2: How do board composition and organizational context shape directors' perspectives on AI in governance? This question investigates whether board composition—digital literacy, professional backgrounds, cognitive diversity—influences how directors perceive AI, and how factors like industry context, organizational size, and governance culture affect their reasoning about AI adoption.

RQ3: What governance implications and practices do directors identify as important for responsible AI adoption? This question explores anticipated or experienced changes in board dynamics, accountability structures, and deliberation quality, as well as what governance practices directors believe necessary for responsible adoption. It examines both the transformations AI may introduce and the safeguards boards should implement.

Addressing these questions through systematic qualitative research with Ukrainian board directors fills the empirical gap, extends and refines theoretical understanding, and provides practical insights for boards navigating AI adoption in challenging contexts.

CHAPTER 2: LITERATURE REVIEW AND THEORETICAL FRAMEWORK

2.1 Theoretical Framework: Understanding AI Through Multiple Lenses

Understanding AI in governance requires theoretical frameworks that explain both what AI promises and what challenges it introduces. Like viewing a complex landscape through different camera lenses—wide-angle, telephoto, macro—each theoretical perspective reveals patterns invisible from other vantage points. This research employs three primary theories as its core analytical framework, each illuminating a fundamental dimension of AI in boardrooms: how it affects human cognition, how differences in board composition are associated with its impact, and how it reshapes governance as a system. Four supporting theories extend and operationalize these foundations, providing specific conceptual tools for examining accountability mechanisms, temporal dynamics, and practical engagement strategies.

Together, these perspectives provide the conceptual foundation for investigating the three research questions stated above. They illuminate what AI means for board judgment, whose voices gain or lose influence when algorithms enter deliberations, and what new forms of governance practice emerge when technology and human wisdom must coexist.

Three Primary Theories: The Core Framework

Bounded Rationality: AI as Cognitive Prosthetic. Simon (1947, 1997) argued that bounded rationality theory fundamentally challenged economics' assumption of the perfectly rational decision-maker. Simon observed that humans don't optimize; we "satisfice." We make decisions that are good enough given our cognitive limitations, incomplete information, and time constraints. A board director facing a complex strategic choice cannot simultaneously

consider all options, forecast all consequences, or process all relevant data. Bounded rationality is the human condition.

AI promises to push back the boundaries of rationality, to extend what Simon called our "computational limits." Algorithms can process datasets no human could read, identify correlations no analyst could spot, model scenarios too complex for spreadsheets. In this view, AI is a cognitive prosthetic—a tool that makes boards smarter by overcoming the very constraints Simon identified.

But AI may also extend some cognitive boundaries while creating new ones. When an algorithm recommends a strategic course of action, directors face questions about whether they can truly understand why, validate the recommendation, identify its blind spots, and take genuine intellectual ownership of the choice. The “black box” problem means that AI doesn’t simply expand rationality—it may transform it, replacing one set of cognitive constraints (limited information processing) with another (limited algorithmic interpretability). This study uses the working label “bounded algorithmic rationality” to describe the possibility that AI simultaneously extends and constrains directors’ cognitive capacity.

Recent scholarship continues to extend Simon's framework to contemporary contexts. Kahneman (2011) built on bounded rationality to demonstrate systematic cognitive biases affecting decision-making, earning a Nobel Prize for showing how predictable deviations from rationality shape judgment. More recently, researchers have applied bounded rationality specifically to algorithmic decision contexts. Čaić et al. (2018) demonstrate how AI systems can both alleviate and exacerbate cognitive constraints, while Jarrahi (2018) examines how human and artificial intelligence can complement each other within bounded rationality frameworks. These recent developments show that Simon's foundational insight—that human rationality is constrained and situated—remains highly relevant for understanding how decision-makers engage with AI systems that promise to extend cognitive capacity while introducing new limitations.

Upper Echelons Theory: Why Board Composition Matters. Hambrick and Mason's (1984) upper echelons theory suggests that organizations reflect their leaders' cognitive bases and values. They articulated a foundational insight that has generated hundreds of empirical studies:

to understand why organizations make certain choices, look to the characteristics—backgrounds, expertise, experiences—of the people at the top. Executive characteristics serve as proxies for interpretive frameworks shaping strategic choices.

Applied to AI adoption, upper echelons theory generates predictions. A board populated with technology executives may perceive AI differently than a board of traditional manufacturing leaders. Directors with quantitative training may be more comfortable with algorithmic recommendations than those whose expertise is qualitative and experiential. Cognitive diversity—the variety of perspectives and problem-solving approaches on a board—may affect how critically boards interrogate AI outputs rather than accepting them at face value.

Upper echelons theory has continued to evolve and generate substantial empirical research. Hambrick (2007) revisited the theory, noting that over 200 studies had built on the original framework, confirming its explanatory power while identifying boundary conditions. Recent applications extend the theory to digital contexts specifically. Bughin et al. (2017) demonstrate that CEO digital literacy significantly predicts organizational AI adoption, while Haffke et al. (2016) show how CIO background characteristics shape digital transformation success. Most relevant for this study, Carpenter et al. (2004) meta-analyzed 37 studies confirming that top management team characteristics indeed predict strategic choices, lending robust empirical support to applying upper echelons logic to AI adoption decisions. These contemporary developments demonstrate that executive characteristics remain powerful predictors of organizational choices even as the nature of strategic decisions evolves with technology.

If AI's impact depends on who sits around the boardroom table, then understanding AI in governance requires understanding board composition. The same technology introduced into different boardroom contexts may yield profoundly different outcomes—not because the algorithms change, but because the humans interpreting them do. Upper echelons theory reminds us that technology effects are never purely technical; they depend fundamentally on the humans interpreting and applying that technology. This theory provides the foundation for RQ2's examination of how board composition influences director perspectives on AI sensemaking. While upper echelons theory focuses specifically on the characteristics of individuals in top management positions, organizational context—industry dynamics, firm size, resource availability—provides the environmental setting within which these individual characteristics operate. RQ2 explores both board composition effects (as upper echelons theory predicts) and broader organizational contextual influences on director sensemaking.

Socio-Technical Systems: AI as Governance Disruptor. Trist's (1981) socio-technical systems theory offers a crucial insight: technology and social arrangements are not independent. You cannot optimize technology without considering the social system it inhabits, and you cannot understand social dynamics without acknowledging the technologies embedded within them. The core principle—that optimal performance requires joint optimization of technical and social elements—has been validated across numerous organizational contexts.

For boardrooms, this means AI is not a neutral tool that directors simply use. It's an element of a broader governance ecosystem that will reshape—and be reshaped by—that ecosystem. When AI enters the boardroom, it may affect who has influence (directors who understand algorithms may gain power), how deliberation occurs (discussions may become more data-centric), what information flows matter (real-time dashboards may displace periodic reports), and how accountability is structured (responsibility for AI-informed decisions may become ambiguous).

Socio-technical systems thinking has remained influential in understanding organizational technology adoption. Baxter and Sommerville (2011) updated socio-technical principles for contemporary contexts, emphasizing that technology design must account for social, organizational, and human factors to be effective. More recently, scholars have applied socio-technical frameworks specifically to AI and algorithmic systems. Benbya et al. (2020) demonstrate how AI fundamentally reconfigures organizational work practices and structures, requiring careful attention to the interplay between technical capabilities and social arrangements. Raisch and Krakowski (2021) extend socio-technical thinking to artificial intelligence, arguing that successful AI implementation requires treating AI not as a technical fix but as a socio-technical system that transforms organizational relationships and work practices. These contemporary applications confirm that Trist's core insight—that technical and social systems must be jointly optimized—remains essential for understanding AI adoption even decades after the original formulation.

Supporting Theories: Extending the Framework

While the three primary theories establish the core dimensions of analysis, four additional theoretical contributions provide specific conceptual tools that operationalize and extend this

foundation. Think of these as specialized instruments that help us examine particular mechanisms and processes more precisely.

Agency Theory: The Accountability Challenge. Fama and Jensen's (1983) agency theory framework explains how governance structures address the separation of ownership and control. Their key insight—that effective governance separates decision management (initiation and implementation) from decision control (ratification and monitoring)—has become foundational to understanding corporate board functions. But this framework becomes deeply problematic when AI enters the picture.

If boards cannot fully understand AI recommendations due to algorithmic opacity, questions arise about whether they can truly exercise decision control. They may ratify decisions, but can they genuinely evaluate them? If management controls the configuration, training data, and inputs of AI systems that boards rely on for oversight, questions arise about whether the crucial separation between management and control has subtly eroded. Agency theory provides the conceptual vocabulary for articulating these governance challenges that emerge when algorithmic systems mediate the principal-agent relationship. It supports our bounded rationality and socio-technical analyses by highlighting specific governance mechanisms—monitoring, evaluation, ratification—that AI may disrupt.

AI Decision-Making Typology: Operational Precision. Shrestha et al.'s (2019) framework provides essential clarity for applying bounded rationality theory to AI contexts. They distinguish three modes: AI-assisted decision-making (where AI supports information gathering but humans retain authority), AI-augmented decision-making (where AI generates recommendations humans consider alongside other inputs), and AI-automated decision-making (where algorithms both recommend and execute with minimal human involvement).

For strategic board decisions—inherently ambiguous, value-laden, and consequential—full automation is inappropriate. But both assisted and augmented modes are increasingly relevant. A board using AI for risk assessment operates in assisted mode; a board considering AI-generated strategic recommendations operates in augmented mode. These distinctions matter profoundly because governance implications differ. Assistance preserves traditional deliberation structures; augmentation raises questions about how directors maintain genuine authority when sophisticated algorithms propose specific courses of action. The risk is not that directors lose formal authority—they retain the power to decide—but that they gradually defer

to algorithmic recommendations they cannot fully evaluate, transforming decision ratification into ritualistic approval. This typology supports bounded rationality analysis by enabling precise specification of which cognitive functions AI extends and which remain human, and where new cognitive constraints emerge in each mode.

Professional Engagement Strategies: Practical Insight. Lebovitz et al.'s (2021) study of radiologists using AI diagnostics provides crucial empirical grounding. They examined professionals facing high-stakes decisions with partial understanding of the AI systems informing those decisions—precisely the situation boards may face. They identified three engagement strategies: leveraging (treating AI as valuable input while maintaining critical distance), dismissing (selectively rejecting AI when confidence is low), and working around (extracting value while compensating for limitations).

The key insight is "calibrated trust"—knowing when to trust AI and when to question it. Radiologists developed this through repeated observation: seeing many cases where AI succeeded or failed taught them its reliability boundaries. But boards face different conditions. Strategic decisions occur infrequently and often in novel contexts, providing fewer opportunities to calibrate trust through experience. Strategic decisions occur infrequently—major acquisitions happen every few years, not every few hours. Each decision is relatively unique rather than repetitive. A board cannot observe AI performance across hundreds of similar cases to develop calibrated trust. Whether and how directors might develop appropriate trust without repeated observation remains an empirical question.

Technology Structuration: The Temporal Dimension. Orlikowski's (1992) structurational model extends Trist's socio-technical systems thinking by emphasizing technology's temporal dynamics. Her concept of technology's "duality"—simultaneously shaped by and shaping human action—is essential for understanding that AI adoption is not a discrete implementation event but an ongoing process of mutual adaptation.

Initial choices about how to use AI seem provisional and changeable. But these choices gradually become routinized, written into governance procedures, embedded in director expectations, taken for granted. What began as experimental becomes "how we do things here," and at that point, the technology constrains future action. A board that starts using AI dashboards for strategy discussions may find, six months later, that directors no longer bring other forms of analysis because "everything's in the dashboard." The technology has structured

governance practice. This temporal perspective supports our socio-technical framework by highlighting that AI's governance impact unfolds through recursive cycles of human choice and technological constraint, not through one-time implementation. Methodologically, Orlikowski's work exemplifies how qualitative research can build rigorous theory from careful observation of technology use in organizational contexts.

Integration: How the Theories Work Together

These theoretical perspectives—three primary, four supporting—form an integrated analytical framework that enables nuanced examination of AI in boardroom governance.

The three primary theories establish the core dimensions of analysis. Bounded rationality explains the fundamental tension: AI extends cognitive capacity (processing more information, modeling complexity) while creating new constraints (interpretability challenges, validation difficulties). This is AI's double-edged nature. Upper echelons theory explains why this tension plays out differently across boardrooms: board composition—digital literacy, professional backgrounds, cognitive diversity— helps explain why directors in different boards perceive and respond to AI differently. Socio-technical systems theory explains that AI's impact emerges not from the technology itself but from the recursive interaction between algorithms and governance practices, power structures, and deliberative norms.

The four supporting theories operationalize these broad insights into specific mechanisms. Agency theory specifies which governance accountability mechanisms AI threatens (decision control, monitoring, evaluation). The AI decision typology specifies which cognitive tasks are affected in different implementation modes (assisted versus augmented). Professional engagement research provides empirical insight into how individuals develop working relationships with opaque algorithms through strategies like leveraging, dismissing, and working around. Technology structuration provides a dynamic view of how initial AI adoption choices become institutionalized governance practices over time.

2.2 Critical Challenges: Explainability and Performativity

The theoretical frameworks articulated above provide a conceptual architecture and sensitizing framework, rather than a fixed set of hypotheses, for understanding AI in governance. Two additional empirical insights from the AI literature—both critical for boardroom contexts—

complement this framework by highlighting specific technical and social challenges boards must navigate.

The Explainability Imperative: Navigating the Accuracy-Interpretability Tradeoff. Adadi and Berrada's (2018) survey on explainable AI reveals a fundamental tension central to bounded rationality concerns. The most accurate machine learning models—deep neural networks, ensemble methods—are often the least interpretable. They work, but we cannot easily explain how. Simpler models—decision trees, linear regressions—are transparent but may sacrifice predictive power. This is the accuracy-interpretability tradeoff.

For operational decisions, organizations often accept opacity in exchange for performance. A marketing algorithm that increases conversion rates by 15% doesn't require explanation if it delivers results. But governance is different. Directors have fiduciary duties requiring them to understand the basis of their decisions. They must explain choices to shareholders, defend them to regulators, and take responsibility if things go wrong. "The algorithm recommended it" is not a defensible governance position—legally, ethically, or practically.

This suggests that boards may face pressure to prioritize interpretability over maximum accuracy—accepting AI systems that are good enough rather than optimal but opaque. Yet this choice itself raises questions about when the preference for transparency becomes excessive caution that forgoes AI's full benefits, how much accuracy boards should sacrifice for interpretability, and who decides—directors, management, regulators. Adadi and Berrada's work clarifies that post-hoc explanations—attempts to interpret complex models after training—are approximations rather than complete accounts, a limitation directors must understand when evaluating AI recommendations. This technical constraint directly shapes the bounded rationality dynamics boards experience: AI promises to extend cognitive capacity, but only if directors can actually understand what the AI is doing. RQ1 investigates how directors reason about this tradeoff.

AI as Performative: How Algorithms Shape What Matters. Faraj et al.'s (2018) conceptual work challenges the view of AI as a neutral tool that organizations simply apply. They argue that AI is performative—it doesn't just process existing problems but shapes how problems are framed and what counts as relevant information. When an organization adopts an AI system, it implicitly adopts that system's assumptions about what matters, what variables to consider, what outcomes to optimize.

For boards, this raises questions that extend beyond the technical challenges of explainability. If AI systems emphasize quantifiable metrics—revenue growth, cost reduction, market share—strategic deliberation may gradually marginalize qualitative considerations like organizational culture, stakeholder relationships, or long-term sustainability. If algorithms optimize for financial outcomes, social and environmental concerns may receive less attention not through explicit choice but through gradual normalization of algorithmic framing.

The risk is not that AI makes wrong decisions but that it subtly narrows the frame within which decisions are made—and directors may not notice this narrowing because it happens gradually as AI becomes normalized in governance processes. A board might not consciously decide to focus less on employee wellbeing; rather, discussions naturally center on the metrics the AI dashboard highlights, and over time, what's not measured becomes invisible. This performative dimension connects directly to socio-technical systems concerns about how technology reshapes organizational practices and power structures in ways participants may not intend or recognize.

Together, these challenges—the technical constraint of explainability and the social dynamic of performativity—highlight why AI in governance is not simply a matter of adopting better tools. Boards must navigate fundamental tensions between accuracy and transparency, between algorithmic efficiency and human judgment, between quantitative optimization and qualitative wisdom. How directors understand and manage these tensions remains an open empirical question that this research addresses.

2.3 The Gap: What We Don't Know About AI in Boardrooms

This literature review reveals an important gap. While substantial research examines AI in organizational contexts (Brynjolfsson & McAfee, 2017; Davenport & Ronanki, 2018; Shrestha et al., 2019), and sophisticated theories explain corporate governance (Fama & Jensen, 1983; Hambrick & Mason, 1984), research at their intersection—specifically examining how directors make sense of AI for governance—remains limited.

Recent reports from professional organizations document organizational responses to AI. The Institute of Directors (2024) surveyed UK board members about AI governance practices—committee structures, policy frameworks, and oversight mechanisms. Deloitte (2024) examined whether boards have established AI governance processes. PwC (2024) provided

guidance on board oversight structures for AI systems. These practitioner studies establish that boards are creating governance structures for AI and provide valuable descriptive insights about what organizational governance arrangements are being put in place.

However, these practitioner reports address different research question than this study. They document organizational governance structures and practices—what committees exist, what policies have been adopted, what oversight mechanisms are in place. This research instead examines individual director sensemaking—how directors cognitively process AI's implications, how they reason about tradeoffs between accuracy and interpretability, how professional backgrounds and board composition shape these perceptions, how directors navigate the tension between algorithmic recommendations and human judgment. The distinction parallels the difference between studying organizational charts (governance structures) and studying how executives think (director cognition and sensemaking).

Moreover, existing work has important scope limitations. First, practitioner reports focus primarily on mature Western markets (UK, US, Western Europe) with limited attention to emerging markets or post-Soviet contexts where governance traditions, regulatory frameworks, and economic conditions differ significantly. Second, they employ closed-ended surveys that capture breadth (how many boards have AI committees?) but sacrifice depth of understanding about the interpretive processes through which directors make sense of AI. Questions like "Has your board established AI oversight?" reveal organizational structures but not directors' reasoning, cognitive tensions, or sensemaking dynamics. Third, they provide descriptive snapshots without theoretical grounding in cognitive theories (bounded rationality), compositional theories (upper echelons), or systemic theories (socio-technical systems) that would illuminate why directors perceive AI differently and how these perceptions vary across contexts.

The empirical gap in director sensemaking. We lack systematic evidence about how directors make sense of AI for strategic governance. The literature reviewed in Section 1.3 focuses on AI applications in operational and professional contexts—Lebovitz et al. (2021) examine radiologists, Faraj et al. (2018) discuss algorithmic work more broadly, and Shrestha et al. (2019) analyze organizational decision structures without specifically addressing board-level director cognition. Practitioner reports tell us what governance structures exist but not how directors reason about AI's benefits and risks, how they weigh competing values like accuracy versus interpretability, how their professional backgrounds and board composition

influence these perceptions, or what cognitive and deliberative dynamics emerge when algorithms enter boardroom discussions. RQ1 addresses this empirical gap by examining current engagement patterns and, crucially, director sensemaking—how directors think about and interpret AI.

The theoretical gap: Applying cognitive and compositional theories to director sensemaking about AI. Bounded rationality, upper echelons theory, and socio-technical systems theory all offer insights about how individuals make sense of complex technologies, how top management characteristics shape strategic perceptions, and how technologies reshape organizational dynamics. But these theoretical predictions remain untested in the specific context of board director sensemaking about AI. Do cognitive constraints manifest as bounded rationality theory predicts when directors encounter opaque algorithms? Does board composition shape AI perceptions as upper echelons theory suggests? How do socio-technical dynamics unfold when AI enters boardroom deliberations? These remain open theoretical questions requiring investigation with board directors themselves, examining not what governance structures they've built but how they think, reason, and make sense.

The Ukrainian context amplifies the gap. For the Ukrainian context specifically, the empirical and theoretical gaps are even more pronounced. No research—neither practitioner reports nor academic studies—examines how Ukrainian directors make sense of AI. Ukrainian directors operate under distinctive conditions that may differ from Western contexts—including wartime pressures, post-Soviet governance legacy, evolving regulatory frameworks, resource constraints, and EU integration dynamics. What aspects of board composition, organizational setting, and broader context Ukrainian directors themselves perceive as shaping their AI perspectives remains an open empirical question that this study explores through open-ended inquiry. Rather than imposing predetermined contextual categories, the research adopts an exploratory approach (particularly through RQ2) that allows directors to identify in their own words which aspects of their board composition, organizational context, and broader environment they believe matter for AI sensemaking. Directors may emphasize technical expertise on the board, industry dynamics, company resources, regulatory pressures, wartime conditions, governance culture, or entirely different factors that researchers have not anticipated. This methodological choice reflects the study's theory-informed but genuinely exploratory stance—using theoretical frameworks as sensitizing concepts to guide initial

inquiry while remaining fundamentally open to discovering what actually matters in directors' own sensemaking processes.

The practical gap matters particularly for Ukrainian boards. While Western practitioner reports recommend governance structures (AI committees, oversight policies), they provide limited insight into how directors should think about AI—how to develop informed perspectives when technical understanding is limited, how to weigh competing considerations, how to maintain genuine oversight rather than ritualistic approval, how to integrate AI recommendations into deliberative processes that preserve human judgment. Ukrainian boards need not just structural guidance but cognitive frameworks for director sensemaking. Without research examining how directors actually reason about AI—and what aspects of their composition and context they themselves perceive as salient—recommendations remain abstract. Understanding director sensemaking provides the foundation for practical guidance grounded in how directors actually think rather than how consultants believe they should think.

CHAPTER 3: RESEARCH METHODOLOGY

3.1 Research Design: Why Qualitative?

Methodology begins with epistemology—with a stance about what we're trying to know and how we can know it (Morgan, 2007). For this research, several considerations dictate a qualitative, exploratory approach.

First, the phenomenon itself is young and undertheorized. AI in boardroom governance is not a well-worn research path with established constructs and validated instruments. Quantitative research excels at testing hypotheses, but it requires knowing what hypotheses to test (Patton, 2002). Qualitative research excels at discovering what questions we should be asking—at generating and elaborating theory (Gioia et al., 2013). Second, the research questions are fundamentally about meaning and interpretation. This research does not ask "How much does AI improve decision quality?" (a quantitative question requiring objective measurement). It asks "How do directors make sense of AI? What does AI mean to them? How do they reason about its implications?" These are questions about subjective experience, about how intelligent actors interpret a new and ambiguous technology. Answering them requires hearing directors' voices, understanding their reasoning in their own terms, not reducing their perspectives to pre-determined scales (Lincoln & Guba, 1985).

Third, context is not noise to be controlled but signal to be understood. A board's experience with AI is inseparable from that board's composition, the organization's industry, the regulatory environment, the strategic challenges faced. These contextual factors don't confound analysis; they're central to it (Yin, 2018). Qualitative methods preserve this contextual richness, enabling us to see how AI engagement varies across contexts rather than averaging away the variation.

For these reasons, this research adopts a qualitative, exploratory design aimed at rich description, grounded understanding, and theory generation rather than hypothesis testing. The study is grounded in a pragmatist paradigm (Morgan, 2007), which prioritizes practical understanding over epistemological dogma. The goal is not to establish universal truths but to systematically document and analyze board directors' perceptions, experiences, and reasoning regarding AI in governance contexts.

3.2 Data Collection Methods

To balance depth of understanding with breadth of coverage, this study employs a sequential qualitative design using two data collection methods—semi-structured interviews and written qualitative surveys, following the logic of qualitative survey research described by Jansen (2010). This design leverages the strengths of each method while mitigating their respective limitations.

Semi-structured interviews serve as the primary data collection method, providing the rich, contextual, episode-based data essential for theory-building about sensemaking, power dynamics, and governance practice. Written qualitative surveys complement interview data by reaching additional directors who cannot commit to an interview but are willing to provide written responses, enabling triangulation by comparing patterns across data sources (Lincoln & Guba, 1985), and providing broader coverage of the Ukrainian director population to assess whether interview findings reflect wider patterns.

Both methods use an identical 9-question core instrument organized into four sections aligned with the research questions. Section A (Background) includes 2 questions establishing director characteristics relevant to upper echelons analysis. Section B (Current Engagement and Sensemaking) includes 3 questions addressing RQ1 by exploring whether and how boards discuss AI and how directors reason about benefits and risks. Section C (Board Characteristics and Context) includes 2 questions addressing RQ2 by probing how board composition and organizational context shape AI perceptions. Section D (Governance Implications) includes 2 questions addressing RQ3 by exploring anticipated or experienced impacts on board dynamics, accountability, and practice.

Surveys present the same 9 open-ended questions with simple written response-length guidance (for example, 3–10 sentences depending on question complexity). No additional sub-prompts are shown; respondents answer in their own words based on what they consider most relevant. The survey is implemented via Google Forms, enabling anonymous response collection and direct export for analysis. The survey is available in both Ukrainian and English. Estimated completion time is 25-35 minutes.

Interviews were conducted via Zoom to accommodate directors' schedules and geographic distribution across Ukraine. Each interview lasted 35-45 minutes. The researcher offered

interviews in either Ukrainian or English based on participant preference. With participant consent, interviews were recorded using Zoom's built-in recording feature. Recordings were transcribed verbatim using professional transcription services. For interviews conducted in Ukrainian, transcripts were professionally translated to English for analysis, with attention to preserving meaning and nuance. The researcher reviewed all transcripts for accuracy.

The complete interview guide and survey instrument are provided in *Appendix A*.

3.3 Sampling Strategy

The target population consists of individuals currently serving or recently serving (within two years) on corporate boards in Ukraine. This geographic focus enables depth of understanding within a specific governance context while acknowledging that Ukrainian boards operate under distinctive conditions—including wartime circumstances, regulatory frameworks, economic pressures, and corporate governance traditions—that may shape AI perceptions differently than in Western contexts.

Within this population, the study employs purposive sampling to recruit directors who can provide information-rich insights (Patton, 2002). Rather than seeking statistical representativeness, the sampling seeks strategic variation across dimensions theory suggests matter. Industry variation ensures the sample includes technology firms, financial services, manufacturing, retail, and other sectors relevant to the Ukrainian economy, enabling examination of whether AI perceptions differ by industry context. Digital literacy variation means recruiting both technologically sophisticated directors and those with traditional backgrounds, allowing exploration of whether tech-savvy directors tend to view AI more favorably, as upper echelons theory would suggest. AI engagement stage variation includes boards at different points—haven't discussed AI, actively considering, piloting tools, regular users—enabling comparison of how perceptions differ across adoption stages. Organizational size variation acknowledges that large firms have different resources and pressures than smaller ones, potentially affecting AI adoption patterns.

Recruitment employed multiple channels within Ukraine: LinkedIn outreach to Ukrainian directors identified through public board listings and business registries, professional networks including Ukrainian governance associations and director networks, the researcher's university network (American University Kyiv alumni and executive education participants), and

personal referrals within Ukrainian business communities. The recruitment message was offered in both Ukrainian and English to accommodate director preferences.

3.4 Data Analysis

Once data was collected, analysis began. This study uses a Gioia-inspired coding approach (Gioia et al., 2013) to structure the move from raw data to more abstract theoretical insight. Because the research starts from an articulated theoretical framework (bounded rationality, upper echelons theory, socio-technical systems), the analysis follows an abductive logic of theory elaboration rather than purely inductive grounded theory. The same coding procedures are applied to both interview transcripts and survey responses, as both yield narrative text data.

Analysis proceeds in three stages. First-order analysis begins by reading all responses (interview transcripts and survey texts) multiple times to become immersed in the data, then systematically coding using participants' own terms and phrases as codes (Gioia et al., 2013). When a director says "We can't explain AI decisions to shareholders," the code is exactly that—participant language, not researcher interpretation. This stage generates potentially hundreds of first-order codes, capturing the full variety of what directors say before imposing any interpretive structure.

Second-order analysis asks: What are these various first-order codes instances of? What broader categories or themes do they represent? This requires interpretation—moving from participants' language to more abstract concepts (Gioia et al., 2013). For example, first-order codes like "can't explain decisions," "legal liability concerns," and "fiduciary duty requires understanding" might be instances of a second-order theme called "Accountability Ambiguity." This stage produces a more manageable number of themes—perhaps 15-25—each defined clearly with inclusion/exclusion criteria. At this stage, existing theories act as sensitizing concepts rather than fixed categories: some themes may reflect anticipated ideas (for example, bounded rationality concerns), while others may extend, refine, or challenge the initial framework.

Third-order aggregation further abstracts themes into aggregate dimensions—overarching theoretical concepts that represent the study's core contributions (Gioia et al., 2013). For instance, themes about accountability ambiguity, interpretability challenges, and director expertise gaps might aggregate into a dimension called "Governance Capability Constraints"—

the idea that AI strains boards' ability to perform their governance functions effectively. This dimension is interpreted in dialogue with the theoretical frameworks (bounded rationality, socio-technical systems) while remaining grounded in what directors actually said, consistent with an abductive, theory-elaborating approach.

The Gioia-inspired approach used in this study emphasizes transparency through data structure representations that show connections between first-order codes, second-order themes, and aggregate dimensions (Gioia et al., 2013). These visualizations help the researcher see relationships and refine the analytic structure, while enabling readers to trace the logic from raw data to theoretical claims. For responses provided in Ukrainian, translation to English was performed during the coding phase, with attention to preserving meaning and nuance. Key quotes used in findings note when translation has occurred.

3.5 Research Quality

Qualitative research requires quality criteria appropriate to its interpretivist foundations. Lincoln and Guba (1985) propose four criteria: credibility, transferability, dependability, and confirmability.

Credibility addresses whether findings accurately represent the phenomenon. Multiple strategies enhance credibility: triangulation through multiple participants and two data collection methods (interviews and surveys) provides cross-validation (Lincoln & Guba, 1985); thick description with rich, detailed accounts enables readers to assess whether interpretations are warranted (Geertz, 1973); negative case analysis prevents premature conclusions by examining instances that don't fit patterns (Lincoln & Guba, 1985); and potential member checking by sharing preliminary findings with selected participants provides a reality check on whether interpretations resonate with their experiences (Lincoln & Guba, 1985).

Transferability addresses whether findings apply beyond this sample. Qualitative research doesn't claim statistical generalizability but rather analytic generalization (Yin, 2018)—showing how findings inform theory applicable to other contexts. Transferability is enhanced through rich contextual description enabling readers to assess similarity to other settings, and through theoretical connections that suggest patterns may operate beyond the specific sample (Lincoln & Guba, 1985).

Dependability addresses whether the research process has been transparent and systematic. This is established through audit trails (records of raw data, coding decisions, analytic memos), methodological transparency through detailed procedure description, and reflexivity acknowledging researcher positionality and potential influences on interpretation (Lincoln & Guba, 1985).

Confirmability addresses whether findings emerge from data rather than researcher biases. This is enhanced through data grounding with extensive verbatim quotes demonstrating claims arise from what participants said, reflexive awareness documented in memos helping the researcher notice when preconceptions might influence interpretation, and peer review by advisors providing external perspectives that challenge interpretations (Lincoln & Guba, 1985).

3.6 Ethical Considerations

Informed consent is obtained at the beginning of each interview and through the survey's opening page. Participants receive clear explanation of the research purpose, what participation involves, how data will be used, and confidentiality protections. For interviews, verbal consent is obtained and recorded. For surveys, proceeding past the consent page constitutes consent. Participants may withdraw at any time without penalty.

Anonymity is essential for honest responses. No identifying information is collected beyond what's necessary for analysis. Interview participants are assigned numerical codes (R01, R02, etc.). Survey responses contain no names, company names, or email addresses (except optionally and separately for those wanting research findings). Quotes are attributed only by participant number. Confidentiality extends beyond anonymity—any details that could identify a company are removed or disguised when quoting responses.

Data security means interview recordings and transcripts are stored in password-protected files accessible only to the researcher and thesis advisors. Survey data collected via Google Forms is immediately downloaded and stored in encrypted files, then deleted from Google servers. All data will be retained for five years as per IRB requirements, then permanently deleted.

Respect for participants' time and expertise shapes every design aspect. Interview questions are clear and well-crafted, interview length is explicitly stated (35-45 minutes), survey length is clearly communicated (25-35 minutes), and invitations emphasize genuine interest in directors'

insights. Reciprocity acknowledges that directors give time and insights without compensation; in return, the research commits to producing knowledge that benefits boards collectively and offers participants a summary of findings if desired.

Special consideration is given to the Ukrainian context. Given ongoing wartime conditions, questions avoid topics that could create security concerns or require disclosure of sensitive strategic information. The interview guide and survey are available in both Ukrainian and English to ensure directors can respond in their preferred language. The researcher, based in Ukraine, understands the local context and can conduct interviews with cultural sensitivity.

3.7 Limitations

All research has boundaries. Sampling limitations are inherent in non-probability purposive sampling—the sample is not statistically representative, and participants likely differ from non-participants, potentially being more interested in AI or more comfortable with technology.

Geographic concentration in Ukraine means findings reflect the Ukrainian governance context specifically. Ukrainian boards operate under distinctive conditions—including ongoing wartime circumstances, post-Soviet governance legacy, emerging regulatory frameworks, and specific economic pressures—that may differ significantly from Western European or North American contexts. While this focus enables deep contextual understanding, transferability to other geographic contexts should be assessed carefully (Lincoln & Guba, 1985). The study contributes to understanding AI in governance within Ukrainian and potentially broader post-Soviet and emerging market contexts, but may not fully reflect patterns in mature Western markets.

Self-report bias means directors' responses reflect perceptions that may be shaped by social desirability or hindsight rather than objective reality. Directors may present themselves and their boards in favorable light, or retrospectively rationalize decisions. The study mitigates this through assured anonymity and by focusing on directors' genuine sensemaking rather than claiming to measure objective board effectiveness.

Cross-sectional design provides a snapshot rather than observing attitude evolution over time. The study captures directors' perspectives at one moment, unable to track how perceptions change as experience with AI accumulates or as the technology itself evolves.

Response depth variability is inevitable—some participants will provide rich detail while others are briefer, some interviews will generate deep discussion while others remain surface-level, and some survey responses will be expansive while others are terse. The research design anticipates this through multiple participants ensuring sufficient rich data even if some responses are less detailed.

Method effects may emerge between interviews and surveys. Interviews enable probing and rapport-building that may elicit different content than written surveys. The study acknowledges this by treating each method's data appropriately and noting method effects if they appear during analysis.

Researcher interpretation is unavoidable in qualitative research—different researchers might identify somewhat different themes or emphasize different aspects of the data. The study enhances interpretive rigor through a systematic, Gioia-inspired coding procedure, peer debriefing, and transparent data presentation enabling readers to assess whether interpretations are warranted (Gioia et al., 2013).

Language and translation introduce potential for nuance loss. For interviews and surveys conducted in Ukrainian, translation to English for analysis and reporting may not perfectly capture original meanings. The study mitigates this through professional translation services, researcher review (the researcher is fluent in both languages), and noting when quotes are translated.

These limitations don't invalidate findings but bound them. The research provides valuable exploratory insights into an understudied phenomenon within a unique governance context, generating frameworks and hypotheses that future research can test more rigorously.

CHAPTER 4: FINDINGS

4.1 Analytical Approach

This chapter presents findings from 22 Ukrainian board directors across banking, energy, technology, healthcare, FMCG, logistics, and public sectors. Analysis followed Gioia methodology (Gioia, Corley, & Hamilton, 2013), moving systematically from directors' own language to interpretive themes to theoretical insights.

The approach progressed through three stages. First, directors' exact phrases were captured as first-order concepts. When one director stated, *"For me, the biggest benefit is cutting through information overload,"* this became a distinct concept preserved in their words. This stage generated approximately 390 concepts across all responses, maintaining fidelity to how directors themselves made sense of AI.

This study did not aim for theoretical saturation in the classic grounded-theory sense, which would require continued sampling until no new concepts emerge (Glaser & Strauss, 1967). Given the exploratory nature of the research and practical constraints on elite access, the goal was sufficient diversity and depth to surface meaningful patterns and generate empirically grounded propositions. The sample's variation across industries, technical backgrounds, and AI engagement levels provided multiple perspectives on each research question. By the final responses, similar themes were recurring across participants, suggesting adequate coverage—though findings should be understood as propositions warranting further investigation rather than conclusions generalizable to all directors.

Second, these concepts were systematically compared within and across respondents to identify patterns. When multiple directors used phrases like "cutting through overload," "summaries of documents," and "highlighting important data," these were grouped into the interpretive theme "Information Filtering." This iterative refinement—testing labels against raw data, splitting or merging themes as needed—continued until achieving stable, coherent groupings. The process yielded approximately 20 themes across three research questions.

Third, relationships among themes were examined to connect them with theoretical constructs. For example, themes around filtering information, generating scenarios, and integrating data all described ways AI might extend board cognitive capacity. These patterns aligned with

established theory—in this case, Simon's (1947,1997) bounded rationality—though directors themselves used different language. This stage produced six aggregate dimensions grounded in empirical patterns while engaging theoretical frameworks.

Data structure tables showing this progression appear in Appendix B. Several steps ensured rigor: maintaining detailed audit trails, regularly returning to raw data to verify interpretations stayed grounded, actively seeking contradictory evidence rather than ignoring it, and using extensive quotes so readers can assess whether interpretations fit evidence.

Analysis compared patterns across interviews (n=5) and surveys (n=17) to assess potential method effects. Core themes emerged consistently across both data sources. Accountability concerns appeared in 100% of both surveys and interviews. Wartime context influences appeared in 88% of surveys and 100% of interviews. Regulatory pressures, security concerns, and the dialectical tension between AI benefits and risks appeared with similar frequency in both pools, increasing confidence that central findings reflect genuine patterns rather than method artifacts. Minor differences in emphasis emerged. Survey respondents more frequently used phrases like “human in the loop” (82% vs 40%) and offered metaphors such as 'magic number on screen' and 'PowerPoint stories.' Interview participants elaborated more on reconstruction priorities (80% vs 6%). These differences reflect variation in emphasis rather than contradictory findings—no theme appeared exclusively in one source that contradicted themes from the other. Given this consistency, interview and survey data were analyzed as an integrated pool. The convergence across data sources strengthens confidence in the findings' trustworthiness.

4.2 Participant Characteristics

Data collection involved written surveys (n=17) and interviews (n=5) with Ukrainian board directors selected for diversity across industries, experience, and technical backgrounds. Six participants came from banking and financial services, three from technology and IT, three from energy and utilities, three from manufacturing and logistics, two from healthcare, two from FMCG and retail, and three from public sector organizations. Board experience ranged from five to over twenty years, with most serving on multiple boards and sitting on audit, risk, or strategy committees.

AI familiarity varied considerably. Three participants described basic knowledge—general awareness without deep understanding. Twelve reported moderate familiarity—understanding AI concepts in business terms, overseeing implementations, using AI tools regularly, but not being technologists themselves. Seven reported high expertise, including current or former CTOs who had built AI systems. This range provided insights into how technical capability shapes what directors see, worry about, and recommend.

All participants governed organizations operating in or significantly exposed to Ukrainian and Central/Eastern European markets. Most explicitly referenced navigating wartime conditions—damaged infrastructure, displaced teams, volatile demand, heightened security threats. This shared context shaped many perspectives, making AI simultaneously urgent (for better forecasting under volatility) and risky (given degraded data and adversarial threats).

4.3 How Directors Make Sense of AI (RQ1)

Analysis of directors' responses regarding AI benefits, concerns, and current board discussions revealed a consistent pattern: directors simultaneously see AI as potentially solving real problems boards face while creating new governance challenges. This tension—AI as both helpful and threatening—structures their cautious approach.

4.3.1 AI as Practical Solution to Board Challenges

Directors consistently described AI as potentially addressing three problems boards regularly encounter.

First, too much information. Board materials arrive days before meetings, sometimes hundreds of pages across finance, risk, operations, markets, and strategy. Reading everything thoroughly is difficult; determining what actually matters is harder. One banking director captured this directly:

"For me, the biggest benefit is cutting through information overload. AI can help sort, summarize, and highlight what's really important, instead of us digging through huge packs."
(R01)

An FMCG director elaborated, connecting information management to concerns about narrative control:

"The main benefit is getting clearer, faster insight from a lot of messy data. In FMCG, we deal with many SKUs, channels, and regions, especially now with constant changes due to the war. AI can help highlight where we lose margin, where promotions really work, and where demand is shifting. This can make board discussions more concrete and fact-based, instead of relying only on PowerPoint stories." (R03)

The phrase "PowerPoint stories" suggests directors wonder whether management presents information selectively. AI filtering might give boards more independent access to what's actually happening, not just what management chooses to highlight.

Second, too much backward focus. Current board processes typically explain last quarter's performance, analyze historical trends, understand what already happened. Directors described wanting to look forward instead, especially in volatile environments where past patterns offer limited guidance for future decisions. A technology director explained the desired shift:

"Boards can move from backward-looking reports to more forward-looking scenarios, with probabilities and ranges instead of single numbers. AI can help detect weak signals—emerging customer behavior, operational stress, or market shifts—that we might otherwise miss until it's too late." (R04)

Directors used phrases like "seeing around corners," "early warning signals," and "detecting weak signals"—all suggesting desire to anticipate rather than merely react. Ukrainian wartime context amplified this need. When infrastructure gets destroyed overnight and supply chains collapse unpredictably, boards need tools for navigating fundamental uncertainty rather than extrapolating stability from history.

Third, too much fragmentation. Finance reports arrive separately from operational data, market intelligence through different channels than risk dashboards, strategic initiatives discussed apart from resource constraints. Traditional reporting structures fragment information by function, forcing directors to mentally assemble pieces into coherent pictures. This integration becomes increasingly difficult as organizational complexity grows. A banking executive described the challenge:

"AI can help us integrate huge volumes of data—credit, liquidity, customer behavior, macro indicators—and translate that into scenarios the board can actually work with. For a bank

operating in and around Ukraine, early warning on risk concentrations, customer distress, or operational bottlenecks is extremely valuable." (R07)

An energy director described similar struggles where multiple risk domains intersect—commodity prices, geopolitical events, operational incidents, climate exposures—with AI potentially helping "digest large data sets from trading, logistics, production, and external sources and highlight patterns or stress points that traditional reports miss."

These three patterns—information overload, backward orientation, data fragmentation—explain AI's appeal. Directors experience boards as facing specific, recurring challenges. AI offers potential help with all three: filtering what matters from overwhelming flows, enabling forward-looking scenario analysis, and synthesizing insights across organizational silos.

These patterns align with what Simon (1947, 1997) described as bounded rationality—the observation that decision-makers face cognitive limits on processing information, making calculations, and integrating complex systems. But directors didn't use this academic language. They talked about being "buried in board packs," needing to "see around corners," and wanting to "connect the dots." The theoretical frame helps interpret their practical experience without imposing concepts foreign to how they actually think.

Importantly, directors consistently described AI as augmenting human judgment rather than replacing it. They used phrases like "support tool," "additional input," "helps us see"—never suggesting AI should make decisions independently. This distinction becomes critical when examining their governance recommendations.

4.3.2 AI as Source of Governance Problems

While recognizing AI's potential benefits, directors simultaneously articulated concerns about new governance risks. These concerns emerged around three questions: Can boards explain AI-influenced decisions? Who is accountable when things go wrong? Will directors lose critical thinking capacity?

The explanation problem surfaced repeatedly. Boards must justify major decisions to regulators, shareholders, and other stakeholders. This accountability requires articulating

reasoning in understandable terms. When questioned—"Why did you approve this investment?" or "Why accept this risk?"—directors need answers comprehensible to intelligent generalists, not just technical specialists.

AI threatens this explainability. A banking CTO with extensive regulatory experience stated the problem directly:

"My biggest concern is the black box problem—if a director can't explain how a decision was reached, how can we fulfill our fiduciary duties? Banking regulators, especially the National Bank, will ask us to justify every material decision, and 'the algorithm said so' won't be acceptable." (R05)

The phrase "the model said so" recurred across multiple directors as shorthand for governance failure—a scenario where boards defer to algorithmic authority they neither understand nor control. Directors distinguished sharply between AI as tool (acceptable) and AI as opaque oracle (unacceptable), worrying that technical complexity could collapse this distinction in practice. One noted simply: "If we can't explain how a recommendation was produced, it's hard to stand behind it when something goes wrong."

The accountability problem differed but connected. Even if directors could explain what AI recommended, questions remain about who bears responsibility when AI-informed decisions produce bad outcomes. Traditional governance assumes clear accountability chains—boards delegate to management, management implements, boards oversee. AI complicates this clarity:

"The accountability question troubles me most: if a board approves a strategy based on AI analysis and it fails, who is responsible? The directors who relied on the tool? The management who selected it? The vendor who built it? Our legal and regulatory frameworks don't have clear answers yet." (R05)

Healthcare directors raised similar concerns where AI recommendations could affect patient outcomes but liability frameworks remain unclear. Who answers when treatment decisions follow flawed algorithmic advice? The treating physician? The hospital board that approved the system? The technology vendor? This ambiguity troubled directors because it potentially undermines the accountability structures that make governance meaningful.

The judgment erosion problem involved AI's effects on board behavior and culture rather than its technical properties. Directors worried about over-reliance—that boards might gradually treat AI outputs as authoritative rather than provisional, losing their own critical thinking capacity. An e-commerce CEO used vivid language:

"Biggest worry is that people start trusting the 'magic number' on the screen more than common sense. In e-commerce, data can be biased by promotions, supply issues, or random events, and AI doesn't always know the story behind the spike." (R13)

The "magic number" metaphor captured concern that sophisticated presentations might suppress critical engagement. If AI dashboards look scientifically rigorous, directors without technical confidence might hesitate to question them, particularly when challenging requires admitting limited understanding. Several directors described potential "deskilling"—boards becoming less capable of independent judgment as they grow accustomed to algorithmic decision support.

These three concerns—explanation difficulties, accountability ambiguity, judgment erosion—converge on a fundamental worry: AI might undermine conditions enabling effective board functioning. Boards must articulate reasoning to stakeholders, maintain clear responsibility lines, and preserve critical judgment that questions rather than validates. Directors saw AI as potentially threatening all three.

Yet importantly, directors presented these concerns not as reasons for rejection but as design criteria for responsible adoption. The governance risks they articulated became, in their later recommendations, the governance safeguards they insisted upon. Understanding AI as both helpful and threatening positioned directors to approach adoption neither naively enthusiastic nor reflexively resistant, but thoughtfully experimental.

4.3.3 AI as Data and Security Vulnerability

Beyond governance process concerns, directors emphasized how environmental and infrastructural challenges could undermine AI's value or create new risks. These concerns centered on data quality and cybersecurity, both particularly acute in Ukrainian wartime contexts.

Data quality emerged as persistent worry. AI systems require high-quality, representative data, but organizational environments—especially in Ukraine—often provide data that is incomplete, inconsistent, or systematically distorted. An FMCG director captured the core problem: "If the data is wrong or incomplete, AI can push us in the wrong direction very confidently."

Wartime conditions particularly exacerbate data challenges. One supervisor noted:

"The biggest concern is that models may look convincing but rely on bad or incomplete data, especially in Ukraine where the war distorts statistics. There is a risk that some directors will over-identify AI with objective truth and underestimate human judgment." (R18)

The problem extends beyond missing data to include systematic bias and "model drift"—where models trained on historical patterns become invalid when underlying realities shift. Pre-war training data becomes irrelevant for post-invasion conditions. A banking CTO summarized bluntly: "Feed garbage in, get garbage out, no matter how sophisticated your AI is." Directors across sectors described how legacy systems with inconsistent data standards, war-disrupted patterns, and incomplete regional coverage create data quality problems that limit AI reliability.

Cybersecurity concerns operated at multiple levels. Directors worried about AI systems becoming high-value targets for adversaries seeking to steal information, manipulate decisions, or disrupt operations. A security executive with military background articulated the most comprehensive threat model:

"Any AI system processing board-level information becomes a high-value target—not just for cyber attacks, but for intelligence collection by hostile actors. If AI systems are cloud-based or involve third-party vendors, you're creating pathways for information to leak or be intercepted. There's also the question of system integrity—how do we know the AI hasn't been compromised and is feeding us manipulated information?" (R06)

Healthcare directors emphasized data privacy risks: "Any data breach due to AI could lead to catastrophic consequences." A fintech executive raised concerns about external AI services: "Using tools like ChatGPT raises questions about where prompts and outputs are stored and how easily confidential ideas could leak."

The cybersecurity worry has three dimensions: operational risk (system compromise leading to data theft), integrity risk (adversaries biasing AI outputs to influence decisions), and availability risk (systems disabled when boards have become dependent). For Ukrainian directors operating under ongoing military conflict with sophisticated cyber warfare, these represent operational realities rather than theoretical concerns.

These data and security challenges revealed a sobering reality: AI's governance value depends fundamentally on information quality and infrastructure security. In stable, well-digitized environments with mature data governance, these challenges might be manageable. But Ukrainian directors operate differently—occupied territories create data blind spots, war disruptions invalidate historical patterns, legacy infrastructure compounds technical problems, and sophisticated adversaries actively target decision-making systems. Under these conditions, data and security vulnerabilities shift from manageable risks toward fundamental constraints on what AI can reliably deliver.

4.3.4 The Dialectic of Promise and Risk

Directors see AI as both promise and problem. They described boards struggling with information overload, backward focus, and fragmented data—AI could help with all three. But AI also creates new risks: boards might not be able to explain decisions, accountability becomes unclear, directors might stop thinking critically, data quality might be poor, and systems might be vulnerable to attack.

This wasn't enthusiasm turning to worry over time. Directors expressed both views simultaneously, often in the same response. One captured it perfectly: boards see the value but don't want decisions that look like they were made "because the model said so."

Current engagement reflects this ambivalence. Most boards have discussed AI, typically as operational tool whose outputs feed into materials, but few have adopted it widely for governance purposes. The typical stance is "interested but careful"—willing to explore but not ready to commit fully. Directors want better understanding and proper governance frameworks before moving faster.

This caution makes sense given their operating context. Directors must explain decisions to regulators, maintain clear accountability, and keep systems working under pressure. Ukrainian

wartime conditions amplify these demands—when organizations face physical attacks, data problems, talent losses, and sophisticated adversaries, introducing complex AI dependencies requires careful justification. The benefits must survive encounter with real constraints.

The patterns directors described connect to established theories, though they didn't use academic language. When they talked about information overload and difficulty processing everything, this aligns with Simon's (1947, 1997) bounded rationality—the idea that decision-makers face cognitive limits. When they worried about explaining AI decisions and maintaining accountability, this connects to corporate governance scholarship on boards' oversight function. When they emphasized how wartime conditions make AI both more urgent and more risky, this reflects socio-technical systems thinking—technology effects depend profoundly on context.

But these theoretical connections emerged from analyzing what directors said, not from imposing frameworks beforehand. Directors spoke about "board packs," "magic numbers," "garbage in, garbage out," and "adversaries targeting systems." Theory helps interpret their experience without replacing it.

The next question becomes: why do some directors embrace AI faster than others? The answer involves who sits on the board and what pressures the organization faces.

4.4 Research Question 2: How Composition and Context Shape Perspectives

While all directors grapple with AI's promise-and-risk tension, their specific stances vary considerably. Analysis revealed this variation stems systematically from two sources: who sits on the board and what environment the organization operates in.

4.4.1 Board Composition: Technical Expertise as Double-Edged Sword

Directors universally emphasized that board composition fundamentally shapes AI discussions—their depth, quality, outcomes. This influence operates through complex dynamics involving technical expertise, traditional governance backgrounds, and the tension between capability and exclusion.

Even one or two directors with genuine technical expertise transforms board AI discussions. These individuals perform critical translation functions, converting technical concepts into

business language others can engage with. They also provide quality control, helping boards separate genuine AI capability from marketing hype. One director explained:

"When we have a couple of directors with digital or data backgrounds, the conversation becomes more concrete and less buzzword-driven—they can translate tech talk into business language for everyone else." (R01)

This translation capacity matters because AI proposals often arrive wrapped in technical jargon and optimistic framing. Directors without technical backgrounds struggle to evaluate claims, understand assumptions, or assess implementation risks. An IT executive described the quality difference: "Where technical experience is high, we discuss model architectures and data quality. Where it is low, the conversation quickly devolves into generalized discussions about 'digitalization' without concrete plans."

But technically capable directors do more than translate—they also shape agendas and influence outcomes. Directors comfortable with AI naturally drive adoption discussions, frame alternatives, steer decisions. This influence can be positive, preventing naive adoption or identifying valuable opportunities. But it also creates power dynamics within boards. An e-commerce CEO observed: "Tech people are excited and want to roll out new tools quickly; the ex-bankers and lawyers ask about risk, bias, and contracts."

Traditional directors—those with primarily accounting, legal, regulatory, or operational backgrounds—play equally critical balancing roles. These individuals typically approach AI more cautiously, focusing on risks, controls, regulatory implications. A banking accountant explained:

"Boards dominated by experienced bankers, accountants, and regulators tend to be cautious. I've seen meetings where they talk about innovation, and the traditional side keeps asking simple questions about controls, audit trails, and compliance." (R02)

This skepticism often proves valuable in retrospect. Directors described situations where enthusiastic AI proposals were slowed after traditional directors raised fundamental questions about controls or accountability—decisions that later proved wise when technical challenges emerged or performance disappointed. A security executive emphasized the value of risk-focused questioning: "Directors with operational or military backgrounds tend to be skeptical

of technology promises because we've seen things fail under pressure." He recalled asking about backup procedures when systems fail, questions that seemed paranoid until major disruption forced reliance on manual processes.

Effective AI governance appears to require both technical capability and governance discipline. Without technical expertise, boards risk naive adoption or uninformed rejection. Without governance discipline, boards risk moving too fast without adequate safeguards.

However, expertise diversity creates its own challenges. When technical directors dominate discussions with jargon or move quickly through complex material, less technical directors may disengage. One director noted candidly: "Tech-comfortable directors may start dominating the discussion because they understand the models, while others, even very experienced, may talk less."

A healthcare director with primarily clinical background described this dynamic vividly:

"Our board is mostly doctors and healthcare managers, so when technology comes up, we tend to look at each other waiting for someone else to explain it. We had one presentation about digital health tools, and honestly most of us were lost after the first few slides." (R11)

This "digital asymmetry" creates governance risks. If experienced directors with deep industry knowledge withdraw from discussions because they cannot follow technical details, the board loses valuable perspectives. One worried: "Conversations might shift from understanding the business to arguing about what the model says."

Given these dynamics, directors emphasized board diversity as essential for effective AI governance. Diversity here means specifically mixing technical, financial, operational, legal, and industry expertise to enable multi-angle evaluation. One summarized: "The mix of tech, finance, and commercial backgrounds means we can look at AI from several angles: feasibility, risk, and business value."

Some organizations formalized this through governance structures. An IT executive explained: "We created an AI Advisory Council with independent experts to bridge the knowledge gap for non-technical directors, which has helped significantly." A healthcare director described establishing a "Bioethics and Technology Committee to ensure both ethical and technical aspects are addressed."

What emerges is a nuanced picture: effective AI governance is unlikely from homogeneous boards—whether uniformly technical or uniformly traditional. Rather, it requires deliberately bringing multiple forms of expertise into productive tension, with structural mechanisms enabling communication across knowledge domains. Technical directors provide translation and quality control; traditional directors ensure governance discipline. The challenge lies in managing digital asymmetry so expertise differences enrich rather than fragment deliberation.

This pattern aligns with what Hambrick and Mason (1984) described as upper echelons effects—organizational outcomes reflect top management team characteristics. But directors didn't reference this theory. They spoke practically about "the mix of people around the table" and how "tech people drive agendas while bankers slow things down." The theoretical connection illuminates their observations while remaining grounded in their experience.

4.4.2 Organizational Context: War, Regulation, and Sector

Beyond composition, organizational and environmental context powerfully shapes directors' AI perspectives. Three dimensions emerged as particularly influential: wartime conditions creating paradoxical pressures, regulatory intensity demanding explainability, and sector-specific risk logics generating different assessments.

Wartime conditions create distinctive paradox: war simultaneously increases hunger for better forecasting (making AI attractive) and undermines data reliability while raising error stakes (making AI risky). A banking director explained:

"Our exposure to Ukraine means we operate in very volatile environment, so anything improving forecasting and risk sensing gets attention. At the same time, we work in regulated sectors, so supervisors expect explainable decisions. And given war and reconstruction context, there's strong focus on resilience, so AI is seen both as opportunity and as something that must be handled very responsibly." (R01)

A security executive described the mindset shift most starkly:

"Operating during wartime fundamentally changes your perspective on everything, including technology. We've had facilities hit by missiles, staff killed or wounded, massive infrastructure disruptions. The board's mindset is about resilience and survival, not optimization or

innovation. Every decision gets evaluated through 'does this help us stay operational under attack or does it create new vulnerabilities?'" (R06)

Directors described how war creates urgency for AI (need better forecasting under extreme uncertainty) while simultaneously undermining AI's foundations (degraded data quality, talent losses, security threats). This produces what one called "restrained appetite"—recognition of potential value coupled with awareness that practical constraints make implementation difficult and risky.

Regulatory intensity shapes AI adoption profoundly in supervised industries. In banking, energy, and healthcare, boards operate knowing major decisions must be justified to external authorities with enforcement power. This makes explainability mandatory, not merely desirable. A banking CTO explained:

"The regulatory environment is paramount—Ukraine's banking sector is heavily supervised, and the National Bank has become quite strict. Any new technology needs to fit within regulatory frameworks that weren't written with AI in mind, which creates uncertainty." (R05)

Healthcare directors described similar constraints: "In HealthTech, we deal with sensitive patient data, meaning our AI adoption must adhere to highest standards of data security, privacy, and clinical validation. This necessitates much slower, more deliberate, and risk-averse approach."

The regulatory imperative creates what directors called a "high bar" for AI adoption. Unlike consumer tech companies that can experiment rapidly and tolerate occasional failures, regulated entities must demonstrate—before deployment and continuously after—that AI systems meet safety, soundness, privacy, and explainability standards. This regulatory reality fundamentally shapes board risk appetite.

Sector-specific logics generate divergent AI assessments. While detailed sector analysis would be lengthy, several patterns merit highlighting. Banking directors emphasized model risk frameworks, regulatory constraints, and fraud pressures. Energy directors highlighted grid stability, physical infrastructure protection, and cyber threats. Healthcare directors emphasized patient safety, ethics, and liability. Public sector directors focused on corruption risks and

transparency requirements. Technology directors described AI as core competitive capability. FMCG and logistics directors emphasized supply chain optimization and margin pressures.

These sector differences mean identical AI capabilities get assessed very differently. A demand forecasting model might be mission-critical in FMCG (where volatility threatens margins), potentially valuable but risky in banking (where model risk and regulatory scrutiny loom large), and inappropriate in healthcare clinical decisions (where patient safety and liability dominate).

Ukrainian context specifically combines multiple pressures creating unique conditions. Directors explicitly identified this as distinctive—the combination of wartime stress, post-Soviet governance legacies, EU integration pressures, reconstruction imperatives, and economic volatility creates a configuration rarely experienced simultaneously elsewhere. A banking supervisor explained:

"The fact that we work with Ukrainian market in state of war very much shapes our view. AI is seen as chance to better manage risks and plan reconstruction, but we understand that data is fragmented and behavioral patterns have changed. We are also influenced by EU regulatory context, which expects responsible and transparent AI use." (R18)

An investment manager noted: "Since we represent international investors, the board's view is strongly shaped by investor expectations for responsible management and transparency. We also consider lack of stable legal framework for AI in Ukraine, which requires us to apply international standards to maximum extent."

What emerges is complex interplay between composition and context. Technical directors might push AI adoption aggressively in lightly regulated technology sectors but proceed cautiously in heavily supervised banking. Traditional directors' skepticism proves particularly valuable in high-stakes sectors like healthcare or critical infrastructure. Wartime conditions amplify both AI's appeal (for navigating uncertainty) and AI's risks (from degraded data and adversarial threats). Diverse boards operating in complex, high-stakes contexts must synthesize multiple forms of expertise and stakeholder concerns to navigate AI adoption responsibly.

These patterns reflect socio-technical systems thinking (Trist & Bamforth, 1951)—technology effects depend fundamentally on social and institutional context. AI is not neutral tool whose value can be assessed abstractly; rather, its implications emerge from interaction between

technical capability and environmental demands. Directors intuitively grasped this, even without using academic language. They described how "context shapes everything" and "what works in Silicon Valley won't work here."

The variation observed across directors thus reflects not arbitrary differences but systematic alignment between governance approaches and contextual demands. Understanding these patterns helps explain why AI adoption proceeds unevenly across organizations and suggests that prescriptive "best practices" divorced from context may prove unhelpful.

4.5 Research Question 3: Governance Implications and Practices

Directors described both how AI might change board dynamics and what practices would enable responsible adoption. The anticipated changes cluster around power, deliberation, and documentation. The recommended practices cluster around maintaining human authority, formalizing frameworks, ensuring transparency, controlling data, and building capability.

4.5.1 How AI Might Change Boards

Directors anticipated power shifts toward those who can interpret AI outputs. An FMCG director predicted: "Board discussions will be more centered around dashboards and scenarios. People who can read and question these AI-based insights will have more influence."

A technology director elaborated:

"AI will make boards more data-centric by default. Instead of long debates based on anecdotes, we'll spend more time challenging scenarios, assumptions, and model outputs. Directors who can bridge tech and business will naturally gain influence, so we need to make sure that doesn't silence others—everyone should feel able to question the numbers." (R04)

Directors worried about negative dimensions—that AI could create information asymmetries favoring management or specific board members. A security executive explained: "If AI is providing analysis and recommendations, whoever manages that AI effectively shapes what the board considers and how decisions are framed."

Directors anticipated deliberation processes becoming more scenario-driven and analytics-intensive. A banking director described expected changes:

"AI will change the tone of discussions more than formal structure at first. People comfortable with data may become more influential because they can interpret outputs and steer conversation. Meetings could become more driven by dashboards and scenarios, which is helpful but can also make softer issues—like culture or leadership quality—easier to overlook."
(R01)

An e-commerce CEO described anticipated shift toward "live" meetings: "I expect board meetings to become more 'live,' with data and scenarios refreshed on the spot instead of fixed PDFs." An IT services director noted: "Debates will shift from 'what is happening?' to 'given these modeled scenarios, which trade-offs do we accept?'"

But directors worried about negative side effects. A banking accountant expressed concern: "Conversations might shift from understanding business to arguing about what the model says." A logistics director worried: "The risk is that softer issues—like culture, supplier relationships, community impact—get less airtime because they're harder to quantify."

Directors anticipated enhanced documentation requirements. A banking director explained: "We may see more emphasis on documenting how AI inputs were used or challenged, which changes how accountability is recorded." A technology director elaborated: "I think we'll see more explicit documentation of how AI inputs were used in major decisions, which will become part of the board's 'paper trail.'"

These anticipated changes present both opportunities and risks. Better information, clearer trade-offs, more proactive oversight—these are potential benefits. But marginalization of non-technical voices, loss of qualitative judgment, accountability obscurity—these are potential costs. The governance practices directors recommend aim to capture benefits while mitigating risks.

4.5.2 What Practices Boards Need

Directors universally emphasized that AI must remain tool supporting human judgment, never replacing it. This "human-in-the-loop" principle emerged as the single most consistent

recommendation. A banking director stated: "The mood is 'interested but cautious': people see the value, but nobody wants decisions to look like they were made 'because the model said so.'"

A technology director explained the reasoning: "My main concern is misunderstanding, not the technology itself... we must be clear that AI is a tool to support judgment, not a way to outsource accountability."

Directors translated this into operational requirements. A banking accountant insisted: "Human decision-makers must remain fully accountable, with clear record of how AI outputs were used or challenged." A banking CTO specified: "There must be mandatory human review of any AI output before it influences board decisions."

Directors called for formal AI governance frameworks defining scope, responsibilities, standards, and processes. A banking director outlined core elements:

"I'd start with clear policy on how AI can and cannot be used in decision-making and reporting. Boards and management should know who is responsible for AI oversight, and key models should be validated, monitored, and periodically reviewed by independent experts." (R01)

A banking COO described organizational integration: "We need formal AI and model risk framework sitting alongside credit, market, and operational risk frameworks. That includes clear ownership, standards for development and validation, independent model risk management, and regular internal audit review."

Some directors described governance structures evolving to accommodate AI oversight: "A key safeguard is establishing dedicated AI Risk Committee within existing Risk or Technology Committee structure."

Directors emphasized transparency about how AI influences decisions. A banking director explained: "For major strategic and risk decisions, I'd keep clear 'human in the loop' principle and make sure minutes reflect how AI insights were used, not just final decision."

A fintech director described the practice: "On big strategic or risk decisions, we aim to capture in minutes how AI informed debate, where we overrode it, and why."

Directors also called for maintaining traditional analytical methods alongside AI, at least initially, to enable validation and build confidence. A banking CTO explained: "I'd insist on maintaining parallel traditional analysis methods initially, so we can cross-check AI outputs and build confidence gradually."

Directors emphasized data governance and infrastructure security as inseparable from AI governance. A security executive articulated the most stringent approach:

"First principle: complete control over infrastructure. No cloud services, no third-party data processing, everything on systems we own and can physically secure. Any AI touching board information must be treated like classified material." (R06)

A fintech director described selective restrictions: "We decided not to let external LLMs process any sensitive customer data without strict controls. We put in place AI policy covering data use and acceptable use of public tools like ChatGPT, especially around confidential information."

Directors recognized that AI governance failures often stem from data governance failures—biased training data, inadequate security, unclear provenance, inappropriate use of external services with sensitive information.

Directors emphasized building AI literacy as prerequisite for effective oversight. A banking director stated: "Directors need basic training so they can read AI-based materials without being either dazzled or intimidated."

A banking CTO elaborated: "Board members need training, not to become data scientists, but to ask right questions and understand limitations."

Directors emphasized training should enable informed questioning rather than making directors technical experts. One explained: "The board needs plain-language explanations of how models work and where they can fail."

These practices converge on coherent vision: maintain human authority through human-in-the-loop principles, develop formal frameworks with clear ownership and independent validation, ensure transparency about AI's influence on decisions, control data governance and infrastructure security, and build director capability through training.

Directors demonstrated sophisticated understanding that AI governance requires adapting existing governance principles—accountability, transparency, independent oversight, capability-building—to specific challenges AI presents. The emerging model treats AI neither as business-as-usual technology (requiring no special governance) nor as radically ungovernable (requiring rejection), but rather as powerful capability requiring thoughtful, structured, but feasible governance innovation.

These recommendations align with best practice emerging in model risk management (particularly in banking) and with corporate governance scholarship emphasizing boards' dual monitoring and advising roles. But again, directors didn't reference academic frameworks. They spoke practically about "keeping humans accountable," "having proper controls," and "making sure directors can ask intelligent questions." Theory helps organize their recommendations without imposing foreign concepts.

Directors see AI as both promise and problem. They described boards struggling with information overload, backward focus, and fragmented data—AI could help with all three challenges. But AI also creates governance risks: explanation difficulties, accountability ambiguity, potential judgment erosion, data quality concerns, and security vulnerabilities. This tension structures their cautious, exploratory approach—"interested but careful"—willing to explore but not ready to commit fully without better understanding and proper governance frameworks.

Board composition and organizational context fundamentally shape AI perspectives. Technical directors enable informed adoption through translation and quality control but can create power asymmetries. Traditional directors ensure governance discipline through skeptical questioning but may slow innovation. Diverse boards with effective communication mechanisms produce most balanced discourse. Contextually, wartime conditions create paradoxical pressures (AI more urgent yet more risky), regulatory intensity demands explainability (slowing adoption in supervised sectors), sector-specific logics generate divergent assessments (same AI capability valued very differently across industries), and Ukrainian context combines multiple stressors creating unique governance laboratory.

Directors anticipate AI will transform board functioning through power shifts toward data-fluent members, deliberation changes toward scenario-driven discussions, and enhanced documentation requirements. In response, they converge on governance practices: maintaining

human-in-the-loop as foundational principle, developing formal AI governance frameworks, ensuring transparency about how AI influences decisions, controlling data governance and infrastructure security, and building director AI literacy through training.

A coherent pattern emerges across findings: the directors in this study navigate genuine tensions between AI's capabilities and risks with sophistication reflecting neither naive enthusiasm nor reflexive resistance, but thoughtful experimentation grounded in governance realism. Ukrainian wartime context—extreme volatility, degraded data, security threats, reconstruction imperatives—serves as stress test for AI governance, revealing dynamics that may emerge more gradually elsewhere. The next chapter interprets these findings through theoretical lenses, connects patterns to broader debates, discusses limitations, and proposes directions for future research and practice.

CHAPTER 5: DISCUSSION

5.1 Introduction

The findings reveal directors in this study navigating a fundamental paradox: AI offers potential solutions to real cognitive constraints boards face while simultaneously threatening core governance capabilities. This dialectical understanding—AI as both cognitive extension and governance challenge—shapes their cautious, exploratory approach to adoption.

Three major patterns emerged from analysis. First, directors simultaneously recognize AI's benefits and risks, reflecting sophisticated sensemaking under technological uncertainty. Second, board composition and organizational context systematically shape where directors land on the enthusiasm-caution spectrum, with technical expertise, traditional governance backgrounds, wartime conditions, regulatory intensity, and sector-specific logics all shaping perspectives. Third, directors' governance recommendations demonstrate adaptive strategy preserving accountability principles through practice innovation rather than wholesale rejection or uncritical adoption.

The chapter addresses each research question in turn while maintaining thematic coherence. Section 4.2 interprets findings regarding directors' dialectical understanding of AI (RQ1), examining how bounded rationality, agency theory, and sensemaking perspectives illuminate their simultaneous recognition of benefits and risks. Section 4.3 analyzes how board composition and organizational context systematically shape perspectives (RQ2), drawing on upper echelons, socio-technical, and institutional theories. Section 4.4 discusses governance implications and emerging practices (RQ3), examining adaptive strategies directors propose. Section 4.5 then integrates insights across research questions, revealing interconnections among directors' understanding, contextual influences, and governance responses.

5.1.1 Distinguishing Theory-Led and Data-Led Insights

The theoretical framework guided initial analysis, but findings both confirmed and complicated prior expectations. Transparency about this interplay guards against confirmation bias.

Theory-confirmed patterns.

Several findings aligned with theoretical predictions from the study's core frameworks.

Directors' descriptions of cognitive challenges mapped onto Simon's (1947, 1997) bounded rationality framework as anticipated. The three recurring problems directors identified—information overload, backward-looking focus, and data fragmentation—correspond directly to the cognitive constraints Simon described. Directors' interest in AI as a tool to extend their analytical capacity, while remaining wary of creating new comprehension difficulties, reflects the bounded rationality logic that decision aids both extend and reshape cognitive limits.

The influence of professional backgrounds on AI perceptions confirmed upper echelons theory predictions (Hambrick & Mason, 1984). Directors with technical backgrounds perceived AI opportunities more readily and expressed greater confidence in managing implementation risks. Directors with traditional governance backgrounds focused more on accountability concerns and proceeded more cautiously. These patterns align with the theory's core premise that executive characteristics serve as filters shaping how leaders interpret strategic issues.

The context-dependency of AI assessments confirmed socio-technical systems theory expectations (Trist, 1981). Directors consistently emphasized that AI's value and risks cannot be assessed abstractly but depend on organizational circumstances—data quality, security environment, regulatory pressures, and existing governance culture. The way wartime conditions transformed AI from efficiency tool into potential vulnerability illustrates precisely the kind of technical-social interdependence the theory predicts.

Data-led surprises.

Three patterns emerged that extended beyond initial expectations.

First, the universality of dialectical thinking was stronger than predicted. Theory suggested directors might cluster into enthusiasts or skeptics; instead, nearly all held both orientations simultaneously regardless of background. This “interested but careful” stance was more pervasive than the framework anticipated.

Second, security and adversarial concerns were more prominent than expected. The specific ways wartime conditions shaped risk perceptions—particularly concerns about AI as attack vector—emerged inductively rather than from theoretical prediction.

Third, traditional directors' contributions proved more nuanced than upper echelons theory suggested. Rather than simply resisting AI, non-technical directors often identified governance

risks that technical directors underestimated—a complementary rather than oppositional dynamic.

These data-led insights suggest refinements to theory rather than rejection, and emerged from attending to patterns the initial framework did not anticipate.

5.2 Understanding AI's Dual Nature (RQ1)

Directors' simultaneous recognition of AI as cognitive extension and governance threat addresses the study's first research question: how do boards currently engage with AI and make sense of its benefits and risks? Three theoretical perspectives illuminate this dialectical understanding.

5.2.1 Bounded Rationality and Cognitive Extension

Directors' descriptions of AI's cognitive appeal align with Simon's (1947) bounded rationality theory. The three challenges directors identified—information overload, backward orientation, and data fragmentation (see Section 3.3.1)—map directly onto what Simon termed “computational limits”. This alignment suggests directors intuitively grasp bounded rationality's core insight without using academic language.

Recent scholarship extends this foundational insight in ways that illuminate directors' experiences. Čaić et al. (2018) demonstrate how AI systems can simultaneously alleviate and exacerbate cognitive constraints—reducing information overload through filtering while potentially creating new comprehension deficits if outputs become opaque. Jarrahi (2018) argues that effective human-AI collaboration depends on complementarity rather than substitution, with AI augmenting human judgment in ways that preserve rather than replace cognitive agency. Directors' emphasis on AI as "support tool" rather than decision-maker reflects intuitive grasp of this augmentation logic.

However, directors' recognition of governance risks complicates simple bounded rationality framing. While AI might extend information processing capacity, it simultaneously threatens oversight capacity if decision logic becomes opaque. Raisch and Krakowski (2021) describe this as the "comprehension-efficiency trade-off"—gains in computational efficiency may come

at cost of human understanding. Directors seeking AI that filters information without obscuring reasoning navigate this trade-off, attempting to capture cognitive benefits while preserving explainability necessary for accountability.

This finding extends bounded rationality theory in important way. Simon focused primarily on individual decision-makers; this study reveals how bounded rationality operates at collective level in board settings. Boards face unique cognitive challenges: multiple individuals must achieve shared understanding despite different expertise and limited time together. The complexity of modern organizations exceeds any individual director's comprehension, yet boards must somehow construct adequate understanding to provide effective oversight. AI offers potential to address this collective bounded rationality through mechanisms unavailable in Simon's era—real-time data integration, multi-scenario modeling, pattern detection across massive datasets.

Yet cognitive extension tools that reduce comprehension create new constraints, potentially substituting information overload with explanation deficit. The challenge for boards involves identifying forms of AI augmentation that address specific cognitive limitations—information filtering, scenario generation, data integration—without undermining capabilities boards need for effective oversight.

5.2.2 Agency Concerns and the Monitoring Function

Directors' governance concerns connect to agency theory (Jensen & Meckling, 1976), which views boards as monitoring mechanisms reducing information asymmetries between shareholders and management. AI introduces new agency complications that contemporary scholarship has begun examining. Packin and Lev-Aretz (2018) argue that AI in corporate boardrooms could either enhance monitoring (through independent data analysis) or undermine it (if management controls algorithmic tools), depending on governance arrangements.

Directors worried explicitly about this double-edged possibility. AI filtering might give boards more independent information access or might simply provide management with more sophisticated narrative control. The difference depends on who selects, configures, and operates AI systems—questions directors cannot always observe or verify. Kellogg et al. (2020) describe this as the "algorithmic accountability problem": when algorithms mediate

decisions, responsibility lines blur across humans (directors, management, data scientists) and artifacts (models, datasets, platforms).

More fundamentally, AI challenges agency theory's assumption that monitoring works through directors' independent judgment. If algorithmic recommendations become difficult to question—either because directors lack technical capability or because outputs appear scientifically authoritative—the board's monitoring function erodes. Lebovitz et al. (2021) document this dynamic in professional settings, showing how AI adoption can shift human roles from critical evaluators to validators of algorithmic outputs, particularly when expertise asymmetries exist.

Directors' emphasis on explainability, accountability, and human-in-the-loop principles directly addresses these agency concerns. By insisting AI remain tool rather than authority and building capability to question outputs, directors attempt to preserve monitoring function despite technological change. This suggests agency theory remains relevant but requires explicit attention to how information technologies reshape monitoring effectiveness and information asymmetries.

The findings also reveal new agency challenge: algorithmic intermediation of board-management information flows creates meta-level information asymmetry. Traditional agency theory assumes boards can observe management behavior through reports and presentations. But when AI systems controlled by management filter and synthesize information for boards, directors face uncertainty not just about operations but about the filtering process itself. This "second-order" information asymmetry—uncertainty about how uncertainty is being managed—compounds traditional agency problems.

5.2.3 Sensemaking Under Technological Disruption

Weick's (1995) sensemaking theory offers additional insight into directors' dialectical understanding. AI triggers sensemaking by disrupting familiar board processes and creating uncertainty about appropriate responses. Sensemaking, in Weick's framework, involves retrospective interpretation of ambiguous situations through which actors construct plausible narratives guiding action.

The language directors used reveals this active meaning construction. Phrases like "magic numbers," "PowerPoint stories," "the model said so"—these aren't neutral descriptions but meaning-laden metaphors helping directors grasp AI's implications. Beane (2019) describes how organizational members engage in "shadow learning" when confronting new technologies, developing unofficial knowledge through experimentation that may contradict official narratives about technology's benefits. Directors' metaphors suggest similar process—constructing working theories about AI's governance implications through experience rather than simply accepting vendor or consultant promises.

"Magic numbers" frames AI outputs as potentially deceptive despite scientific appearance—something that looks authoritative but may mislead. "PowerPoint stories" frames management presentations as potentially selective narratives rather than objective information. "The model said so" frames algorithmic authority as governance failure, establishing that deference to algorithms violates fiduciary responsibilities. These interpretive frames guide action: if AI produces magic numbers, directors should question rather than trust; if "the model said so" is unacceptable, human authority must be preserved.

Faraj et al. (2018) argue that AI's organizational implications depend fundamentally on how actors make sense of and enact AI capabilities within specific social and institutional contexts. Sensemaking theory helps explain variation across directors. Those with technical backgrounds possess different interpretive resources—experience building systems, debugging models, seeing implementations fail—shaping how they understand AI's governance implications. Those with traditional governance backgrounds draw on different experiences—regulatory investigations, audit failures, control breakdowns—leading to different concerns. Diverse boards bring multiple interpretive frames into conversation, enabling richer collective sensemaking than homogeneous boards achieve.

This sensemaking perspective suggests directors' understanding will continue evolving through experience. Weick emphasizes that sensemaking is ongoing and retrospective—meaning emerges through action and reflection rather than preceding it. Current "interested but careful" stance may represent temporary equilibrium pending additional encounters with AI systems. As Beane (2019) demonstrates, sensemaking around new technologies often produces multiple conflicting interpretations before settling into relatively stable shared understanding—a process requiring time and experimentation that directors are just beginning.

Directors' dialectical understanding of AI reflects sophisticated sensemaking combining multiple theoretical perspectives. Bounded rationality theory illuminates why AI appeals cognitively but must be qualified to account for comprehension-efficiency trade-offs. Agency theory explains governance concerns but requires extension to address algorithmic intermediation of information flows. Sensemaking theory reveals how directors actively construct meaning through metaphorical language and interpretive frames shaped by professional backgrounds. Together, these perspectives explain directors' simultaneous recognition of AI as both cognitive savior and governance threat—a tension they navigate through cautious exploration rather than resolving prematurely through wholesale adoption or rejection.

5.3 Shaping Forces: Composition and Context (RQ2)

While all directors grapple with AI's dual nature, where they land on the enthusiasm-caution spectrum varies systematically. These patterns address the second research question: how do board composition and organizational context shape perspectives? Three theoretical lenses prove particularly valuable.

5.3.1 Upper Echelons Effects and Epistemic Power

Findings strongly support upper echelons theory (Hambrick & Mason, 1984), which posits that organizational outcomes reflect top management team characteristics. Hambrick and Mason argued that executives' cognitive bases, values, and perceptions—shaped by backgrounds and experiences—systematically influence strategic choices.

This study extends upper echelons theory to board contexts and digital governance specifically. Directors with technical backgrounds perceived AI opportunities more readily, drove adoption discussions more actively, and expressed more confidence managing implementation risks. Directors with traditional governance backgrounds focused more on controls, expressed more skepticism, and slowed adoption until accountability mechanisms became clear. These weren't random individual differences but systematic patterns reflecting how professional backgrounds shape cognition and values.

The study also reveals mechanisms through which upper echelons effects operate in board AI discussions: digital asymmetry and what might be termed epistemic power—influence derived

from knowledge rather than formal authority. When technical expertise concentrates in few directors, those individuals gain disproportionate influence through informational advantage. They can evaluate proposals others cannot, challenge assumptions others miss, frame alternatives others don't see. This creates influence that operates independently of and sometimes more powerfully than formal board positions.

However, findings complicate simple upper echelons predictions. The theory might suggest boards with more technical directors adopt AI faster and more extensively. But directors emphasized that diverse boards—mixing technical and traditional expertise—produce better AI governance than homogeneous boards. Pure technical boards might move fast but miss governance risks; pure traditional boards might maintain controls but miss opportunities. Optimal board composition appears to involve productive tension between different backgrounds rather than dominance of any single profile.

This suggests upper echelons theory needs refinement for board contexts. While team research often treats cognitive homogeneity as reducing conflict and improving coordination (Hambrick, 2007), board scholarship suggests cognitive heterogeneity proves valuable for oversight function. Boards exist partly to provide independent perspective management may lack; too much cognitive similarity between boards and management potentially undermines this independence. The challenge involves managing heterogeneity so diverse perspectives enrich rather than fragment deliberation—explaining why directors emphasized communication mechanisms like AI advisory councils and plain-language briefings.

5.3.2 Socio-Technical Embeddedness and Context Dependency

Findings align with socio-technical systems theory (Trist & Bamforth, 1951), which emphasizes that technology effects depend fundamentally on social organization and environmental context. Early socio-technical research examined how coal mining technology required corresponding changes in work organization; contemporary scholarship applies similar logic to information systems.

Directors described experiencing this socio-technical interdependence acutely. AI capabilities that might prove valuable in stable Western contexts face fundamentally different constraints in Ukrainian wartime environments. Incomplete data from occupied territories, war-disrupted patterns invalidating historical models, legacy infrastructure compounding technical

challenges, sophisticated adversaries targeting systems—these aren't merely implementation obstacles but fundamental constraints reshaping what AI can reliably deliver.

This context-dependency extended beyond Ukraine specifically to encompass sector and regulatory environment. Banking directors emphasized regulatory frameworks demanding explainability; healthcare directors emphasized patient safety and liability concerns; public sector directors emphasized corruption risks and transparency imperatives. The same AI forecasting capability gets assessed very differently across these contexts based on sector-specific institutional logics.

Orlikowski (2000) argues that technology should be understood not as fixed object with inherent properties but as enacted through practice within specific social and material arrangements. Directors' experiences illustrate this enactment perspective—AI's governance value emerges through interaction with data quality, regulatory frameworks, security conditions, and organizational capabilities rather than flowing deterministically from technical features. AI proves "good" or "bad" for governance not in abstract but through specific instantiation within particular contexts.

The socio-technical perspective also illuminates why Ukrainian context proves theoretically interesting. Most AI governance research examines stable Western settings where data infrastructure is mature, institutional frameworks well-developed, and physical security assumed. Ukrainian wartime conditions remove these background assumptions, revealing dependencies that stable contexts obscure. When data quality degrades fundamentally, when legal frameworks lack clarity, when adversaries actively target systems—AI's governance value looks very different. Ukraine functions as natural experiment or stress test revealing socio-technical dynamics that matter everywhere but show up most clearly under extreme conditions.

This finding extends socio-technical systems theory by demonstrating how extreme contexts illuminate dependencies normally invisible. Just as Trist and Bamforth's (1951) coal mining studies revealed socio-technical interdependencies through examining unusual work conditions, Ukrainian directors' experiences reveal AI governance dependencies through examining unusual environmental conditions. The insight generalizes beyond Ukraine: any board operating under stress—organizational crisis, regulatory investigation, market

disruption—faces similar constraints making AI simultaneously more urgent and more problematic.

5.3.3 Institutional Pressures and Legitimacy Seeking

Institutional theory (DiMaggio & Powell, 1983) offers additional insight into contextual effects, particularly regulatory intensity's role. Institutional theory argues that organizations adopt structures and practices not purely for technical efficiency but also for social legitimacy. Organizations face institutional pressures—coercive (regulatory requirements), mimetic (imitating peers), and normative (professional standards)—shaping their choices.

Directors in heavily regulated sectors described experiencing strong coercive institutional pressures around AI. Banking supervisors demand explainability; healthcare regulators require clinical validation; energy regulators expect infrastructure protection. These aren't purely technical requirements but institutional expectations defining what counts as legitimate AI use. Directors adopt explainability mechanisms and human-in-the-loop principles partly because these practices satisfy institutional demands, not just because they improve decision quality.

The study also reveals mimetic institutional pressures—directors look to peers and international standards when local frameworks lack clarity. Directors representing organizations with international investors or parent companies described applying "international corporate governance standards to maximum extent" even when Ukrainian law doesn't require it. This reflects mimetic isomorphism—adopting practices legitimated elsewhere to signal sophistication and reduce uncertainty.

Normative pressures appear less developed currently. Directors didn't reference emerging professional standards around board AI governance (perhaps because few exist yet), but their convergence on similar practices—human-in-the-loop, formal frameworks, transparency—may represent early normative consensus forming within governance community. As professional associations develop AI governance guidance and directors gain shared training, normative pressures may strengthen.

Institutional theory thus helps explain governance practices directors recommend. These practices aren't purely rational responses to technical challenges but also symbolic responses to institutional expectations. By adopting visible governance mechanisms—AI risk

committees, independent validation, explicit documentation—boards signal legitimacy to regulators, investors, and other stakeholders. This doesn't mean practices are merely ceremonial; rather, it recognizes that governance serves both functional (improving oversight) and institutional (demonstrating appropriateness) purposes simultaneously.

The Ukrainian context particularly highlights institutional complexity. Directors navigate multiple, sometimes conflicting institutional pressures: Ukrainian regulators with incomplete AI frameworks, EU regulators with evolving standards that Ukrainian organizations must meet for market access, international investors expecting global best practices, and local stakeholders with corruption concerns. This institutional pluralism creates both confusion and opportunity—confusion about which standards apply, but opportunity to selectively adopt practices from different institutional environments.

Board composition and organizational context systematically shape AI perspectives through multiple theoretical mechanisms. Upper echelons theory explains how professional backgrounds shape cognition, with technical expertise enabling informed adoption but creating epistemic power asymmetries requiring deliberate management through board diversity and communication structures. Socio-technical systems theory explains how AI's governance value depends on contextual factors—data quality, institutional frameworks, security conditions—with Ukrainian wartime environment serving as stress test revealing dependencies stable contexts obscure. Institutional theory explains how regulatory intensity and legitimacy concerns shape adoption patterns, with coercive, mimetic, and emerging normative pressures all influencing governance practices directors recommend. Together, these theoretical perspectives explain substantial variation observed across directors and organizations in how they balance AI's promise against its perils.

5.4 Governance Innovation and Adaptation (RQ3)

Directors' governance recommendations address the third research question: what implications and practices do directors identify as important for responsible adoption? Analysis reveals sophisticated adaptive strategy preserving core principles through practice evolution.

5.4.1 Enduring Principles, Evolving Practices

The core principles directors insisted on preserving are familiar from corporate governance scholarship: boards must maintain accountability (clear responsibility for decisions), transparency (explainable reasoning), and independent judgment (critical evaluation rather than rubber-stamping). These principles trace to fundamental agency and stewardship concerns—boards exist to protect stakeholder interests through oversight requiring these capabilities (Fama & Jensen, 1983).

What's novel is how directors propose instantiating these principles given AI's specific challenges. Traditional accountability assumes decision-makers can articulate reasoning in human-understandable terms; AI accountability requires additional mechanisms—documenting how algorithmic inputs were used or overridden, maintaining parallel traditional analyses for validation, establishing clear ownership of model governance. Traditional transparency assumes straightforward logic; AI transparency requires translation mechanisms—plain-language briefings for directors, independent expert review, explicit disclosure of model limitations and failure modes.

Rai (2020) describes this shift as moving from outcome accountability to process accountability—making visible not just what was decided but how algorithmic tools influenced decisions. Directors' emphasis on documenting "how AI insights were used, not just the final decision" reflects this process orientation. Similarly, their insistence on maintaining parallel traditional analyses alongside AI outputs creates what might be termed "redundant transparency"—multiple paths to understanding that reduce dependence on any single (potentially opaque) analytical method.

This pattern—enduring principles, evolving practices—demonstrates what organizational scholars term dynamic capabilities (Teece et al., 1997). Organizations succeed not through static fit with environment but through capacity to adapt practices while maintaining core identity. Directors demonstrate similar dynamic capability at governance level: adapting how boards work to accommodate AI while preserving what boards fundamentally do (oversight, accountability, strategic guidance).

The Ukrainian wartime context particularly highlights this adaptive capacity. Directors couldn't simply copy Western AI governance models because contextual differences—data quality problems, security threats, institutional weaknesses—make direct transfer inappropriate. Instead, they adapted principles to constraints: if external AI services pose security risks given

adversarial threats, use on-premise systems; if data quality is questionable due to war disruptions, maintain strong human override authority; if regulatory frameworks lack clarity, voluntarily apply international standards to demonstrate legitimacy. This contextual adaptation while preserving governance fundamentals demonstrates sophisticated rather than mechanical thinking.

5.4.2 Human-in-the-Loop as Foundational but Insufficient

Directors' universal emphasis on human-in-the-loop deserves particular attention. This wasn't presented as temporary measure pending AI improvement but as permanent governance requirement—AI must always support rather than replace human judgment in material board decisions.

This norm reflects multiple theoretical concerns. From agency perspective, human-in-the-loop preserves directors' monitoring function—they remain ultimate decision-makers rather than becoming validators of algorithmic outputs. From legal perspective, it maintains clear accountability—humans, not algorithms, bear responsibility when decisions produce bad outcomes. From cognitive perspective, it prevents deskilling risks directors worried about—maintaining human judgment capacity through continued exercise rather than allowing atrophy through non-use.

However, human-in-the-loop also creates practical tensions that recent scholarship has examined. What constitutes meaningful "human in loop" when algorithms process data humans cannot personally verify, apply models humans don't fully understand, and generate recommendations humans lack expertise to evaluate comprehensively? Lebovitz et al. (2021) document this challenge in healthcare contexts, showing how mandating "physician in the loop" for AI diagnostics doesn't guarantee meaningful oversight if physicians lack capability to critically evaluate algorithmic inputs. They found physicians sometimes defer to AI recommendations they cannot adequately assess, with human-in-the-loop becoming procedural requirement rather than substantive safeguard.

Directors recognized this tension, explaining why they paired human-in-the-loop requirements with director capability-building recommendations. They need training not to become data scientists but to ask intelligent questions—understanding basic AI concepts, recognizing common failure modes (overfitting, data bias, concept drift), probing assumptions and

limitations. Without this capability, human-in-the-loop risks becoming what one director called "symbolic oversight"—procedurally correct but substantively empty.

This finding connects to broader debates about algorithmic accountability in legal and information systems scholarship. Scholars increasingly recognize that simply requiring "human decision" doesn't guarantee meaningful oversight if humans lack capacity to evaluate algorithmic inputs. The combination directors proposed—human-in-the-loop plus director literacy plus explainability requirements plus independent validation—creates layered accountability more robust than any single mechanism alone.

Kellogg et al. (2020) describe this as "accountability infrastructure"—multiple mechanisms working together to maintain responsibility despite algorithmic mediation. Directors' recommendations align with this infrastructural approach: no single safeguard suffices, but combination of human authority, technical capability, transparent processes, and independent review creates resilient accountability even when individual mechanisms prove imperfect.

5.4.3 Governance as Collective Learning Process

Directors' cautious, experimental stance—"interested but careful"—suggests they view AI governance not as problem to solve once but as ongoing learning process. Current governance practices reflect current understanding, which experience will likely evolve. This perspective aligns with organizational learning theory (Argyris & Schön, 1978), which distinguishes single-loop learning (adjusting practices within existing frameworks) from double-loop learning (questioning frameworks themselves).

Directors appear positioned for potential double-loop learning about board governance. Initial AI encounters may produce single-loop adjustments—adding AI-specific policies, establishing review procedures, providing training. But accumulated experience may eventually prompt deeper questions: What is board's role when algorithms generate strategic insights management didn't anticipate? How should board composition evolve if digital literacy becomes prerequisite for effective oversight? Do traditional quarterly meeting rhythms work when AI enables real-time risk monitoring?

The Ukrainian wartime context may accelerate this learning given intensity of pressures directors face. When organizations operate under extreme stress—war, reconstruction,

volatility—governance inadequacies become visible faster than in stable environments. Ukrainian boards may therefore develop governance innovations that prove valuable elsewhere as AI adoption spreads globally. This suggests value in longitudinal research tracking how director understanding and practices evolve through experience.

Directors' governance recommendations reveal sophisticated adaptive strategy rather than binary choice between traditional governance and technological transformation. They preserve core principles—accountability, transparency, independent judgment—while evolving practices to accommodate AI's specific challenges. Human-in-the-loop emerges as foundational norm but proves insufficient without complementary mechanisms: director capability-building, explainability requirements, independent validation, process documentation. Together these create accountability infrastructure resilient to individual mechanism failure. Directors' experimental stance suggests they view governance as ongoing learning process, with current practices representing provisional equilibrium subject to revision through accumulated experience. Ukrainian context, by intensifying pressures and constraints, may accelerate governance innovation with broader applicability.

5.5 Integrating Insights Across Research Questions

The three research questions, while analytically distinct, reveal interconnected dynamics that become visible through cross-question synthesis.

5.5.1 How Understanding Shapes and Is Shaped by Context

Directors' dialectical understanding (RQ1) doesn't emerge in vacuum but reflects board composition and organizational context (RQ2). Technical directors emphasize cognitive benefits more than governance risks; traditional directors emphasize risks more than benefits. This isn't because technical directors ignore risks or traditional directors ignore benefits, but because professional backgrounds provide different lenses making certain aspects more salient.

Moreover, organizational context amplifies or dampens both dimensions of the dialectic. Wartime volatility makes cognitive benefits more urgent (need better forecasting under extreme uncertainty) while simultaneously making governance risks more acute (degraded data, security threats). Regulatory intensity in supervised sectors shifts balance toward risk emphasis—explainability requirements and accountability concerns loom larger when

supervisors scrutinize major decisions. Sector-specific logics further modulate the balance: patient safety concerns in healthcare, systemic stability concerns in banking, corruption concerns in public sector all shape which dimension of AI's dual nature receives more attention.

This suggests directors' sensemaking (RQ1) and contextual influences (RQ2) form recursive loop rather than linear causation. Context shapes what directors attend to and how they interpret AI, but directors' interpretations also shape how they enact context—choosing which institutional pressures to prioritize, how to frame AI within regulatory requirements, whether to emphasize opportunities or constraints when discussing organizational strategy.

5.5.2 How Context and Composition Interact

The relationship between board composition and organizational context proves more complex than simple additive effects. Technical expertise matters more in some contexts than others. In highly volatile, data-rich environments like technology companies or financial services, technical directors' ability to evaluate AI capabilities and interrogate algorithmic assumptions proves especially valuable. In contexts where regulatory scrutiny is intense, traditional directors' control orientation and accountability focus proves especially valuable.

Diverse boards navigate context more effectively than homogeneous boards precisely because different contexts foreground different aspects of AI governance. When discussing AI implementation in customer-facing operations, technical directors contribute expertise about capabilities and limitations. When discussing AI's governance implications with regulators, traditional directors contribute expertise about compliance frameworks and accountability structures. Neither expertise alone suffices; both prove necessary for navigating complex, high-stakes environments.

Ukrainian wartime context particularly highlights this interaction. Extreme conditions make both technical and governance expertise essential simultaneously—technical expertise to assess feasibility given data and infrastructure constraints, governance expertise to maintain accountability despite operational stress. Boards lacking either dimension struggle: those with only technical expertise may deploy AI that proves ungovernable under pressure; those with only governance expertise may miss opportunities to use AI for survival-critical functions like early warning or resource optimization.

5.5.3 How Practices Reflect and Reinforce Understanding

Governance practices directors recommend (RQ3) both reflect their dialectical understanding (RQ1) and attempt to address contextual challenges (RQ2). Human-in-the-loop principles reflect understanding that AI threatens accountability; director training reflects understanding that digital asymmetry creates governance risks; transparency requirements reflect understanding that explainability proves essential for oversight.

But practices also shape understanding through feedback loops. As boards experiment with human-in-the-loop mechanisms, they gain experience refining what "meaningful" human involvement means—discovering, for example, that pro forma review doesn't constitute genuine oversight. As they invest in director training, they develop more sophisticated understanding of AI's capabilities and limitations, enabling more nuanced assessment of proposals. As they implement transparency requirements, they learn which aspects of AI decision processes actually need documentation for accountability versus which create excessive bureaucracy.

This feedback suggests dynamic rather than static relationship among the three research domains. Directors' initial understanding shapes governance experiments; governance experiences refine understanding; refined understanding informs governance evolution. This learning loop explains why directors described current stance as provisional—"interested but careful" pending more experience—rather than settled conviction.

5.5.4 Theoretical Integration: Toward Contingent AI Governance

Synthesizing across research questions suggests need for contingent theory of AI governance—one recognizing that effective approaches depend on interaction among director cognition, board composition, organizational context, and accumulated experience.

Simple universalist prescriptions—"boards should adopt AI" or "boards should reject AI"—miss crucial contingencies. Whether AI enhances or undermines governance depends on: what cognitive challenges boards face (information overload, temporal myopia, fragmentation), what expertise boards possess (technical versus traditional, diverse versus homogeneous), what context boards operate in (stable versus volatile, regulated versus unregulated, adequate versus

inadequate data quality), and what governance mechanisms boards implement (human-in-the-loop, transparency, capability-building).

This contingent perspective aligns with configurational thinking in organization theory (Meyer et al., 1993), which emphasizes that organizational effectiveness depends on fit among multiple elements rather than optimizing any single dimension. Effective AI governance similarly requires fit among understanding (recognizing dialectical nature), composition (balancing expertise), context (adapting to constraints), and practices (preserving principles through evolved mechanisms).

Ukrainian directors' experiences prove valuable precisely because extreme context makes contingencies visible that stable environments obscure. Their adaptive strategies—preserving principles while evolving practices, maintaining human authority while leveraging computational power, balancing enthusiasm with caution—offer considerations for boards navigating similar tensions under less extreme but still challenging conditions.

CHAPTER 6: CONCLUSION

6.1 Summary of Key Insights

The board directors in this study reveal that AI in governance is neither straightforward opportunity nor simple threat, but rather a complex socio-technical challenge requiring sophisticated navigation. Their experiences illuminate three insights about how boards might responsibly engage with algorithmic decision support.

AI demands dialectical thinking. Directors resist binary framings—adoption versus rejection, enthusiasm versus resistance. Instead, they maintain productive tension, recognizing AI's genuine potential to address cognitive constraints boards face while remaining alert to governance risks it introduces. This sophisticated ambivalence—what directors called being "interested but careful"—emerges as strength rather than weakness, enabling boards to explore AI's benefits without surrendering oversight capabilities. The dialectic proves essential: pure enthusiasm risks governance capture by algorithmic authority; pure skepticism risks missing valuable cognitive augmentation.

Context and composition determine outcomes. Whether AI enhances or undermines governance depends fundamentally on who sits around the board table and what pressures the organization faces. Technical directors see opportunities traditional directors miss; traditional directors see risks technical directors underestimate. Wartime conditions make AI simultaneously more urgent and more dangerous. Regulatory intensity shifts calculations toward caution. Sector-specific logics generate radically different assessments of identical capabilities. These contingencies mean no universal prescription suffices—effective AI governance requires adaptation to specific circumstances rather than mechanical application of best practices developed elsewhere.

Governance innovation preserves rather than abandons principles. Directors chart middle path between governance paralysis and reckless transformation. They insist on accountability, transparency, and independent judgment—familiar governance principles—while recognizing these must be instantiated differently when algorithms mediate decisions. Human-in-the-loop, explainability requirements, director capability-building, independent validation, process documentation—these practices create what might be termed accountability infrastructure, layered mechanisms ensuring resilient oversight even when individual safeguards prove

imperfect. This adaptive strategy demonstrates that boards can embrace technological capability without abandoning governance responsibility.

These insights matter beyond Ukraine. As AI adoption accelerates globally, boards may confront similar tensions between cognitive enhancement and governance challenges, similar needs to adapt governance to context, similar imperatives to innovate practices while preserving principles. Ukrainian directors' experiences under extreme conditions illuminate dynamics that will become increasingly visible elsewhere as algorithmic decision support becomes ubiquitous.

6.2 Theoretical Contributions

This research advances theoretical understanding across multiple domains, with implications extending beyond immediate AI governance questions.

The bounded rationality paradox. Simon's (1947,1997) theory posits that decision aids help overcome cognitive constraints. This research reveals more complex reality: AI addresses some cognitive limits (information processing capacity) while creating new ones (comprehension deficits when logic becomes opaque). This suggests theoretical refinement: cognitive extension tools don't simply expand capacity but rather trade one form of boundedness for another. The question becomes not whether AI reduces bounded rationality but whether it trades constraints boards can manage (too much information) for constraints boards cannot (unexplainable recommendations). Future bounded rationality research should examine these trade-offs rather than assuming decision aids straightforwardly improve decisions.

Epistemic power in governance contexts. Scholars have recognized AI's role as an epistemic tool that shapes what counts as authoritative knowledge (Faraj et al., 2018). Upper echelons theory (Hambrick & Mason, 1984) explains how executive backgrounds shape strategic choices but focuses primarily on operational teams. This study documents how professional backgrounds create influence in oversight contexts through epistemic authority—directors who can interpret AI outputs gain disproportionate voice regardless of formal position. The contribution lies not in identifying epistemic power as a novel concept, but in documenting how it operates specifically within board AI deliberations. For boards serving monitoring functions, cognitive diversity proves more valuable than cognitive homogeneity, reversing the

logic that applies to operational teams. This distinction—between teams that execute and teams that oversee—deserves more theoretical attention.

Extreme contexts as theoretical microscopes. Socio-technical systems theory emphasizes context-dependency but typically examines stable organizational settings. This research demonstrates how extreme contexts—war, volatility, institutional uncertainty—function as theoretical microscopes, making visible dependencies that stable environments obscure. Ukrainian directors' experiences with degraded data, security threats, and institutional gaps reveal socio-technical interdependencies crucial everywhere but most apparent under stress. This suggests methodological implication: researchers seeking to understand technology's social embeddedness might deliberately study extreme contexts rather than assuming stable settings provide clearer views. Extremity illuminates rather than distorts.

Algorithmic intermediation and meta-level agency problems. Agency theory addresses information asymmetries between principals and agents. This research reveals second-order complication: when AI systems controlled by management filter information for boards, directors face uncertainty not just about operations but about the filtering process itself. This meta-level information asymmetry—not knowing what's being hidden by algorithmic selection—compounds traditional agency concerns. Future agency research should examine how information technologies reshape rather than simply reduce information asymmetries, recognizing that some technologies solve traditional problems while creating novel ones.

Governance as adaptive practice. Corporate governance scholarship often treats board practices as relatively stable responses to regulatory or institutional pressures. This research reveals boards capable of sophisticated adaptation—preserving core principles while substantially evolving practices when confronting technological disruption. The concept of accountability infrastructure—layered, redundant mechanisms creating resilient oversight—offers framework for understanding governance innovation more broadly. Rather than viewing governance as choosing between traditional practices and wholesale reinvention, this perspective suggests continuous adaptation maintaining continuity through change.

6.3 Practical Implications

These theoretical insights translate into concrete guidance for practice.

For boards: Embrace productive tension. The instinct to resolve AI's dialectical nature—deciding it's primarily opportunity or primarily threat—should be resisted. Maintaining tension between recognition of benefits and awareness of risks proves more valuable than premature resolution. Boards should cultivate what directors called "interested but careful" stance: exploring applications while preserving skepticism, piloting implementations while insisting on safeguards, building capability while maintaining oversight authority. This requires deliberate effort because organizational and institutional pressures often push toward binary positions—enthusiastic adoption to appear innovative or cautious rejection to appear prudent. Productive ambivalence, while cognitively taxing, enables learning through controlled experimentation.

Invest in composition and capability simultaneously. Board diversity around AI requires both structural and developmental interventions. Structurally, boards need mix of technical expertise (to evaluate capabilities) and traditional governance expertise (to maintain accountability). Neither alone suffices; productive tension between perspectives generates better outcomes than homogeneity. But structure alone proves insufficient—digital asymmetries create exclusion risks even on diverse boards. Developmental interventions—training for less technical directors, plain-language communication norms, translation mechanisms—enable diverse boards to function effectively. The goal isn't making all directors technical experts but rather ensuring all directors can participate meaningfully in AI discussions through combination of structural diversity and capability development.

Build accountability infrastructure, not single safeguards. No individual governance mechanism proves sufficient for AI oversight. Human-in-the-loop becomes symbolic without director capability to question outputs. Explainability requirements prove empty without independent expertise to validate explanations. Transparency mechanisms generate paperwork without insight unless someone can interpret technical documentation. The solution involves layered, redundant mechanisms: human authority plus technical capability plus explainability plus independent validation plus process documentation. This infrastructure approach accepts that individual mechanisms will sometimes fail but ensures governance doesn't collapse when they do. Redundancy that might seem inefficient for stable technologies proves essential for novel, uncertain ones.

Adapt to your context, don't copy solutions. Effective AI governance depends profoundly on organizational circumstances. Data quality, regulatory intensity, security threats,

institutional maturity, sector-specific risks—these contingencies shape what works. Boards should resist copying practices from organizations operating under different constraints. Instead, identify core principles requiring preservation (accountability, transparency, independent judgment) and adapt practices to instantiate these principles given specific context. Ukrainian boards couldn't copy Western approaches because wartime constraints—degraded data, security threats, institutional gaps—made direct transfer inappropriate. Most boards face less extreme but still significant contextual variations requiring adaptation rather than imitation.

For regulators: Enable rather than prescribe. Findings suggest regulators should focus on governance outcomes—ensuring accountability, transparency, explainability—rather than prescribing specific technical implementations. Given profound context-dependency, prescriptive rules risk being too rigid for some contexts and too permissive for others. Principles-based regulation allowing boards to adapt implementation to circumstances while maintaining accountability for outcomes proves more robust. Regulators can support effective governance by developing director education resources, facilitating information sharing about governance innovations across organizations, and providing safe harbors for controlled experimentation rather than mandating specific practices.

For technology providers: Design for governance. Findings suggest AI systems should embed governance capabilities rather than treating them as afterthoughts. Explainability mechanisms enabling plain-language description of model logic, interfaces allowing directors to interrogate assumptions and test sensitivities, audit trails documenting how algorithms influenced decisions, human override capabilities preserving final authority—these features make systems governable. Vendors might resist this as constraining innovation or adding complexity, but governable AI proves more sustainable than opaque systems boards eventually reject due to oversight impossibility. Design for governance becomes competitive advantage as boards become more sophisticated about oversight requirements.

6.4 Limitations and Boundary Conditions

The insights this research generates come with important qualifications.

Ukrainian context provides both advantage and limitation. Wartime conditions create stress test revealing dependencies normally invisible, offering theoretical value. But extreme context

also limits direct generalizability—boards operating under stability face different pressures. The challenge involves extracting transferable insights from context-specific findings. The claim isn't that all boards should govern AI exactly as Ukrainian directors recommend, but rather that dynamics visible under Ukrainian stress exist latently elsewhere, becoming apparent as boards encounter their own forms of pressure even if less extreme. Readers should ask not "do these findings apply to my stable context?" but rather "what dependencies might these findings help me anticipate before stress makes them obvious?"

Cross-sectional design captures moment, not movement. Directors' perspectives reflect current understanding shaped by limited experience. As boards accumulate direct encounters with AI—successes, failures, surprises—their views will evolve. What appears as sophisticated caution now might shift toward either deeper skepticism (if early experiments disappoint) or greater confidence (if implementations succeed). Longitudinal tracking would reveal this evolution, but current findings offer only snapshot. This temporal limitation means conclusions should be framed provisionally: this is how directors currently make sense of AI, not necessarily how they will ultimately view it after accumulated experience reshapes understanding.

Self-reports may diverge from actual behavior. Directors described what they believe should happen regarding AI governance, not necessarily what actually happens in practice. Stated intentions face implementation challenges—governance recommendations that sound reasonable in interview may prove impractical when confronting real resource constraints, time pressures, political dynamics. Moreover, social desirability may shape responses—directors perhaps emphasizing governance discipline more in interviews than they prioritize in actual decisions. Observational research examining real board AI discussions would provide crucial validation or complication of findings about governance practices.

Sample selection toward engagement. Voluntary participation likely attracted directors already thinking seriously about AI governance. The sample may thus over-represent sophisticated perspectives and under-represent directors avoiding or dismissing these questions. This selection effect means findings illuminate how engaged directors make sense of AI but may miss barriers preventing engagement in first place. Future research deliberately sampling AI-skeptical or disengaged directors could reveal different dynamics—perhaps simpler framings, more categorical rejection, less governance innovation. The sophisticated

dialectical thinking observed here may be characteristic of directors who choose to grapple with AI questions rather than universal response.

Limited examination of actual outcomes. The research examines perspectives and stated practices but not whether AI actually improves board effectiveness or whether recommended governance mechanisms actually prevent anticipated problems. Directors believe AI could address cognitive constraints and worry it might undermine accountability, but do these predictions prove accurate? Do governance practices they recommend actually work? Outcome research linking AI adoption and governance practices to board performance metrics would provide essential evidence currently absent. Without outcome data, conclusions necessarily remain at level of theory and aspiration rather than validated effectiveness.

6.5 Future Research Directions

The patterns this research reveals open several promising avenues for deeper understanding.

The story doesn't end with stated intentions. Directors described sophisticated governance strategies, but how do these play out through actual implementation? Following boards longitudinally as they move from AI curiosity to experimentation to adoption would reveal whether dialectical tensions persist or resolve, whether governance mechanisms work as intended or require unexpected adjustments, whether experience confirms initial concerns or reveals surprises directors didn't anticipate. Some boards might discover that AI proves easier to govern than feared; others might encounter unanticipated challenges. Tracking this evolution would transform the current snapshot into developmental narrative, revealing how sensemaking and governance practices mature through accumulated experience.

Ukrainian extremity illuminates, but comparison would clarify boundaries. The stress test Ukrainian conditions provide makes dependencies visible, but which patterns reflect universal dynamics versus context-specific responses? Examining boards across radically different settings—peacetime stability, mature institutions, varied regulatory regimes, different sectors—would establish what generalizes and what remains locally contingent. Do directors everywhere exhibit dialectical thinking or does Ukrainian volatility uniquely produce this ambivalence? Do diverse boards always outperform homogeneous ones or does this depend on specific governance challenges? Such comparison would reveal whether insights from extreme

context transfer to calmer waters or whether different conditions generate entirely different governance dynamics.

Words and actions sometimes diverge. Directors articulated thoughtful governance intentions, but observational research examining actual boardroom discussions would reveal whether practice matches principle. How do epistemic power dynamics actually unfold when technical directors dominate discussions? How do boards really negotiate conflicts between innovation pressure and accountability concerns? What governance mechanisms get implemented versus which remain aspirational given resource and time constraints? Ethnographic methods—observing meetings, analyzing documents, tracking decisions—would provide reality check on self-reported perspectives, potentially revealing gaps between what directors believe should happen and what organizational politics, resource constraints, and ingrained habits allow.

Does AI actually deliver on its promise and threats? The research captures directors' predictions about AI's effects, but outcome studies would test whether these predictions prove accurate. Does AI adoption actually improve board decision quality, or do cognitive benefits prove illusory? Do governance mechanisms directors recommend actually prevent accountability erosion, or do problems emerge despite safeguards? Which specific practices correlate with successful adoption versus failed implementation? Linking AI use and governance approaches to board performance metrics—decision quality, strategic foresight, risk oversight—would move discussion from theoretical speculation to empirical evidence about what works under what conditions.

Political dynamics deserve deeper examination. This research identifies epistemic power but only scratches the surface of political processes through which boards negotiate AI adoption. Process studies examining conflicts, coalition formation, persuasive rhetoric would reveal how boards manage disagreements between technical enthusiasts and governance skeptics. When do conflicts prove productive, generating creative solutions? When do they devolve into paralysis or premature capitulation to one side? What communication strategies help boards navigate digital asymmetries constructively? Understanding political micro-dynamics would complement current structural insights about composition effects.

Practices diffuse through fields, not just within organizations. How do governance innovations spread from pioneering boards to broader adoption? Professional associations, consultants, regulators, peer networks all potentially facilitate diffusion, but through what mechanisms and

with what transformations? Do practices directors develop locally get codified into universal standards, or do they remain adapted to specific circumstances? How quickly does experimental diversity give way to normative convergence around "best practices"? Studying governance innovation at field level would illuminate institutional dynamics shaping how board AI oversight evolves collectively rather than just organizationally.

These research directions share common thread: moving from understanding how directors currently think about AI governance toward understanding how governance actually evolves through implementation experience, across varied contexts, under political pressures, with measurable consequences, and through institutional diffusion. Current research provides foundation—sophisticated snapshot of director sensemaking under extreme conditions. Future research can build dynamic, comparative, behavioral, outcome-oriented, and institutional understanding that transforms foundation into comprehensive theory of AI governance in practice.

6.6 Concluding Reflections

This research began with straightforward question: how do board directors make sense of AI's implications for governance? The answer proves more complex and more interesting than simple adoption-or-rejection binary.

Directors in this study demonstrate that effective AI governance may require holding contradictions rather than resolving them prematurely. AI addresses real cognitive constraints boards face while threatening core governance capabilities. Both realities matter; neither should be dismissed. This dialectical stance—maintaining productive tension between recognition of promise and awareness of peril—enables learning through cautious experimentation rather than forcing premature commitment or rejection.

Directors also demonstrate that governance innovation needn't mean governance abandonment. Core principles—accountability, transparency, independent judgment—can be preserved while substantially evolving practices to accommodate technological change. Human-in-the-loop, explainability requirements, capability-building, independent validation—these mechanisms instantiate familiar principles in novel circumstances. The adaptive strategy directors exhibit suggests path between two extremes: governance paralysis treating all change as threatening versus governance capitulation treating technology as inevitably transformative.

Perhaps most importantly, directors demonstrate that context fundamentally shapes what works. Board composition, organizational circumstances, institutional pressures, sector-specific risks—these contingencies mean no universal prescription suffices. Effective governance requires adaptation to specific situations rather than mechanical application of supposed best practices. This contextual sensitivity, visible in how directors in this study navigate extreme constraints, offers a proposition for other boards: understand your circumstances, preserve core principles, adapt practices accordingly.

The theoretical insight unifying these practical observations involves recognizing AI governance as socio-technical challenge requiring simultaneous attention to technological capabilities, human cognition, organizational context, and institutional pressures. Neither pure technical analysis (what can AI do?) nor pure governance theory (what should boards do?) suffices alone. Understanding requires integrating multiple perspectives—bounded rationality, agency theory, upper echelons, socio-technical systems, institutional theory—because AI governance sits at intersection of multiple domains. This theoretical eclecticism, far from representing confusion, reflects the genuine complexity of the phenomenon.

As AI becomes ubiquitous in organizational decision-making, boards cannot avoid engagement. The question becomes not whether to use AI but how to do so responsibly. Directors in this study provide an instructive example: cautiously, experimentally, adaptively, preserving accountability fundamentals while evolving practices to capture cognitive benefits. Their experiences under extreme stress—war, volatility, institutional uncertainty—illuminate challenges that other boards may increasingly face as algorithmic decision support spreads.

One director's phrase captures the essential tension boards must navigate: "People see the value, but nobody wants decisions that look like they were made «because the model said so»." This tension—between leveraging computational power and maintaining human accountability—defines effective board governance in the algorithmic age. These directors show this tension can be navigated successfully, neither embracing AI uncritically nor rejecting it categorically, but rather engaging thoughtfully with both its promise and its peril. Their sophisticated ambivalence, adaptive strategies, and contextual sensitivity offer propositions worth considering for boards confronting similar challenges under their own forms of pressure.

The future of corporate governance involves not choosing between human judgment and algorithmic analysis but rather finding ways to productively combine both—capturing AI's

cognitive advantages while preserving governance capabilities boards need for effective oversight. This integration proves difficult but achievable, requiring deliberate attention to board composition, director capability, governance mechanisms, and contextual adaptation.

The experiences of directors in this study suggest one possible path forward: embrace dialectical thinking, invest in diverse capable boards, build layered accountability infrastructure, adapt to context, treat governance as ongoing learning. These practices, forged under extreme conditions, offer considerations for boards navigating AI's transformation of corporate governance.

REFERENCES

- Argyris, C., & Schön, D. A. (1978). *Organizational learning: A theory of action perspective*. Addison-Wesley.
- Adadi, A., & Berrada, M. (2018). Peeking inside the black box: A survey on explainable artificial intelligence (XAI). *IEEE Access*, 6, 52138-52160. <https://doi.org/10.1109/ACCESS.2018.2870052>
- Baxter, G., & Sommerville, I. (2011). Socio-technical systems: From design methods to systems engineering. *Interacting with Computers*, 23(1), 4-17. <https://doi.org/10.1016/j.intcom.2010.07.003>
- Benbya, H., Pachidi, S., & Jarvenpaa, S. (2020). Special issue editorial: Artificial intelligence in organizations: Implications for information systems research. *Journal of the Association for Information Systems*, 22(2), 281-303. <https://doi.org/10.17705/1jais.00662>
- Brynjolfsson, E., & McAfee, A. (2017). The business of artificial intelligence. *Harvard Business Review*. <https://hbr.org/2017/07/the-business-of-artificial-intelligence>
- Beane, M. (2019). Shadow learning: Building robotic surgical skill when approved means fail. *Administrative Science Quarterly*, 64(1), 87-123. <https://doi.org/10.1177/0001839217751692>
- Bughin, J., Hazan, E., Ramaswamy, S., Chui, M., Allas, T., Dahlström, P., Henke, N., & Trench, M. (2017). *Artificial intelligence: The next digital frontier?* McKinsey Global Institute.
- Čaić, M., Odekerken-Schröder, G., & Mahr, D. (2018). Service robots: Value co-creation and co-destruction in elderly care networks. *Journal of Service Management*, 29(2), 178-205. <https://doi.org/10.1108/JOSM-07-2017-0179>
- Carpenter, M. A., Geletkanycz, M. A., & Sanders, W. G. (2004). Upper echelons research revisited: Antecedents, elements, and consequences of top management team composition. *Journal of Management*, 30(6), 749-778. <https://doi.org/10.1016/j.jm.2004.06.001>

- Cycyota, C. S., & Harrison, D. A. (2006). What (not) to expect when surveying executives: A meta-analysis of top manager response rates and techniques over time. *Organizational Research Methods*, 9(2), 133-160. <https://doi.org/10.1177/1094428105280770>
- Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108-116.
- 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>
- Fama, E. F., & Jensen, M. C. (1983). Separation of ownership and control. *Journal of Law and Economics*, 26(2), 301-325. <https://doi.org/10.1086/467037>
- Faraj, S., Pachidi, S., & Sayegh, K. (2018). Working and organizing in the age of the learning algorithm. *Information and Organization*, 28(1), 62–70. <https://doi.org/10.1016/j.infoandorg.2018.02.005>
- Geertz, C. (1973). *The interpretation of cultures*. Basic Books.
- 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.
- 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>
- Haffke, I., Kalgovas, B., & Benlian, A. (2016). The role of the CIO and the CDO in an organization's digital transformation. *Proceedings of the International Conference on Information Systems (ICIS)*.
- Hambrick, D. C. (2007). Upper echelons theory: An update. *Academy of Management Review*, 32(2), 334-343. <https://doi.org/10.5465/amr.2007.24345254>

Hambrick, D. C., & Mason, P. A. (1984). Upper echelons: The organization as a reflection of its top managers. *Academy of Management Review*, 9(2), 193-206. <https://doi.org/10.5465/amr.1984.4277628>

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)

Jansen, H. (2010). The logic of qualitative survey research and its position in the field of social research methods. *Forum: Qualitative Social Research*, 11(2), Article 11. <https://doi.org/10.17169/fqs-11.2.1450>

Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons*, 61(4), 577-586. <https://doi.org/10.1016/j.bushor.2018.03.007>

Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366-410. <https://doi.org/10.5465/annals.2018.0174>

Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.

Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2021). To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organization Science*, 33(1), 126-148. <https://doi.org/10.1287/orsc.2021.1549>

Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. Sage Publications.

Packin, N. G., & Lev-Aretz, Y. (2018). Learning algorithms and discrimination. In W. Barfield & U. Pagallo (Eds.), *Research handbook on the law of artificial intelligence* (pp. 88-113). Edward Elgar Publishing.

Morgan, D. L. (2007). Paradigms lost and pragmatism regained: Methodological implications of combining qualitative and quantitative methods. *Journal of Mixed Methods Research*, 1(1), 48-76. <https://doi.org/10.1177/2345678906292462>

- Orlikowski, W. J. (1992). The duality of technology: Rethinking the concept of technology in organizations. *Organization Science*, 3(3), 398-427. <https://doi.org/10.1287/orsc.3.3.398>
- Patton, M. Q. (2002). *Qualitative research & evaluation methods* (3rd ed.). Sage Publications.
- Rai, A. (2020). Explainable AI: From black box to glass box. *Journal of the Academy of Marketing Science*, 48(1), 137-141. <https://doi.org/10.1007/s11747-019-00710-5>
- Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation-augmentation paradox. *Academy of Management Review*, 46(1), 192-210. <https://doi.org/10.5465/amr.2018.0072>
- Shrestha, Y. R., Ben-Menahem, S. M., & Von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. *California Management Review*, 61(4), 66-83. <https://doi.org/10.1177/0008125619862257>
- Simon, H. A. (1947). *Administrative behavior: A study of decision-making processes in administrative organizations*. Macmillan.
- Simon, H. A. (1997). *Administrative behavior: A study of decision-making processes in administrative organizations* (4th ed.). Free Press.
- 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<509::AID-SMJ882>3.0.CO;2-Z](https://doi.org/10.1002/(SICI)1097-0266(199708)18:7<509::AID-SMJ882>3.0.CO;2-Z)
- Trist, E. L., & Bamforth, K. W. (1951). Some social and psychological consequences of the longwall method of coal-getting. *Human Relations*, 4(1), 3-38. <https://doi.org/10.1177/001872675100400101>
- Trist, E. (1981). *The evolution of socio-technical systems: A conceptual framework and an action research program* (Occasional Paper No. 2). Ontario Quality of Working Life Centre.
- Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). Sage Publications.

Deloitte. (2024). AI in the boardroom: Progress on adoption but room to accelerate. Deloitte Global. <https://www.deloitte.com/global/en/issues/trust/progress-on-ai-in-the-boardroom-but-room-to-accelerate.html>

Institute of Directors. (2024). AI governance in the boardroom. IoD Policy Report. <https://www.iod.com/app/uploads/2025/06/IoD-AI-Governance-in-the-Boardroom-1-7f341cd46a780f216ad68d88d4cddc8f.pdf>

PwC. (2024). Board oversight of AI. PwC Governance Insights Center. <https://www.pwc.com/us/en/services/governance-insights-center/library/board-oversight-ai.html>

Weick, K. E. (1995). Sensemaking in organizations. Sage Publications.

APPENDIX A: DATA COLLECTION INSTRUMENTS

This appendix presents the core 9-question instrument used for both semi-structured interviews and written surveys. The identical question structure enables direct comparison across methods while respecting the strengths of each format.

For interviews: Questions are presented conversationally with opportunities for probing, follow-up questions, and natural dialogue. The interview begins with an informed consent script and concludes with thanks and offer to share findings.

For surveys: The same open-ended questions are presented in written format via Google Forms with simple response-length guidance (for example 3–10 sentences depending on question complexity). The survey begins with an informed consent page and concludes with optional questions about future outlook and whether participants want to receive findings.

CORE 9-QUESTION INSTRUMENT

SECTION A: BACKGROUND (2 questions)

Question 1: Board Experience

Please briefly describe your board experience, including industries you've served in, approximate years on boards, and types of organizations (public companies, private firms, nonprofits, etc.).

For interviews - Probe: What roles have you held? What kinds of strategic decisions have you been involved in?

For surveys - Guidance: 3-5 sentences. Include information about your roles and types of strategic decisions you've been involved in.

Question 2: AI Familiarity

How would you characterize your familiarity with AI and data analytics technologies? Please describe any relevant professional experience, education, or personal interest.

For interviews - Probe: Have you worked with AI in other contexts? Do you have technical training? What's your comfort level with technology generally?

For surveys - Guidance: 2-4 sentences. Include your comfort level with technology and any specific exposure to AI.

SECTION B: CURRENT ENGAGEMENT & SENSEMAKING (3 questions) → RQ1

Question 3: Board AI Discussions

Has your board discussed using AI for strategic decision-making or governance purposes? Please describe any experiences or conversations you have had, or, if there have not been any, how you understand this situation.

For interviews - Probe: Can you walk me through a specific instance? Who raised it? What was the context? How did the discussion unfold? What was decided or what happened next?

For surveys - Guidance: 5-10 sentences. Please be as specific as possible about actual discussions (who raised it, what was the context, how the discussion unfolded, what was decided) or why they haven't occurred and whether you expect them to.

Question 4: Potential Benefits

From your perspective as a director, what do you see as the most significant potential benefits of using AI for board-level decision-making?

For interviews - Probe: Can you rank these benefits? Which matters most and why? Are there any benefits that might be especially relevant to your industry or organization?

For surveys - Guidance: 5-8 sentences. Please explain your reasoning for each benefit you identify and, if possible, indicate which benefits you consider most important and why.

Question 5: Concerns and Challenges

What concerns or challenges, if any, do you see regarding AI use in board-level decision-making or governance?

For interviews – Probe: Which concerns worry you most? Can you give specific examples? How do you personally think about any trade-offs you see when using AI in board work?

For surveys - Guidance: 5-8 sentences. Please be specific about each concern and why it matters. Feel free to indicate which concerns you find most significant.

SECTION C: BOARD CHARACTERISTICS & CONTEXT (2 questions) → RQ2

Question 6: Board Composition Effects

In your experience, how does the composition of the board influence how directors approach AI? Please describe any situations that illustrate this.

For interviews - Probe: Thinking about a specific AI discussion: Who spoke up? Who was quiet? Whose arguments carried weight? Why? Did technical expertise give some directors more influence? Did it create power dynamics?

For surveys - Guidance: 6-10 sentences. Concrete examples are especially valuable. If you've observed specific AI discussions, please describe who spoke up, who was quiet, whose arguments carried weight, and whether technical expertise created particular dynamics.

Question 7: Organizational Context

Are there aspects of your organization or broader context that you think shape your board's perspective on AI adoption? Please describe them in your own words.

For interviews - Probe: How does your specific context—your industry, your company's situation—affect how you think about AI? Would a board in a different sector or country think about this differently?

For surveys – Guidance: 5–8 sentences. Please explain how these contextual factors influence thinking about AI in your board.

SECTION D: GOVERNANCE IMPLICATIONS (2 questions)

Question 8: Governance Implications

How do you think AI might change board dynamics, deliberation processes, or accountability structures? Please answer in your own words.

For interviews – Probe: Can you give concrete examples of changes you anticipate or have observed? Have you seen situations where people rely more—or less—on particular sources of information when technology is involved? How would you maintain genuine oversight if AI tools became part of board work?

For surveys – Guidance: 7–10 sentences. Please describe any changes you anticipate or have observed and explain your reasoning. If possible, mention both positive and negative implications.

Question 9: Governance Practices and Safeguards

What governance practices or safeguards do you believe are necessary for boards to use AI responsibly and effectively?

For interviews - Probe: Which safeguards are most critical? Who should be responsible for implementing them? How would these work in practice?

For surveys - Guidance: 6-10 sentences. Please explain why each practice you recommend is important and how it would work in practice. Which safeguards do you consider most critical? Who should be responsible for implementing them?

APPENDIX B: DATA STRUCTURE TABLES

This appendix presents the data structure showing the systematic progression from first-order concepts (directors' own language) through second-order themes (interpretive groupings) to aggregate theoretical dimensions, following Gioia methodology (Gioia et al., 2013).

The complete analysis involved approximately 390 first-order concepts identified across 22 respondents, organized initially by survey questions (Q1-Q9) and subsequently through cross-cutting thematic analysis. Tables B.1, B.2, and B.3 below display representative examples for each research question, illustrating the analytical logic through which directors' authentic language was systematically organized into interpretive themes and connected to theoretical constructs.

Appendix shows representative examples; full structure available upon request.

TABLE B.1: Data Structure for RQ1 – Current Engagement and Sensemaking

Representative First-Order Concepts	Second-Order Themes	Aggregate Dimensions
COGNITIVE EXTENSION PATTERNS:		
• "Cutting through information overload" (R01, R04)	Theme 1.1: Information Filtering & Anti-Overload Mechanism	DIMENSION 1: AI as Cognitive Extension (Bounded Rationality)
• "Summaries of board packs" (R01, R02)		
• "Highlighting what's important" (R01, R03)		
• "Making discussions fact-based" (R03)		
• "Reducing manual work" (R02, R07)		
• "Scenario analysis & what-if testing" (R04, R14)	Theme 1.2: Scenario Generation & Early Warning System	
• "Early warning signals" (R04, R07, R14)		
• "Forward-looking vs backward-looking" (R04, R05)		
• "Weak signal detection" (R04)		
• "Seeing around corners" (multiple respondents)		
• "Integration of fragmented data" (R07, R09)	Theme 1.3: Systems Integration Tool	
• "Connecting operational silos" (R09)		
• "Synthesizing macro & micro data" (R07)		
• "Connecting the dots" (multiple respondents)		
GOVERNANCE RISK PATTERNS:		
• "Can't explain how model works" (R01, R05)	Theme 2.1: Explainability Gap & "The Model Said So" Problem	
• "The model said so' won't work" (R01, R05, R07)		
• "Black-box distrust" (R05, R18)		
• "Regulators won't accept algorithms" (R05, R07)		
• "Fiduciary duty requires explanation" (R05)		
• "Unclear who is accountable" (R01, R05)	Theme 2.2:	
• "Who is responsible if AI fails?" (R05, R10)		

• "Directors vs management vs vendor?" (R05)	Accountability Fog & Liability Uncertainty	DIMENSION 2: AI as Black-Box Governance Challenge	
• "Legal frameworks unclear" (R05, R10)			
• "Accountability fog" (R18)			
• "Over-reliance on tools not understood" (R02, R05)	Theme 2.3: Over-Reliance & Deskilling Risks		
• "Deskilling of directors" (R05)			
• "Magic number on screen" (R13)			
• "Automation bias" (R10)			
• "Losing critical judgment" (R05, R07)			
DATA & SECURITY PATTERNS:			DIMENSION 3: Data & Security Vulnerabilities
• "Garbage in, garbage out" (R03, R05)	Theme 3.1: Data Fragility & Model Drift		
• "War-distorted statistics" (R18)			
• "Pre-war models no longer valid" (R20)			
• "Incomplete regional data" (R12, R17)			
• "Legacy system inconsistencies" (R05)			
• "AI systems as attack targets" (R06)	Theme 3.2: Cybersecurity & Adversarial Threats		
• "Cloud/third-party risks" (R06, R16)			
• "Data breach consequences" (R10)			
• "System integrity concerns" (R06)			
• "Intelligence collection by hostile actors" (R06)			

TABLE B.2: Data Structure for RQ2 – Board Composition and Organizational Context

Representative First-Order Concepts	Second-Order Themes	Aggregate Dimensions
COMPOSITION PATTERNS:		
• "Tech directors translate to business language" (R01, R04)	Theme 4.1: Tech-Savvy Directors as Translators & Power Centers	DIMENSION 4: Board Composition & Digital Asymmetry (Upper Echelons Theory)
• "Digital backgrounds improve debates" (R01, R04)		
• "Can ask the right questions" (R04, R05)		
• "Tech people drive AI agenda" (R04, R13)		
• "Understanding model architectures" (R05, R19)		
• "Traditional directors bring skepticism" (R02, R05)	Theme 4.2: Traditional Profiles as Skeptics & Control Function	
• "Bankers focus on controls" (R02, R07)		
• "Asking about audit trails" (R02)		
• "What if model is wrong?" (R06)		
• "Regulatory compliance questions" (R05, R07)		
• "Less tech-confident directors tune out" (R03, R11)	Theme 4.3: Digital Asymmetry & Exclusion Dynamics	
• "Eyes glaze over during presentations" (R11)		
• "Fear of asking basic questions" (multiple)		
• "Power centralization around systems" (R06)		
• "Directors who can't follow stay silent" (R03)		
• "Mix of tech, finance, operations gives best debates" (R04)	Theme 4.4: Diverse Boards as Prerequisite for Balance	
• "AI Advisory Council bridging gaps" (R19)		
• "Balance of enthusiasm and caution" (R01, R04)		
• "Translation mechanisms needed" (multiple)		
• "Plain-language briefings essential" (R01, R05)		
CONTEXT PATTERNS:		
• "War drives hunger for forecasting" (R01, R03)	Theme 5.1: War & Volatility as Paradoxical Pressure	DIMENSION 5: Contextual Pressures (Socio-Technical Systems)
• "Volatility makes AI attractive" (R01, R04)		
• "War-distorted data undermines models" (R03, R18)		
• "Resilience prioritized over innovation" (R06)		
• "Every decision through 'survival' lens" (R06)		
• "Supervisors demand explainability" (R01, R05)	Theme 5.2: Regulatory Intensity Demanding Controls	
• "Regulators won't accept algorithms" (R05, R07)		
• "Tight oversight in banking/energy/health" (R05, R10, R14)		
• "National Bank strict after crisis" (R05)		
• "Must justify every material decision" (R05, R07)		
• "Banking = model risk + regulation" (R02, R05, R07)	Theme 5.3: Sector-Specific Risk Logics	
• "Healthcare = ethics + patient safety" (R10, R11)		
• "Public sector = corruption concerns" (R21, R22)		
• "Tech = competitive necessity" (R04, R13, R16)		
• "FMCG = margin optimization critical" (R03, R17)		
• "Ukraine-specific volatility" (R01, R18)		

• "Reconstruction imperatives" (R01, R21)	Theme 5.4: Ukrainian/CEE Context as Unique Laboratory	
• "EU integration pressure" (R18, R20)		
• "Post-Soviet governance legacy" (R21)		
• "Institutional gaps and uncertainty" (R18, R20)		

TABLE B.3: Data Structure for RQ3 – Governance Implications and Practices

Representative First-Order Concepts	Second-Order Themes	Aggregate Dimensions
ANTICIPATED IMPLICATIONS:		DIMENSION 6: AI's Governance Transformations & Safeguards
• "Data-fluent directors gain influence" (R03, R04)	Theme 6.1: Power Shifts Toward Data-Fluent Members	
• "Directors who interpret AI outputs lead" (R04)		
• "Power centralization around systems" (R06, R14)		
• "Chairs must balance voices" (R04)		
• "Dashboard-driven discussions" (R01, R03)	Theme 6.2: Deliberation Process Changes	
• "Shift from 'what happened' to 'which scenario'" (R19)		
• "More live, less static meetings" (R13)		
• "Risk of hiding behind 'model says'" (R01, R02)		
• "Softer issues get less airtime" (R01, R12)		
• "Minutes must show how AI used" (R01, R04)	Theme 6.3: Enhanced Documentation Requirements	
• "Document how insights challenged" (R01)		
• "Explicit records of AI influence" (R05)		
• "Paper trail for accountability" (multiple)		
RECOMMENDED PRACTICES:		
• "AI as support, not decision-maker" (R01, R03, R04)	Theme 6.4: Human Authority & "Tool Not Authority" Norm	
• "Tool, not authority" (R04)		
• "Human-in-the-loop principle" (R01, R04, R05)		
• "People remain accountable" (R02, R05)		
• "Mandatory human review" (R05, R07)		
• "Clear AI policy needed" (R01, R04, R05)	Theme 6.5: Formal Governance Frameworks	
• "Model risk framework" (R05, R07)		
• "AI Risk Committee" (R10, R11)		
• "Independent validation required" (R01, R05, R07)		
• "Clear ownership of AI oversight" (R01, R07)		
• "Minutes reflect AI use" (R01, R13)	Theme 6.6: Transparency & Decision Trails	
• "Dual confirmation for major decisions" (R04)		
• "AI as evidence, not decision" (R05)		
• "How AI informed debate, where overridden" (R13)		
• "Parallel traditional analysis maintained" (R05)		
• "On-premise for sensitive data" (R06)	Theme 6.7: Data Governance & Infrastructure Control	
• "No external LLMs with confidential info" (R16)		
• "Data provenance tracking" (R10)		
• "Cybersecurity audits mandatory" (R06, R19)		
• "Complete control over infrastructure" (R06)		
• "Mandatory training for directors" (R01, R05)	Theme 6.8: Director AI Literacy & Capacity-Building	
• "Plain-language briefings" (R01, R05)		
• "Avoid symbolic oversight" (R05)		
• "Enable intelligent questioning" (R01, R05)		
• "Not data scientists, but can ask right questions" (multiple)		