INTERNATIONALE HOCHSCHULE

Executive Governance of Generative AI in Knowledge-Intensive Organizations: A Systematic Literature Review and Conceptual Framework for Value Creation and Risk Mitigation

Capstone Thesis for the M.B.A. Major in IT-Management

Carina Sophie Schoppe

IU14132098

Submitted on 13 March 2026

Supervisor: Prof. Dr. Dirk Battenfeld

Acknowledgements

I would like to start with expressing my sincere gratitude to the people who have supported and accompanied me throughout my academic and personal journey.

My deepest thanks go to my parents, Prof. Dr. Marion Klammer Schoppe and Karl Rudolf Schoppe.
My mother, as a role model, guided me toward the world of academia, while my father, Karl Rudolf Schoppe, introduced me to entrepreneurship and consistently encouraged me to pursue my ambitions with confidence.
Together, they shaped my intellectual development and made it possible for me to pursue my dream of living and studying in Australia.

To further extend my gratitude, I want to thank my heartfelt appreciation to Annalena Holz, whose steady support and calm presence gave me the strength to remain focused and determined in achieving my goals.

My sincere thanks go to my supervisor and examiner, Prof. Dr. Dirk Battenfeld, for his continuous guidance and support throughout the development of this capstone thesis.

And finally, to my loyal best friend Percy — wherever you may be now, may you always have something to chew on.

Abstract

Generative artificial intelligence (GenAI) is rapidly reshaping knowledge work and value creation in knowledge-intensive organizations, yet executive oversight frequently lags adoption. This capstone thesis addresses that governance gap by synthesizing what the academic literature currently says about (1) how GenAI creates organizational value and (2) which strategic, operational, regulatory, and ethical risks executives must manage to achieve sustainable benefits. The guiding research question is: *How can executives design governance structures for generative AI in knowledge-intensive organizations to maximize organizational value while minimizing strategic and operational risks?*

Methodologically, the study applies a PRISMA-oriented systematic literature review with mixed evidence like peer-reviewed and high authority grey sources. The analysis uses iterative thematic coding to consolidate evidence on value mechanisms, risk typologies, and governance controls, with explicit attention to the persistent “principles-to-practice” challenge in responsible AI guidance.

The findings indicate that GenAI can deliver significant value through productivity gains, faster onboarding, accelerated innovation cycles, and improved business–IT alignment. At the same time, value realization is constrained by a distinct risk landscape, including hallucinations and quality failures, bias and fairness harms, privacy and data-governance exposure, third-party toolchain opacity, security vulnerabilities, regulatory uncertainty, and organizational side effects such as reduced creativity and governance fatigue. Across the literature, a core gap is the limited availability of operationally actionable, executive-ready governance designs tailored to knowledge-intensive contexts.

Building on the synthesis, the capstone thesis proposes an integrated executive governance framework that aligns GenAI use cases with strategy and risk appetite, institutionalizes policy and accountability structures and embeds human-in-the-loop controls for quality-sensitive work. A complementary KPI and monitoring architecture is developed to enable continuous executive oversight across value, risk, and compliance dimensions.

Limitations relate to the fast-evolving evidence base and the literature-only design; future work should validate the framework empirically through studies.

Keywords: generative AI, executive governance, knowledge-intensive organizations, risk management, responsible AI, policy architecture, AI inventory, auditability, human-in-the-loop, KPI monitoring.

Table of Contents

1. Introduction

Generative Artificial Intelligence (AI) is transforming organizations and the ways in which they compete within the digital economy. These technologies have changed traditional value creation models in knowledge-intensive organizations, providing enhanced efficiencies, innovative methods of customer engagement, and improved decision-making. However, the adoption of generative AI brings various strategic and operational risks, including ethical, regulatory, and operational concerns. This necessitates a holistic governance framework that enables knowledge-intensive organizations to balance the benefits of generative AI with the inherent strategic and operational risks of its deployment. This thesis addresses the following research question: *How can executives design governance structures for generative AI in knowledge-intensive organizations to maximize organizational value while minimizing strategic and operational risks?*

The capstone thesis, "Executive Governance of Generative AI in Knowledge-Intensive Organizations," explores the intersection of technology management, corporate governance, and risk management. Generative AI is used for tasks such as predictive analytics, natural language processing, and autonomous systems. In a knowledge-intensive context, these technologies provide an opportunity to transform the core of value creation for such firms. However, these technologies also present several operational and strategic risks. Furthermore, AI and governance frameworks must align with ethical concerns and regulatory standards. The balance between the strategic value of generative AI for knowledge-intensive organizations and the challenges of its governance highlights the importance of this topic and of this thesis.

This thesis aims to evaluate and synthesize the literature on both the value creation mechanisms and the strategic and operational risks generated by generative AI within organizations. From here, it seeks to construct an executive governance framework that will allow executives to guide the deployment of generative AI technologies effectively and responsibly. The identification of the gaps and limitations in the literature provides the foundation for this conceptual framework, which contributes to the conversation surrounding generative AI governance.

To investigate the research question, the method consists of a systematic literature review of studies, industry reports, and regulatory documents to provide a thorough overview of the topic of interest. This approach provides critical perspectives in multiple domains and enables identification of patterns, common threads, and trends related to generative AI governance in organizations. Evaluation, comparison, and synthesis methods are used throughout, highlighting the most important findings related to the research question. Special attention will be paid to the interrelations of corporate governance principles, responsible AI frameworks, and the specific needs of knowledge-intensive organizations. The established theories, practices, and regulatory standards of corporate governance and responsible AI will influence the development of the governance framework.

Current literature displays significant progress and limitations surrounding the governance of generative AI. Findings from Attard-Frost and Lyons (2024) and Batool, Zowghi, and Bano (2025) highlight the importance of having a governance framework to help companies manage both the benefits and risks tied to the adoption of generative AI. Other literature supports this notion by pointing to the transformative impact that generative AI has in various contexts, while also stating the need to have risk and opportunity management approaches to guide the use of the technology. However, the existing research displays clear limitations in integrating governance principles within generative AI adoption strategies. These works lack the comprehensiveness required to provide executives with practical frameworks for adopting generative AI. Additionally, the fragmented nature of the literature highlights the importance of a comprehensive synthesis on the topic.

This thesis follows a straightforward sequence. The introduction provides context, motivation, objectives, and significance for the capstone thesis.

The theoretical foundation in Chapter 2 summarizes concepts and definitions of generative AI, corporate governance, responsible AI, and the strategic and operational benefits and risks of adopting generative AI for organizations. Chapter 3, *Research Methodology*, describes the methodology, including the literature review process, criteria for literature selection, and the analysis techniques used. Chapter 4, *Findings*, details the key findings about value creation mechanisms and strategic and operational risks that come with the use of generative AI (GenAI) in organizations, as well as gaps in current governance practices. Chapter 5, *Executive Governance Framework of Generative AI*, develops a proposed governance framework for executives that covers value management, risk mitigation, and monitoring and performance practices. The final chapter, *Conclusion*, presents the summary of findings and the managerial and theoretical implications of this thesis. The limitations of the research as well as potential topics for further study are also presented in the *Conclusion* chapter.

1.1 Background and Relevance

Generative AI is a key driver of transformation in knowledge-intensive organizations. This advantage is due to its capabilities to automate time-consuming tasks, optimize curation of large amounts of data, and generate novel insights. Insurance companies use generative AI to speed up claims processing and underwriting, which makes the business more efficient and profitable (Eling et al., 2021). Generative AI also allows for the development of sophisticated predictive models and decision-support tools. This results in timely and more effective decisions in all areas of business (Eling et al., 2021). Thus, productivity improves, allowing companies to personalize and improve their services for all stakeholders. However, the scale of these benefits may vary in different contexts. Generative AI allows for reductions in operating expenses and the streamlining of existing operations. The manual intervention has been reduced in all stages of data handling and includes improved accuracy in document drafting and content creation. These capabilities affect the entire value chain and directly allow for improved customer engagement. Organizations utilizing generative

AI can provide proposals in a much shorter timeframe and with far fewer mistakes (Chandra & Rahman, 2026). Generative AI also opens opportunities for organizations to achieve a competitive advantage by providing adaptive and individualized customer interactions. Customers can co-create functional, emotional, and social benefits through their interactions with generative AI, creating strong and personalized customer relationships and providing real business value in the digital marketplace (Chandra & Rahman, 2026). But as with all AI, context and implementation play a key role in how these benefits can be realized in different organizations.

As generative AI enables automation and augmentation, organizations must redesign their workflows, decision hierarchies, and competencies to ensure appropriate structure for effective generative AI adoption. Organizations must govern these structural changes within the evolving ecosystem's parameters to fully reap the benefits of integrating generative AI. To manage this disruption and achieve the intended goals of implementing generative AI, the organization must establish a novel governance mechanism. These improvements will allow for an overall optimization of the business and its stakeholders. Therefore, the organizational structure must be reconfigured to accommodate and utilize the changes that generative AI brings, as this is one of the most important factors that will enable full generative AI utilization.

On the downside, generative AI exposes organizations to strategic, operational, legal, and ethical risks that can impede business goals. While some studies identify statistically significant associations between AI adoption and macroeconomic indicators such as GDP, evidence also suggests that many data science initiatives do not reach sustained production. The literature frequently discusses governance, risk management, and organizational capability constraints as key barriers to moving from experimentation to deployment (Madanchian & Taherdoost, 2025). Challenges related to the operations of generative AI, such as susceptibility to biases, model hallucinations, and high cybersecurity risks, may lead to the failure to produce high-quality outputs, which can negatively impact an organization's reputation and profitability. This issue is further exacerbated in industries with a high reliance on AI-generated outputs, where incorrect results could potentially cost a company a client or millions of dollars. So, businesses should make sure they know about the AI-related risks in their ecosystem and deal with them before the AI is put into use.

The ethical dilemmas surrounding generative AI use are also growing. Research indicates that most ethical dilemmas can be mitigated through the proficient application of deliberate safeguards, such as fairness and ethics oversight, data protection measures, and anticipatory model training. Some practitioner-oriented sources suggest that implementing safeguards (e.g., fairness monitoring, data security strategies, and proactive model training) may substantially improve ethical outcomes; however, the reported magnitude varies and is not consistently grounded in standardized measurement approaches (Sinclair & Mehta, 2023). Unfortunately, this implementation is often overlooked. Such neglect is often because these methods of implementation have proven to be challenging, with a lack of easily utilized frameworks and a trade-off between ethical consideration and functionality that hinders overall uptake. Additionally, varying laws across national and regional

borders and a lack of consistency in these regulations add to the complexity of governing AI use. Multinational organizations must take special consideration of the differing laws concerning AI explainability, accountability, data privacy, etc., as it can be costly, time-consuming, and operationally difficult to tailor regulatory strategies. These multinational companies must create a flexible governance system that allows them to navigate different local, national, and regional policies while continuing to perform global operations and activities.

Moreover, organizations are struggling to locate qualified people to run and manage AI initiatives, hindering the realization of the technology's full benefits (Soroori Sarabi, 2025). AI development requires a host of highly skilled people to carry out the complex tasks of data science. Talent gaps are present across almost every domain necessary for the development of a functioning and effective generative AI platform (Soroori Sarabi, 2025; Madanchian & Taherdoost, 2025). Generative AI also necessitates people who can oversee, maintain, audit, and use it to produce value. The development of an organizational AI talent pool requires deliberate investments in education and development to equip the workforce with the skills needed to successfully navigate generative AI.

The governance and regulation of generative AI also vary globally. As such, there is no universal governance model that can be applied across organizations worldwide. For example, within the U.S., it is primarily industry and sectoral organizations that control the governance and regulation of AI and algorithms. It's unclear how these rules were made or if they're as strict. In contrast, the EU AI Act emphasizes risk categorization and transparency. These approaches highlight a dichotomy in governing AI but also a tension between regulation and innovation (Baidya et al., 2025). Meta-analyses of regulations governing AI, such as those by Madanchian and Taherdoost (2025) reveal that the two categories most divergent in regulation are those of ethics and regulation procedures. Both areas, despite their differences, face heavy regulation (Madanchian & Taherdoost, 2025). This poses an issue for multinational companies, as well as researchers and consumers, as the guidelines may vary widely depending on the specific domain and region being considered. Additionally, the various regulations and procedures may also provide room for regulatory arbitrage, with organizations potentially moving their operations to regions with less strict guidelines to ensure they can maintain a competitive advantage (Madanchian & Taherdoost, 2025). Finally, internationally, entities such as the OECD, the UN, and several international bodies propose regulations for AI, but these are usually general recommendations and not always actionable (Madanchian & Taherdoost, 2025). Thus, organizations operating across international borders are tasked with understanding the implications of the regulatory environment for their operations, strategies, and business value and creating governance models that best fit their unique circumstances.

AI governance frameworks are influenced by decision-makers who utilize varying processes, structures, and information resources to develop policies, and as such, different political choices may be made that would impact the development of governance measures for these technologies. As AI use is not typically defined in any current framework, the implementation can and will vary depending

on the individual setting. According to Perry and Uuk (2019), this political discretion can impact an AI initiative either positively or negatively, as each of these factors influences which interventions are considered, developed, and adopted. Additionally, for AI regulation to scale, the meta-governance challenges such as stakeholder collaboration, consensus-building, feedback mechanisms, and institutional capacity, need to be addressed (Perry & Uuk, 2019). For example, the governance of autonomous vehicles (AV) can differ across state lines, highlighting that there is no unified regulation standard across regions. Although each state had to create legislation regarding the adoption and use of AVs, their governance measures have varying degrees of effectiveness (Poon Affat, 2025). According to research, to improve this AI governance framework, there must be higher-quality feedback loops and policies to enhance institutional and legislative responsiveness, allowing for more inclusive, integrative policymaking processes (Poon Affat, 2025).

Finally, poor-quality data prevents AI from providing value. It also fails to bring value if the company is unable to put the AI into practice. AI, especially generative AI, needs heaps of specialized and fitting data to train, as well as an organizational system to manage and clean it (Soroori Sarabi, 2025). Data quality concerns can seriously impede the success and value-adding potential of AI programs. Soroori Sarabi (2025) finds data quality to be the biggest failure that impedes the advancement of AI implementation. The same applies to organizations struggling to use the AI to its full potential and not adequately prepared for this changing work environment. For example, in organizations that conduct disaster management, most are unable to train models or fully optimize for operational efficiency because they do not have sufficient human capital available to leverage advanced algorithms (Soroori Sarabi, 2025). For that reason, many lack adequate data resources to implement any substantial use of AI and only use it in limited functions (Soroori Sarabi, 2025). Organizations also lack sufficient human capital to maintain and optimize advanced algorithms, as they require highly skilled, trained professionals (Soroori Sarabi, 2025). Additionally, organizations require resources and talent to oversee the use and implementation of AI; without individuals who understand and effectively utilize AI, these organizations will struggle with capacity for integration and transformation (Soroori Sarabi, 2025).

These cases clearly demonstrate the need for a strong AI organizational framework to have generative AI work effectively. In conclusion, generative AI offers many benefits to organizations in this technology-driven environment. However, the use of generative AI in particular presents significant risks and challenges that require strategic and intentional governance. As these are newly implemented technologies, organizations must ensure their ethical implications are addressed, the regulations they fall under are well understood, and their internal structures and operations are adapted for these technologies to ensure that full adoption and implementation are achieved for competitive advantages (Poon Affat, 2025). This governance must be proactive to avoid potential risks and ensure all benefits that AI has to offer are achieved and used appropriately.

1.2 Problem Statement and Research Gap

The implementation of generative AI in knowledge-intensive organizations has presented various challenges and opportunities, one of which is the requirement for integrated executive governance of generative AI to effectively manage the delicate balance between value generation and risk avoidance. Existing literature identifies a clear gap, as it provides little guidance to executives in such organizations on how to derive business value from the technology while effectively managing the risks involved (Madanchian & Taherdoost, 2025; Perry & Uuk, 2019). Much of existing literature either concentrates on a specific aspect, such as the mitigation of technological risks involving data privacy, algorithm bias, or compliance, or on ways of realizing business value, such as generating operational efficiencies. Despite various existing governance frameworks and models, there is no single framework tailored for executives within knowledge-intensive organizations to navigate this balancing act. This leaves them with fragmented governance strategies, which in turn limits their ability to achieve an effective organizational balance.

Moreover, current governance models found in the literature primarily focus on technical and operational aspects, such as overall model fairness, data protection, and regulatory compliance.

As such, they only serve technologists and not executives. These existing governance models are inadequate for addressing the executive-level issues of organizational cross-functional alignment, governance oversight, and governance infrastructure (Batool et al., 2025). Empirical studies of generative AI operationalization show that only a few organizations were successful in launching their generative AI projects to achieve business impact and organizational-level benefits. Most of these studies cite the failure to achieve success because of inadequate governance practices (Smith, 2025). Furthermore, the literature review showed a lack of frameworks combining risk management capabilities and mechanisms for value creation, adaptive learning, and assurance (Smith, 2025).

A further issue in the literature is the limited treatment of how the simultaneous consideration of commercial, ethical, and regulatory outcomes must be considered for organizations adopting generative AI. As organizations benefit from incremental improvements across the board from the implementation of generative AI, they can run into unprecedented risks if appropriate governance structures aren't in place to mitigate such instances. Examples of unintended consequences include model hallucinations, ethical or public controversies, regulatory scrutiny, and compliance breaches (Gehrmann et al., 2025; Waisam & Silver, 2025). Consequently, the necessity for the evaluation and regulation of both favourable and unfavourable outcomes must be examined (Gehrmann et al., 2025; Waisam & Silver, 2025). The existing literature fails to elucidate how organizations ought to formulate operational strategies for the effective implementation of ethical AI principles, including transparency, privacy, fairness, safety, and accountability. Such an absence also means a lack of actionable guidance to operationalize ethical and responsible AI principles (Madanchian & Taherdoost, 2025; Perry & Uuk, 2019). Existing frameworks provide minimal guidance on embedding ethical AI

principles and often do not include metrics to monitor these principles (Madanchian & Taherdoost, 2025; Perry & Uuk, 2019).

Existing AI standards, such as the ISO 42001 standard, are ineffective for influencing regulatory or executive governance tools and mechanisms. For example, the insurance sector has adopted the NAIC Model Bulletin. However, despite numerous theoretical advancements in the AI regulation space, these standards remain abstract and are not readily translated into concrete action points for executives to adopt (Poon Affat, 2025). This leaves executives uncertain about the application of these standards in complex organizational settings and how they can be effectively used for operational decision-making (Poon Affat, 2025). Frameworks in existing literature do not address several important concerns for executives, such as risk appetites, integrating risk monitoring into existing workflows, and impact assessment (Poon Affat, 2025). There are significant gaps in the literature, as the tolerance of AI risk influences the making and execution of governance decisions.

Furthermore, standards such as the ISO 42001 have been unevenly adopted across geographies and industries. In organizations where incorrect risk assessments can lead to significant financial, operational, legal, or reputational loss, an inconsistent and ineffective operationalization of standard guidelines for executives leaves executives at the mercy of the potential unforeseen consequences of inadequate risk management (Poon Affat, 2025).

Moreover, there is a lack of harmony and convergence of AI governance practices across geographies. Despite the efforts to create ethical governance principles in general areas such as fairness, transparency, accountability, and interpretability, there remain significant differences in emphasis between different jurisdictions (Madanchian & Taherdoost, 2025). These differences include the prioritization and interpretation of these ethical governance standards. For instance, the European Union usually puts more value on AI's ability to be open, audited, explained, and classified by risk than on its ability to be innovative and work in practice. Meanwhile, other regions, such as the United States, tend to prefer sector-specific approaches to AI regulations, which prioritize innovation at the risk of operational uncertainty around regulations and compliance (Madanchian & Taherdoost, 2025). These regulatory differences result in compliance risks in multinational organizations and force executives to confront a variety of often conflicting regulatory requirements. This discrepancy will ultimately lead to uncertainty for companies with multinational operations (Madanchian & Taherdoost, 2025). Further compounding the complexity is the proliferation of AI governance initiatives in various sectors that lead to contradictory approaches. This problem is seen in the insurance sector, as the NAIC Model Bulletin and the various initiatives undertaken in Europe may be inconsistent (Poon Affat, 2025). Similarly, the literature recognizes the difficulty in harmonizing governance approaches on an international scale because of political challenges (Perry & Uuk, 2019). The lack of coherence across organizations requires adaptive governance systems to accommodate different needs and levels of regulatory maturity to deal with geopolitical risks. While not a technological risk or limitation, these issues create uncertainty among executives who may be unsure about how to balance innovation and risk aversion or how to be compliant while delivering

customer value (Madanchian & Taherdoost, 2025; Santos et al., 2025).

Another gap in the existing literature is a lack of treatment of several necessary conditions that are often overlooked but ultimately determine the success of generative AI governance. These include aspects such as data governance, technology infrastructure, and human capital readiness for generative AI. Since using generative AI relies on having the right data and technology, along with skilled workers, leaders need to focus on these areas (Lee et al., 2025; Santos et al., 2025). Most of the research literature only mentions data requirements at a high level (Lee et al., 2025; Santos et al., 2025). In addition, there are no readily available assessment metrics to assess human capital readiness for the utilization of generative AI (Madanchian & Taherdoost, 2025). For example, the lack of interoperability and data lineage remains a common issue (Lee et al., 2025). Data integration and lineage can also lead to legal and reputational risks. The existing literature contains little mention of the requirements for data provenance, protection, reliability, and interoperability (Lee et al., 2025). Furthermore, several research articles in healthcare and disaster relief demonstrate the importance of appropriate data and technology resources to effectively and successfully deploy generative AI for operations (Lee et al., 2025; Santos et al., 2025). Research illustrates how generative AI is incapable of operating at a level of quality without appropriate data and human resources (Lee et al., 2025; Santos et al., 2025). Moreover, none of the reviewed literature focuses on aligning strategic investment for data and infrastructure readiness and human capital talent pipelines in the operationalization of generative AI. This lack of literature calls for a deeper assessment of foundational elements.

Finally, another gap exists regarding the importance of political and administrative processes. Although most of the literature in the AI field does not directly address governance systems, the political and administrative mechanisms through which AI policy and practice are set significantly shape the extent to which governance is realized (Perry & Uuk, 2019). In addition, the processes of policymaking affect the effectiveness of generative AI implementation. Some of these factors are decision authority, resource investment, prioritization of policies, and other processes that shape policy decisions (Perry & Uuk, 2019). The fragmentation of legislative and regulatory systems also shapes policy choices (Perry & Uuk, 2019). Furthermore, AI policy involves a myriad of actors, including legislatures, regulators, courts, law enforcement, and political parties. These groups all vary in priorities for policy and standards. Most frameworks do not acknowledge the political processes and the administrative apparatus in which they function. However, such recognition is essential to ensure appropriate enforcement strategies and to ensure accountability (Perry & Uuk, 2019). Additionally, no literature acknowledges or analyses meta-governance, a broader process in which power is negotiated for setting rules in the field. Meta-governance considers a process for which policies and laws emerge. While often obscured in plain sight, this concept is essential to address who is empowered to decide which issues are discussed, who participates, and how the rules of the game are established (Perry & Uuk, 2019). These meta-governance aspects play an important role in the scalability and durability of governance mechanisms.

In conclusion, the existing body of literature regarding AI governance illustrates multiple notable gaps. These include a fragmented treatment of governance, gaps in data governance and alignment, uneven convergence of standards across the globe, lack of consideration for foundational elements of technology readiness, and a lack of focus on the political processes within which governance is executed.

1.3 Research Objective and Structure

The central research objective of this capstone thesis is the design of an integrated executive governance framework for generative AI in knowledge-intensive organizations to balance maximizing the operational and strategic value of generative AI with mitigating associated risks. The achievement of this objective is required, as the literature reveals that most organizations fail to transition their generative AI systems beyond experimental stages. Smith (2025), citing survey evidence, reports that only about 8% of organizations have integrated AI into core operations, while most initiatives remain in an experimental phase. The literature frequently discusses governance readiness, operating-model maturity, and capability constraints as barriers to scaling beyond pilots (Smith, 2025). The lack of an integrated risk management and value creation framework largely causes such failures. The research objective aims to address this gap in the literature. The study intends to provide advice for executive leaders regarding the governance of generative AI systems by combining interdisciplinary findings from literature with generative AI systems being leveraged across knowledge-intensive organizations. This advice must also incorporate the priorities of organizations (Smith, 2025). Furthermore, it must acknowledge that most organizations are not yet ready for generative AI systems and that skills gaps are a real concern for organizational decision-makers (Bolden et al., 2024).

This research will also aim to address prominent governance gaps for generative AI at the executive level. Many organizations still lack consistent documentation regarding the origin, sourcing, and usage of training data, and the implementation of provenance controls is often incomplete or applied inconsistently—especially in early-stage deployments (Maryala, 2025).

Moreover, there is a lack of attention toward risks after the implementation of AI, which are vital for long-term project sustainability (Strauss et al., 2025). Post-deployment risks can include risks in production or unintended consequences, e.g., model bias, operational disruption, data corruption, and algorithmic underperformance (Batool et al., 2025). In addition, research indicates that the main reasons that stop organizations from scaling generative AI system development are the lack of risk frameworks (Waisam & Silver, 2025), the lack of talent with the skills required, and the lack of processes to address those barriers that prevent generative AI system development from being scaled from pilot project to organization-wide implementation (Waisam & Silver, 2025). Therefore, this research aims to identify and discuss the barriers that might hinder effective executive oversight of generative AI implementation and how the literature suggests they might be mitigated.

Evidence from the real world shows that only a small number of companies are currently getting measurable value from AI projects. Most generative AI use cases are still limited to early-stage applications rather than large-scale, value-realizing deployments (Smith, 2025).

Prior work on enterprise AI adoption indicates that scaling pilots into production remains difficult for many organizations, with only a small minority consistently moving beyond proof-of-concepts (Smith, 2025). The failure to scale is generally due to a lack of expertise in generative AI systems implementation and inadequate KPIs, as well as insufficient post-deployment monitoring (Waisam & Silver, 2025). Bolden et al. (2024) reports that many companies have yet to demonstrate tangible value from AI and that progressing beyond proof-of-concepts remains challenging for a large share of organizations. This persistent execution gap increases the need to strengthen governance and operating-model capabilities to support scaling. It should also specify what KPIs or other metrics organizations should monitor when they deploy AI systems into production to realize the most benefit from this technology. Furthermore, this study will offer a mechanism for moving from a pilot generative AI project into production. The study also aims to analyse the effects of meta-governance and the interactions between varying regulations on the effectiveness of executive governance frameworks for generative AI.

The literature suggests that the fragmented nature of policymaking significantly diminishes the effectiveness of generative AI governance (Perry & Uuk, 2019). The European Union (EU), for example, may have entirely different regulations surrounding generative AI than the United States of America (USA) (Baidya et al., 2025), forcing organizations to comply with multiple local and global regulations regarding the design, development, testing, use, deployment, and governance of generative AI systems. This complexity increases operational inefficiencies and opens organizations up to compliance risks that may hinder generative AI strategy execution (Baidya et al., 2025). Therefore, it is vital to understand how meta-governance influences policymaking processes and thus affects the design and implementation of governance frameworks.

This research will emphasize the adoption and operationalization of responsible AI principles into governance structures. Guidelines on ethics, accountability and transparency through the use and deployment of AI may seem appropriate, yet these general statements do not give organizations much advice on how to integrate the concept of “ethics” into their strategy and execution of artificial intelligence at an executive governance level (Madanchian & Taherdoost, 2025). While the topic of how to translate general governance and AI governance into specific governance and controls is beyond the scope of this study, the executive framework that is proposed seeks to embed the concept of ethics into organizational structures and AI adoption, deployment, and governance. It will aim to provide organizations with clear advice on how to prioritize these factors in their strategy, decision-making, and governance in the adoption and use of generative AI. For example, AI ethics task forces involving key cross-functional representatives can operationalize ethical considerations across an organization to increase transparency and accountability (Waisam & Silver, 2025). This project is supported by the existence of governance and monitoring tools, such as Google’s What-If

Tool (Google, 2018), which allows model-specific testing and feedback, thereby enhancing both transparency and accountability (Waisam & Silver, 2025).

In addition, by incorporating the principles of explainable AI, executive governors may reduce the reputational and legal risks for their organizations and maximize value creation by complying with societal and regulatory expectations (Waisam & Silver, 2025). This research intends to offer pragmatic solutions to these challenges so that the organizational leaders can implement these principles effectively.

Organizations must accept the need for iterative improvement in governance structures. By acknowledging the highly volatile, uncertain, and complex nature of technology, executive governors must adopt governance structures that embrace flexibility and adapt over time. Empirical studies indicate that organizations that operate in highly complex or high-risk domains (e.g., the healthcare industry) update their governance frameworks periodically in response to technological and organizational advancements (Freeman et al., 2025). This research will thus suggest a mechanism for continuous adaptation and improvement of the governance framework.

The capstone thesis aims to contribute a proposed executive governance framework for generative AI, along with the description of the key components that will be used within a model. Through empirical validation (e.g., the testing of governance frameworks for scalability) and a clear articulation of the rationale behind this framework, organizations can apply the proposed methodology, leading to the effective management of both the upside and downside risks involved with generative AI. The framework can be applied to any knowledge-intensive organization (both for profit and non-profit).

2. Theoretical Foundations

This section examines the theoretical foundations of generative AI in the organization. The section will develop a stronger understanding of governance practices by exploring the core theories, which will support the overall goal of creating an executive framework.

2.1 Generative AI in Organizational Contexts

Generative AI emerges as a highly impactful technology for increasing productivity and enabling innovation in knowledge-intensive organizations. Its applications deliver productivity improvements and operational efficiencies across a wide range of business functions. Research suggests that after completing pilot projects and full rollout, organizations see real-world results, from streamlined administration and reduced repetitive tasks to faster service times (Schneider et al., 2024; Smith, 2025). For knowledge-intensive organizations like consultancies, finance firms, or law firms, the reduction in human error and the consistent output of complex tasks enabled by generative AI are a significant benefit. However, researchers should further discuss the challenge of consistently achieving these outcomes in any organization. While organizations have found that generative AI

can lead to substantial performance gains reported in selected contexts (Schneider et al., 2024), how effectively those productivity improvements scale beyond specific domains is unclear and depends on factors such as data infrastructure maturity and skill levels of workers (Smith, 2025).

Generative AI is reshaping how organizations create value for their customers, primarily by influencing the mechanical, cognitive, and affective dimensions of value creation and offering functional, emotional, and social value. Functional value involves the mechanical interactions between the customer and the organization, which is greatly impacted by the efficiency with which generative AI-enabled chatbots and automated content can respond to customers (Chandra & Rahman, 2026). Value is also created cognitively, as generative AI systems can improve customer experiences by providing customized experiences (Chandra & Rahman, 2026). Affective aspects relate to emotions and attitudes exhibited toward the customer during service delivery, especially social value. Anthropomorphic designs and the theory of mind in AI improve how trust and connection are built between customer and provider. Generative AI can also automate tasks that foster empathy and care (Chandra & Rahman, 2026). The result is better interactions across all phases of the customer's journey from pre-purchase to post-purchase. The theories of value creation are thus advanced by generative AI, which enables organizations to interact with the customer earlier, more frequently, and more authentically (Chandra & Rahman, 2026). What must also be asked is: would such implementations of generative AI-based value creation be effective in organizations with less technical infrastructure or expertise? Additionally, while customers regard AI personalization valuable, it should not lead to any biases in outcomes.

Increased availability of compute and larger datasets has fuelled significant advances in generative AI recently but also introduced challenges associated with reliability and authenticity. The computational power needed to train models has increased between fourfold and fivefold annually, with dataset size tripling (Kolt et al., 2025). Such growth has allowed more organizations to access and implement generative AI, thereby improving innovation capacity. AI systems now enable organizations to generate ever-increasing quantities of data, which they then utilize to train more AI models. The downside of this "feedback loop" is that as AI systems ingest more AI-generated content, it will become more difficult to authenticate them (Kolt et al., 2025). Additionally, more AI-based systems can generate content that is nearly indistinguishable from human-created content, creating both challenges and risks for organizations. Content generation also opens an organization up to potential regulatory issues, where synthetic content is misattributed to the organization, as well as reputational and financial risks (Kolt et al., 2025). Organizations are now being driven by these risks to develop solutions that effectively detect synthetic content and that ensure generated content is validated as factual by an internal audit system (Kolt et al., 2025).

Despite its tremendous benefits, the full adoption of generative AI is slow, and only about 8% of organizations have fully embraced its use (Smith, 2025). One major hurdle is the difficulty of scaling the benefits of AI. Implementing generative AI often compromises data security due to its

vulnerability to cyberattacks or unintentional breaches (Smith, 2025). The legal challenges of using new technology, such as the absence of new rules or the inconsistency of existing rules across jurisdictions, add to the complexity of this issue.

While the legal ambiguity of compliance with various laws may be unclear, organizations risk financial loss if non-compliance is revealed later (Smith, 2025). These complex and non-standard requirements are also hindering the widespread use of generative AI. Workforce challenges are one of the most critical impediments to realizing scalable benefits. A major skills gap exists because too few individuals can manage generative AI systems, hindering organizations aiming to transition operations and decision-making to this technology. Organizations that have achieved the most benefit from generative AI have learned that instituting an executive-level management group to lead the transition to AI and establishing an organizational cross-functional governance committee for effective adoption are key steps (Smith, 2025). In the implementation of generative AI in multinational organizations, significant risks arise regarding diversity, ethical, and regulatory considerations. As such, it is important to explore these aspects and how they pertain to AI governance. The regulatory approaches toward AI vary significantly between major geographies, resulting in major challenges for companies with an international presence. The European Union, for example, implements accountability-oriented compliance and transparency in AI governance, while the U.S. focuses on the innovative growth of AI within a compliance-driven framework defined by the sectors for which the AI system will be implemented.

Meanwhile, the Chinese strategy prioritizes rapid technological growth of AI, but with a still-developing set of normative principles to govern it (Hogenhout, 2021). These diverse approaches present large compliance hurdles for international organizations that must develop a governance structure for AI that remains adaptable to each country in real time as regulations are revised (Hogenhout, 2021). Further, for the most crucial ethical dimensions, fairness and privacy, AI practices and interpretations vary widely. Because of the global nature of multinationals, governance models must be responsive to these local variances. AI can introduce operational risks that lead to inequitable and biased treatment if no effective governance controls are in place. Generative AI models display performance variability across languages (Schneider et al., 2024). Specifically, models that have performed very well on standardized tests in English have done poorly in other languages in the tests, which raises some concern about fair delivery for multilingual customers. Indeed, building generative AI systems for low-resource languages is more vulnerable to jailbreaking than building them for widely spoken languages (Schneider et al., 2024). For example, even if an AI tool in English resists jailbreaking tactics by refusing to deliver prohibited content, a tool built in a less widely spoken language it is designed for will likely generate the harmful response.

Bias must also be considered and is often caused by non-representative datasets used during training, especially in high-stakes sectors like the medical and financial industries, which has the potential to cause grave disparities (Schneider et al., 2024). There has been a growing awareness of AI risks from an executive level, as evidenced by some organizations instituting a blanket ban on

generative AI (Schneider et al., 2024). However, there is still not much evidence that organizations have in place proactive mechanisms to detect such risks and mitigate any negative effects on fair and equitable treatment for users. A risk mitigation strategy for generative AI should include auditing the model used to generate the AI, the use of automated bias detection tools, a human-in-the-loop audit review to ensure model quality, and an executive-level steering committee for ongoing and effective AI compliance.

In conclusion, there are several issues that organizations must consider if they plan to effectively implement generative AI. The complexity of the issues demonstrates a crucial need for proper oversight and risk control. For AI to be effectively managed, there must be sufficient governance controls to mitigate the risks associated with its implementation.

2.2 Corporate Governance and Risk Management

Corporate governance frameworks for generative AI in knowledge-intensive organizations must cater to novel types of risks, beyond just conventional IT or operational risks. These include algorithmic bias, loss of transparency, and the evolving nature of compliance. Algorithmic bias has emerged as a serious threat, leading to unfair and discriminatory outcomes in insurance underwriting, employee evaluations, and credit scoring. This undermines an organization's reputation and compliance standing. Mitigation of such risks requires bias detection mechanisms within governance frameworks, ensuring the outputs are equitable and adhere to ethical standards (Batool et al., 2025). The lack of transparency, known as the “black box” effect, is also a critical issue. It compromises an organization’s ability to understand and justify the logic behind decisions driven by AI systems, especially deep learning algorithms. This threatens compliance with regulatory requirements that mandate explainability and auditability (Madanchian & Taherdoost, 2025). Accordingly, organizations must invest in mechanisms that can provide greater clarity to AI models. In addition, the evolving nature of compliance requirements across multiple industries, such as insurance and healthcare, mandates that organizations AI governance frameworks must extend beyond conventional IT controls. At the board level, a risk committee must be set up to oversee compliance matters. Moreover, organizations need to establish clear escalation pathways to report and properly address AI-related incidents (Poon Affat, 2025). The implementation of responsible AI principles into broad strategic goals necessitates the development of internal ethics review boards, extensive auditing mechanisms to uncover biases, and clear reporting pathways. These initiatives can ensure that responsible AI is applied consistently across the organization's operational procedures and strategy (Batool et al., 2025). Effective implementation of these processes, including clear escalation when needed, correlates with greater productivity gains from AI when ethics guidelines are put into place (Madanchian & Taherdoost, 2025).

From an operational standpoint, effective risk management must involve clearly structured processes to identify, analyse, and mitigate AI-related risks. Tools such as the Risk Index Number (RIN) and the ISO/IEC 42001 methodology are valuable for setting priorities based on the level of

risk (Kaul et al., 2025).

The RIN method provides an organization's executive team with a way to define priority AI risks through a score calculated based on the likelihood and severity level of high-impact situations.

Similarly, ISO/IEC 42001, an AI Management System (AIMS) framework, improves risk practices beyond simple risk identification and mitigation by incorporating continuous monitoring and reporting tools. Risk governance practices must be reviewed and updated as needed.

Furthermore, there must be adequate AI risk documentation (Poon Affat, 2025). In addition to RIN and ISO/IEC 42001, continuous monitoring tools like MLOps dashboards must be integrated. These tools detect anomalies, biases, and errors within AI outputs, enabling early corrective actions. These insights can then be shared with executives so they can act quickly and effectively, which lowers the risk of damage to the company (Kaul et al., 2025).

Some studies have shown that when an organization follows a framework such as ISO 42001, they have experienced fewer manual documentation requirements. ISO/IEC 42001 can help structure compliance efforts by standardizing controls and documentation, potentially reducing duplication and audit friction (Kaul et al., 2025). Additionally, with various governments throughout the world starting to recognize international standards like ISO/IEC 42001 in the product and service approvals process, these frameworks must be adopted to enable compliance and continue to operate (Poon Affat, 2025).

Though the need for frameworks is supported by research, the issue arises with how they can be practically applied to corporate governance within an organization.

An illustration of this concern in the insurance sector is evident in the uneven application of the NAIC Model Bulletin on AI Use in Insurance by 24 U.S. states. Although the intention of the framework is consistent across states, its interpretation and execution are not.This makes an organization's risk management incomplete if they are working across several states with differing frameworks (Poon Affat, 2025). Further studies have also shown that despite organizations following all compliance requirements from frameworks such as ISO 42001, they are not actually performing the intended risk management on the organization as described in the framework (Poon Affat, 2025). To counter this, companies are looking for specific practical implementation resources and compliance checklists that translate into the application of frameworks. This gap between the development of frameworks and practical implementation also indicates a trend of slower framework implementation by executive leaders, as academic and standards-based risk models are often being implemented more slowly compared to operational measures. There is also an apparent misalignment of priorities in risk model adoption. Industry leaders and businesses appear to focus on operational benefits and gains, whereas regulatory and government organizations focus on fairness. This situation further exacerbates the gap between published standards and executed AI governance in knowledge-intensive organizations (Poon Affat, 2025).

This also means that because industry executives are slow to implement frameworks, the organization can become unprepared for expected compliance requirements and be susceptible to

unexpected risk events. Risks and compliance issues often force many AI projects to remain in the experimental phase, with little to no actual implementation. The lack of appropriate governance models also hinders the development of new operational approaches for generative AI systems, in turn further restricting their applications across organizations. Additionally, some organizational barriers that prevent generative AI models from going to production phases include talent shortages, a lack of investments into risk framework implementations, and the lack of implementable KPIs that indicate production readiness (Smith, 2025). If these hurdles are not addressed, organizations may fail to realize anticipated returns and may even fall victim to severe financial penalties as well as reputational damage from compliance failure. In addition to these dangers, there is also the possibility of losing valuable AI talent and being excluded from markets due to non-compliance (Poon Affat, 2025; Smith, 2025).

Data governance, infrastructure readiness, and human capital form critical building blocks for generative AI. Without an effective system for data management and compliance, artificial intelligence may limit its usefulness due to inaccuracies in data. Corrupted data damages an organization's competitive standing by providing inaccurate insights and eroding consumer trust (Madanchian & Taherdoost, 2025). The organization needs to invest in quality data storage, proper version control, and clear audit trails to not only meet regulatory demands of transparency but also to increase consumer trust by ensuring AI systems can meet their operational demands. Organizational infrastructure and capacity must be established before generative AI can be properly applied and governed (Madanchian & Taherdoost, 2025). Talent shortages are another barrier when governing AI within an organization, as the issue presents itself as a lack of skill and capabilities from organizational staff to properly manage it. The need for education and skills across several different organizational segments requires appropriate talent recruitment and AI literacy training programs for staff, as this gap represents a significant risk factor. These processes are essential when deploying a generative AI system, but most organizations lack the talent to effectively implement it. Smith (2025) citing McKinsey, reports that only about 8% of organizations have fully integrated AI into their core operations, and most projects are still in the testing stage. The literature points to capability constraints - such as talent availability, operating-model maturity, and governance readiness - as common barriers to scaling beyond pilots. Addressing these constraints requires aligning real-time governance, clear allocation of responsibility and resources, and iterative data/process governance practices to improve regulatory readiness and organizational responsiveness (Ghosh et al., 2025).

When governing generative AI systems that engage in creative processes, organizations must go beyond compliance practices and implement several types of management techniques to effectively manage and govern. AI systems utilized for this purpose should be closely and routinely monitored to ensure they are producing outputs aligned with ethical standards for consumers, businesses, and governmental bodies (Chompunuch & Lubart, 2025). In addition to constant observation and evaluations of AI creative outputs, generative AI must be scenario tested under extreme

circumstances or abnormal contexts to gauge the AI models' ability to hallucinate, fabricate false information, or misuse its outputs for other malicious or dangerous purposes (Kaul et al., 2025). It is also critical to implement a human-in-the-loop process to govern creativity and idea generation. The role of human control is critical to managing risks that cannot be controlled by compliance practices alone. If generative AI is incorporated into processes such as idea generation, for example, a human-in-the-loop protocol requires that all AI outputs that go into development must be thoroughly reviewed and approved by a human worker (Chompunuch & Lubart, 2025). Documentation is key to maintaining integrity for governance purposes in knowledge-intensive organizations. It is crucial to identify the organizational decision-maker when evaluating the output of an AI, as it is important to know who made the decision to choose between the available options.

Such identification is made possible through creating operational audit trails for each AI system and the feedback it provides the users (Kaul et al., 2025). One study highlights the significance of establishing an adequate governance architecture to manage creativity by assessing organizations using different methods. Organizations with effective corporate governance (e.g., participatory governance and clear delegation) reported higher trust in outputs and higher-quality results. The same study indicates fewer damaging or disruptive risk events compared with organizations lacking formal governance processes. This reinforces the need for agile governance architectures that can adapt to rapid technological change (Chompunuch & Lubart, 2025).

One element often missing in current governance models is addressing operational vulnerabilities. Many governance structures have focused heavily on fairness and privacy, which is important, but not enough on the risks of systems that can be exploited, such as adversarial attacks, system outages, and model drift (Batool et al., 2025). Another area that needs to be considered more often is addressing potential issues with generative AI before they occur, as many organizations are focusing on mitigating risks after an incident has taken place rather than implementing strategies and resources beforehand. Many organizations lack proper incident response procedures for disruptions, so their procedures may struggle with unexpected operational issues. Additionally, proper crisis communication and response protocols must support organizational compliance programs (Poon Affat, 2025). To properly evaluate and monitor risk incidents, organizations must implement risk registers and scenario plans to anticipate all potential risk incidents. It’s also imperative that participatory oversight be available so that experts who are not tied to the governance organization can be included in these plans to make them more effective (Kaul et al., 2025).

Also largely unaddressed in current governance models is the transparency of AI governance practices, which are important because they significantly contribute to an organization’s trustworthiness, stakeholder engagement, and accountability. For good oversight, all important decisions made by the organization must be able to be traced (Batool et al., 2025).

Inclusive decision-making processes facilitate increased engagement and feedback from organizational personnel, thereby improving transparency (Poon Affat, 2025).

By creating routine auditing processes through reporting systems or regular data checks that can indicate potential bias or error, stakeholders will be more likely to feel they can trust the AI outputs of the organizations, which can impact engagement and utilization (Batool et al., 2025). Regular reports of operational statuses and regulatory compliance levels can then inform stakeholders and regulatory bodies. This procedure provides further transparency and will further increase consumer trust. When an organization is more transparent in their communication of information and processes, they are more likely to be trusted by consumers, employees, and governing organizations (Poon Affat, 2025). By addressing stakeholder concerns through regular reports and providing clear, concise communication channels for inquiries, organizations can improve transparency levels and increase trust among stakeholders (Batool et al., 2025). Finally, organizations need to foster continuous learning by implementing robust data, testing, monitoring, and model training practices. By actively embracing change and responding appropriately through learning processes, AI governance architecture can then adjust to changes to best achieve desired results (Batool et al., 2025).

2.3 Responsible AI and Technology Governance

Responsible AI implementation necessitates the alignment of ethical principles, operational strategies, and technical execution through multi-layered governance models. These models clarify ethical standards at a high level, enable the operationalization of standards, and drive measurable outcomes. Organizations with multi-layered governance have higher levels of maturity in AI implementation and risk mitigation. More businesses are using multi-layered governance models, but there are still inconsistencies between industries and across geographic areas (Rao et al., 2021). While over half of organizations surveyed report formal policies on AI ethics, less than half have operationalized policies.

For the oversight of multi-layered governance models to be most effective, the layers must be clearly connected. The alignment of board-level policy with operational processes requires that policy oversight include monitoring, bias management, and incident response mechanisms. In the absence of this level of integration, the layers are less effective and unable to reconcile the organization's long-term strategy and daily operational activity.

Rao et al. (2021) indicate that many organizations lack formal AI risk planning and explanatory transparency. Bridging the abstract, board-level ethical standards to tactical decisions improves transparency and ensures organizational accountability, driving positive outcomes in the adoption and scalability of AI. These organizations adopt tools such as algorithmic audits and active risk registers to help govern ethical decisions.

To successfully integrate governance frameworks into business operations, it is necessary to provide staff training and establish operational accountability. Organizations need to set up training programs to help people understand how AI makes decisions and how to follow governance rules. The literature suggests that strong governance and cross-functional enablement can improve

human–AI collaboration outcomes and organizational performance (Nwashili, 2025; Rao et al., 2021). Cross-disciplinary, centralized organizations with clear oversight enable the operationalization of responsible AI governance. Failure to adopt this structure and delegate AI governance to the operations level often results in stalled or failed projects as blind spots and fragmentation result in increased complexity, wasted resources, and diminished organizational value, creating a risk of negative reputational exposure and stakeholder trust issues. Maturity in AI implementation and responsible AI adoption can be determined by measuring operational impact. Such measurements can include metrics such as model bias and explainability of results. Organizations with defined governance models demonstrate improved maturity. In fact, organizations reported having more confidence in their ability to address algorithmic bias, with many stating bias mitigation as the primary area of focus within their policies (Rao et al., 2021). Despite increased adoption and organizational maturity, significant inconsistencies in operationalized models exist across different industries and between different-sized organizations. This imbalance contributes to differences in benefits across regions. The society must acknowledge these discrepancies when evaluating the organizational effectiveness and maturity of governance frameworks. Larger organizations with more resources, those that operate in knowledge-based, higher-income industries, and those operating in the most innovative regions, such as Europe, are more likely to adopt and achieve operational maturity through governance frameworks. Global legal mandates for accountability, transparency, and explainability must govern responsible AI implementation. In addition to the variances in interpretation, the variance in adoption and implementation is also impacted by inconsistent policy compliance across different countries. The EU's AI Act establishes a rights-based governance framework that relies on extensive documentation and transparency requirements. The US adopts a more lenient and sector-specific regulatory framework. Organizations operating internationally and multi-nationally must consider a diverse array of governance models for AI implementation to adhere to varying policies, legal requirements, standards, and customer expectations (Baidya et al., 2025). The lack of a single, global governance policy means executives are responsible for anticipatory regulatory monitoring to ensure compliance across the portfolio. Policies such as EU regulations for all AI systems or the US sector-specific, context-aware risk assessment mandates for finance, health, or energy create challenges for organizations to govern consistently on a global scale.

The lack of coherent global regulatory policy across geographies and industries further highlights the need for harmonization. Differences at the operational level result from inconsistencies in organizational interpretations and the uneven adoption of global frameworks. While organizations such as the OECD and the UN strive for standardization, variations are driven at the national and regional levels, preventing a universal model for responsible AI governance (Baidya et al., 2025). As a result, knowledge-intensive organizations require operational-level frameworks that enable them to align to country or industry-level requirements while managing risk and remaining competitive. Inclusivity, adaptability, and accountability within governance policies must be built in at the

operational layer to provide consistency within organizational applications of AI on a global scale. Because global, regional, and sector-level policies continue to develop, an organization's future risk and operational models must include ongoing identification of legal and ethical issues and changes to prevent compliance failures.

Translating ethical theory into practice is still a challenge. The use of frameworks and guidance to develop responsible AI has increased, but only a few organizations are able to scale beyond AI pilots due to a lack of human-centric governance for the operationalization of responsible AI principles (Madanchian & Taherdoost, 2025). By considering the needs and demands of all stakeholders in any AI implementation—from data collection and system design to testing, deployment, and improvement—organizations can reduce the risk of flawed, biased, and unethical AI outputs. Organizational co-creation with affected stakeholders has shown higher organizational performance (Nwashili, 2025; Taherdoost, 2025). By incorporating input from multiple stakeholders who are directly affected, risks such as biased or inappropriate outputs can be mitigated, and the long-term organizational success is improved. However, this approach is adopted by a small number of organizations due to understaffing, resource constraints, and a lack of organizational processes that require accountability from teams implementing AI tools and policies.

The use of generative AI tools increases operational risk for organizations due to the inherent ambiguity and unknown outcomes of the tools (Nwashili, 2025). The user must exercise caution when introducing these tools in sensitive industries like healthcare (Sezgin, 2024). The adoption of large language models in healthcare virtual assistants, for example, can improve personalization and inclusivity through tailored experiences. These improved efficiencies, however, must not come at the expense of bias (Sezgin, 2024). An imbalance of recommendations to any sub-population could pose health and safety risks. Organizations require privacy preservation, understandable decisions, and robust, cross-disciplinary governance to fully scale these tools to operational benefits. Further, organizations should disclose all these AI applications to their stakeholders and provide appropriate channels for feedback. Organizations must conduct regular audits of inputs and outputs, in addition to ongoing risk assessments, to mitigate risk and ensure customers realize benefits. The success of operationalizing and scaling generative AI models depends on leadership prioritizing the implementation of risk governance by integrating a risk management approach into product development and scaling strategies (Nwashili, 2025). It also relies on creating appropriate monitoring tools to ensure the AI projects are generating the proper results as defined by their scope and stakeholders. These metrics include tracking risk-adjusted ROI and adoption rates for all major features. Such tracking is crucial to driving organizational AI benefits, as these models must monitor operations and incorporate human feedback to reduce failures and enhance benefits. To enhance AI scalability, the incorporation of risk, monitoring, and accountability should be integrated into existing organizational risk and governance programs. Further, organizations can adopt human-AI co-creation models with direct stakeholder input for risk monitoring and assessment. These models implement iterative cycles of learning and human feedback as well as formal checkpoints for risk

and strategy reviews with organizational leadership.

The effectiveness of AI governance implementation depends on executives accepting organizational-level responsibility for the monitoring, evaluation, and accountability practices outlined within the AI ethical standards and governance strategies. Organizations can improve the adoption and operationalization of AI governance practices and their effectiveness through multi-layered frameworks, global compliance programs, and human-centric designs. These approaches align operational execution and ethics with long-term organizational strategy, promoting the benefits of generative AI.

2.4 Strategic Management Perspectives on AI

The strategic management of generative AI within the realm of knowledge-intensive organizations requires the adaptation of existing governance models to meet the unique demands and challenges that this technology introduces. Hilb (2020) posits that the traditional “three-lens” view of business, technology, and society must be transformed to address the multifaceted risks associated with generative AI. The evolution of generative AI has resulted in governance situations that challenge existing assumptions about executive control. The increased autonomy of generative AI means systems are making strategic decisions with less deterministic control, posing new challenges to ensuring alignment with business objectives and societal values.

The development of autonomous AI systems that enhance and supplant human decision-making engenders novel governance frameworks focused on the necessity for transparency and the traceability of AI-generated decisions (Hilb, 2020). These systems introduce new dimensions of strategic governance where executive oversight and audit trails of AI-based outcomes are essential for risk management and building trust. If organizations do not address these factors, they run the risk of operationalizing strategic decision-making through “black boxes” that are not governed appropriately. Insufficient oversight may lead to a lack of trust, regulatory penalties, and a failure to successfully manage technology-driven strategic change.

Generative AI must be governed at an ecosystem-level outcome, requiring the “three-lens” approach to have adaptive feedback capabilities to manage alignment with societal and market norms. Hilb (2020) asserts that incorporating adaptive feedback loops into the “three-lens” approach enables monitoring and response to governance outcomes as the ecosystem evolves and new AI-driven market forces and norms take shape. This adaptation helps ensure the governance framework continues to be relevant and aligns with societal expectations as AI systems increasingly influence broader societal trends and beliefs. A lack of adaptive feedback may result in the technology's output failing to align with expectations of value and organizational goals.

The ability to dynamically monitor and change processes as the AI system evolves in the marketplace and ecosystem is key. Knowledge-intensive sectors are fast-moving environments, and AI technologies such as generative AI can rapidly change the landscape of business competition and regulatory pressures. Perry and Uuk (2019) argue that dynamic governance systems are

needed to be responsive to this technology shift and for organizations to prepare for the future risks and uncertainties in generative AI governance. As it pertains to this research, the rapid pace of generative AI advancements necessitates an ongoing approach to the adoption and deployment of new business applications using this technology. Due to the increased agility, resilience, and overall competitive advantage that these AI systems deliver, the technology change is a strategic imperative that needs to be undertaken. Thus, dynamic governance models are needed to address and mitigate potential risks in real time.

Many executives indicate optimism with respect to the potential returns delivered by generative AI, although there exists a large optimism gap with respect to their capabilities in managing the risks in generative AI (NTT DATA, 2025). NTT DATA's (2025) research shows strong executive optimism - 97% of CEOs anticipate a material productivity impact from GenAI and 44% of the C-suite strongly agree that GenAI's ROI outweighs potential security and legal risks - yet governance capability lags: 86% agree bias remains pervasive, while only 43% strongly agree they have systems to track bias and privacy risks. This indicates a material governance gap because this is a serious gap in strategic management, as the strategic imperatives driving executive optimism and the value delivery from generative AI systems will not be realized if an organization is unable to manage the risks that the same technology will create. Generative AI brings new risks of algorithmic bias, which further complicates strategic governance. Only 43% of organizations have systems in place to address bias and privacy risks in algorithmic outputs (NTT DATA, 2025). The fact that the issue is also a source of concern across operations, research, and technology is a reflection on the lack of preparedness in the overall governance structure. The fact that data science projects and business applications have already been rapidly developed using the technology without the capability to address these critical risks is a concern that must be rectified.

The market is projected to grow at a CAGR of 115.9% between 2023 and 2028 (NTT DATA, 2025), so the technology, and the associated risks, will likely permeate more business functions and operational environments in all sectors and knowledge-intensive organizations. As generative AI technology grows in the marketplace, it will amplify existing risks in organizations that have no risk management strategy in place. According to NTT DATA (2025), projections for the *GenAI services* market indicate extremely rapid growth, rising from US$4.7B in 2023 to US$221.8B by 2028, corresponding to a 115.9% CAGR (2023–2028). Proactive and adaptable governance in strategic management that quickly anticipates and reacts to regulation will be integral to capturing value as technology advances. Divergent and disjointed regulatory environments cause fragmentation and misalignment in strategic governance.

Perry and Uuk (2019) reported divergent regulatory frameworks and that contradictory regulatory outcomes hinder the establishment and enforcement of adequate governance for multinational companies. As an example of this, consider two nations with different AI regulatory agencies that have differing philosophies with respect to the governance of AI systems. This situation creates compliance confusion for organizations as they develop and deploy AI systems in these two nations.

The misalignment causes regulatory complications that further delay deployment and may stifle strategic innovation. Thus, alignment requires proactive engagement with policy changes from regulators. Corporate strategy and external regulations impact governance, and this research has previously noted that both technology and risk surpass regulatory advancements. Global differences in AI ethics also contribute to the growth of regulatory challenges for multinational companies, such as compliance complexity.

Europe leans towards a rights-based perspective on AI ethics, as compared to market-based regulation in the United States (He et al., 2026). Asia adopts a public safety philosophy with strong state coordination. The implication of these differences is the need for customized organizational governance structures within regions as standards emerge in multiple regions. To executives, this means actively monitoring and aligning AI strategy globally to these emergent regional regulatory strategies as differing value and risk expectations occur.

The lack of proactive risk oversight and adequate governance is further evidenced when examining real-world risks with the value-safety-control quadrant. As a real-world view of these risks, He et al. (2026) curate 1,126 real-world GenAI ‘value safety’ incidents and cluster them into 12 value-risk categories, illustrating how value-creation goals can fail under real operational conditions. It can be postulated that incorporating value chain disruption and ecosystem-level risks creates a real need for real-time risk intelligence and forecasting. Most organizations' existing structures cannot adequately manage these risks in their strategic governance.

Responsible AI remains a strategic governance gap for many organizations, reflecting a persistent principles-to-practice disconnect between public commitments and operational control maturity (NTT DATA, 2025). NTT DATA (2025) postulates that a strategic dilemma is arising in many organizations with respect to AI as senior management and boards publicly announce transparency and corporate responsibility in AI deployment, but the organization’s capability to operationalize these statements is still limited. As a result, many data science projects do not reach production, and the literature frequently links this outcome to organizational governance gaps, unclear accountability, and insufficient KPI and control structures (Madanchian & Taherdoost, 2025). Madanchian and Taherdoost (2025) further state that most organizations struggle with integrating AI strategy governance into the strategic planning and measurement process to effectively manage risk. The main inhibitor identified for AI adoption within the context of an enterprise is the lack of a framework that promotes accountability across all functions and levels. This issue occurs due to the low conversion rates of data science projects as they proceed from the pilot phases to production. The underlying reasons are not always obvious, and the implementation is sometimes unachievable (Madanchian & Taherdoost, 2025). Madanchian and Taherdoost (2025) further state that successful AI implementation relies on well-defined organizational structures and responsibilities throughout the organization to reduce uncertainty and promote adoption of such initiatives. At a functional level, each area should have a designated individual with accountability. Executives can place this accountability within a cross-functional governing body, with the CEO and CIO overseeing it. The

ultimate measure of an AI adoption governance system is the ability to align risk considerations and KPIs with overall value objectives.

3. Methodology

This chapter describes the research design and systematic procedures used to identify, screen, analyse, and synthesize academic and high-authority grey literature on executive governance of generative AI (GenAI) in knowledge-intensive organizations. The methodology follows a PRISMA-oriented systematic review (mixed-evidence SLR) and a concept-centric thematic synthesis approach. The chapter documents study identification and screening, data extraction and coding, quality appraisal and limitations, and reliability / validity considerations, providing the methodological foundation for the findings (*Chapter 4*) and the framework derivation (*Chapter 5*).

3.1 Research Design and Literature Selection

This capstone thesis adopts a PRISMA-oriented systematic literature review (SLR) to synthesize evidence on executive governance of generative AI (GenAI) in knowledge-intensive organizations, with a focus on value creation and multidimensional risk mitigation. An SLR design is appropriate because the topic is emergent and fragmented across disciplines and practice. Academic research, authoritative practitioner guidance, standards, and regulatory / industry frameworks not only shape the executive governance of GenAI but also play a crucial role. PRISMA-light is used here as a transparent reporting structure for identifying and screening records across both academic and high-authority practitioner sources.

Google Scholar served as the primary discovery platform due to broad cross-disciplinary coverage. As a supplementary discovery channel, arXiv was used to calibrate terminology and identify early signals; records were retained only if they met the minimum evidence threshold (e.g., traceable provenance and governance relevance) and were treated as grey literature rather than as peer-reviewed studies. In addition, high-authority practitioners and standards-oriented publications (e.g., international standards, regulator guidance, and professional governance frameworks) were included when they provided explicit organizational governance mechanisms relevant to executives.

The search was conducted from the beginning of 01/2026 to the end of 02/2026 using iteratively refined search strings combining GenAI terms (e.g., “generative AI”, “large language model”, “LLM”, “foundation model”) with executive governance terms (e.g., “executive governance”, “board oversight”, “risk appetite”, “controls”, “monitoring”, “assurance”). To strengthen reproducibility in Google Scholar, screening was bounded to the first 100 results per query (sorted by relevance).

Eligibility criteria were defined a priori. Records were included if they (a) addressed GenAI or closely adjacent enabling technologies (LLMs/foundation models) in organizational settings, (b) contained

explicit governance relevance for executive decision-making (e.g., decision rights, accountability mechanisms, control architectures, monitoring and assurance practices), (c) were compatible with knowledge-intensive organizational contexts, (d) were published in English, and (e) met a minimum evidence threshold defined as either:

(i) peer-reviewed academic publication (journal or refereed conference), or

(ii) high-authority grey literature with identifiable provenance (e.g., standards bodies, regulators, professional associations, or well-established industry research publishers) and sufficient transparency to trace claims to methods, references, or documented rationale.

Records were excluded if they were purely technical without governance implications, focused solely on individual/consumer use, lacked organizational context, or were low-authority grey literature without transparent provenance. There were two steps to screening: first, titles and abstracts, and then full-text assessment. Backward reference checking and forward citation tracking were applied to included records, and any additional records identified through snowballing were subjected to the same screening and eligibility criteria.

The initial retrieval yielded n = 376 records. After deduplication (n = 48), n = 328 titles and abstracts were screened, resulting in n = 172 exclusions. Full texts were assessed for eligibility (n = 156), with n = 142 exclusions documented by reason. Backward reference checking and forward citation tracking yielded n = 36 additional records; after screening, n = 3 were retained. The final evidence base included a total of 17 records.

Step Count	Name	Papers left
1	Initial retrieval	376
2	Duplicates removed	48
3	Title/abstract screening	328
4	After title/abstract screening	156
5	Full-text assessment to be excluded	142
6	After full-text assessment	14
7	Snowballing	36
8	After Screening the Snowballing	3
9	Final included studies	17

Table 1: Results of the paper-gathering process

3.2 Data Analysis and Synthesis Approach

Building on the extraction schema, the synthesis applied a concept-centric thematic coding approach across the final included records (n = 17). Relevant text segments were coded into three domains: value creation mechanisms, multidimensional risks, and executive governance mechanisms. Coding proceeded iteratively: an initial open-coding pass captured recurring concepts, which were then consolidated into a stable codebook through cross-study comparison and refinement of category

definitions. To strengthen consistency and reduce interpretive drift, the consolidated codebook was applied in a second pass to all included records.

The synthesis deliberately integrated evidence from diverse study designs - qualitative multi-actor studies, organizational case work, auditing and governance assessments, and cross-country regulatory analyses - because GenAI governance is shaped by socio-technical dynamics, resource dependencies, and external regulatory environments. Evidence from multi-actor and longitudinal contexts was used primarily to identify governance gaps, coordination failures, and feasibility constraints; auditing-oriented studies informed monitoring, escalation, and control requirements; and cross-national analyses informed framework requirements for regulatory heterogeneity and external alignment.

The resulting themes are reported in *Chapter 4*, organized by recurring patterns across the included records (rather than record-by-record summaries). These themes were subsequently mapped into a value–risk–governance structure to provide traceability from the SLR evidence base to the Executive GenAI Governance Framework developed in *Chapter 5*.

3.3 Quality Assessment and Limitations

Because the review includes both academic and high-authority grey literature, quality appraisal explicitly accounts for source type. Peer-reviewed studies were evaluated primarily based on methodological transparency and evidentiary support. Where grey literature did not report empirical methods, it was treated as a governance artifact rather than as an empirical study. This ensures that non-academic sources contribute as governance artifacts with clear legitimacy and are not treated as empirical studies unless they provide documented data and methods.

This systematic literature review synthesizes evidence on executive governance of generative AI in knowledge-intensive organizations. Because the evidence base is emergent, heterogeneous in study designs, and rapidly evolving in practice, explicit quality assessment is necessary to calibrate confidence in the resulting themes and inferences. The purpose of the quality assessment in this thesis is therefore not only to judge individual study “quality” in an abstract sense, but to determine (a) how trustworthy and transparent each study is, (b) how directly it contributes to the executive governance focus of this thesis, and (c) how applicable its insights are to knowledge-intensive organisational settings. In line with that objective, quality assessment is used descriptively to weigh confidence in the synthesis rather than as a purely mechanical pass / fail filter. This is appropriate because GenAI governance scholarship currently includes both conceptual and empirical contributions and excluding all non-empirical studies would risk removing important governance constructs that remain under-tested but are still relevant for executive decision-making.

Quality appraisal approach and criteria

The quality appraisal applied a pragmatic, governance-relevant set of criteria that reflects the needs of an executive-oriented synthesis. Each included study was assessed against five dimensions:

1. Clarity of purpose and scope: Whether the study clearly defines what aspect of GenAI (or adjacent enabling technologies) is being addressed, which organizational setting is assumed, and what governance problem it aims to inform. Explicit scope boundaries and clearly articulated objectives in studies enhance interpretive reliability by precisely positioning their contribution within the review's concept domains.

2. Methodological transparency and appropriateness: Whether the study provides sufficient methodological detail for its design type (e.g., data sources, sampling, analytic steps, and limitations for empirical work; or argument structure, conceptual definitions, and internal logic for conceptual work). For this review, transparency matters at least as much as design type: a modest empirical study with clear procedures may be more useful than a larger study that provides limited detail on how claims were derived.

3. Relevance to executive governance: Whether the study directly addresses executive-level governance mechanisms (e.g., decision rights, accountability structures, control gates, monitoring and assurance practices) rather than only discussing technical performance, generic ethics, or non-organizational use. The review's research question focuses on governance structures that executives can implement, so studies providing operational governance implications were weighted more strongly than those that only offered high-level principles without translating them into organizational mechanisms.

4. Coherence and evidentiary support: Whether claims are supported by data, documented reasoning, or explicit linkage between evidence and conclusions. For empirical studies, this includes the credibility of interpretations relative to the reported data; for conceptual studies, this includes the absence of internal contradictions and the use of clearly defined constructs. Coherence is particularly important in GenAI governance because the same terms (e.g., “governance”, “risk”, “responsible AI”) are often used inconsistently across disciplines.

5. Applicability to knowledge-intensive organizational contexts: Whether the study’s assumptions and examples plausibly generalize to knowledge-intensive work environments (e.g., professional services, R&D, software, analytics, corporate functions) where knowledge, intellectual capital, and decision-making quality are central. Studies rooted in highly sector-specific contexts can still be valuable, but their insights require careful translation. Applicability is therefore assessed not as “whether the study is universally generalisable”, but whether its governance mechanisms can realistically be adapted to the characteristics of knowledge-intensive organisations.

For each study, the above dimensions were reviewed and recorded in a quality appraisal log. The appraisal was then used to calibrate confidence in the derived themes during synthesis. In practice, themes supported by multiple studies with strong transparency and direct executive governance relevance were treated as higher-confidence patterns. Themes supported primarily by single studies, highly sector-specific evidence, or conceptually broad claims without clear operationalization were treated as lower-confidence and were framed more cautiously in the findings and framework derivation.

The quality assessment influenced synthesis in three concrete ways. First, it shaped how strongly particular claims were expressed. Where a theme was consistently supported across multiple studies and aligned with explicit governance mechanisms, it was reported as a robust cross-study pattern. Where evidence was partial, indirectly connected to executive governance, or based on limited transparency, the theme was framed as indicative and positioned as an area requiring further validation.

Second, quality appraisal informed how the review handled the “principles-to-practice” issue. GenAI governance literature frequently includes normative statements about what organisations “should” do. The appraisal therefore distinguished between (a) prescriptive recommendations with explicit operational mechanisms and (b) broad aspirational principles without actionable translation. In the synthesis, actionable governance mechanisms were prioritized because the thesis aims to propose an executive-ready framework. Normative principles were not discarded but were treated as inputs that require conversion into organizational design elements (roles, processes, controls, escalation paths, and monitoring).

Third, appraisal supported the identification of boundary conditions. Studies grounded in specific regulatory environments or sectors can be high quality while still having limited transferability. The appraisal therefore captured contextual constraints and helped avoid overgeneralization. When themes are reported in *Chapter 4*, they are interpreted through the lens of knowledge-intensive organizations, but with explicit recognition that governance implementation must be adapted to sectoral realities, organizational maturity, and regulatory environments.

Several limitations affect both the included evidence base and the review design. These limitations do not invalidate the synthesis, but they constrain the strength and scope of claims and motivate cautious interpretation.

1. Platform and indexing bias: The review relies on a broad scholarly discovery platform with non-transparent indexing and ranking mechanisms. Even with bounded screening, relevant studies can be missed due to how search results are ranked, how venues are indexed, and how terminology differs across disciplines. This limitation is partly mitigated through snowballing (backward and forward citation tracking), but residual coverage bias remains possible.

2. Publication bias and maturity effects: Peer-reviewed publication cycles can lag real-world GenAI adoption. As a result, some of the most practically relevant governance patterns may not yet be fully represented in peer-reviewed outlets, while early academic work may overemphasize certain risks or governance ideals that are still being tested in practice. The review addresses these concerns by distinguishing between themes that reflect mature, repeatedly observed governance mechanisms and those that reflect emerging proposals.

3. Conceptual ambiguity and definitional inconsistency: “AI governance” is not a single stable construct. Across disciplines, governance may refer to corporate governance structures, risk management systems, compliance frameworks, ethics principles, technical controls, or combinations

of these. This variability increases the risk of construct drift during synthesis. The review reduces this risk by operationalizing governance in executive terms (decision rights, accountability, policies, controls, monitoring and assurance), but definitional heterogeneity remains a limitation when comparing studies that use different conceptual frames.

1. Heterogeneity of study designs: The included studies likely span conceptual analyses, qualitative work, surveys, auditing-oriented studies, and cross-country / regulatory comparisons. Such heterogeneity is useful for an executive governance topic because it captures organizational and regulatory complexity; however, it also reduces comparability and limits the feasibility of quantitative aggregation. Consequently, the synthesis necessarily remains qualitative and interpretive, and the resulting framework is conceptually derived rather than statistically validated.

2. Context specificity and transferability constraints: Governance mechanisms that appear effective in one sector or jurisdiction may not transfer directly to another. Organizations that rely heavily on knowledge have very different levels of risk tolerance, compliance exposure, intellectual property sensitivity, and operational speed. Such variability limits the extent to which a single governance design can be presented as universally optimal. The thesis addresses this concern by framing the framework as a modular executive governance architecture that can be adapted based on context, maturity, and risk tiers; nevertheless, transferability remains a limitation until the framework is tested across multiple organizational settings.

3. Rapid evolution of GenAI tooling and regulation: GenAI capabilities, vendor ecosystems, and regulatory expectations evolve quickly. This means that the evidence base can become outdated faster than in more stable domains. The thesis therefore treats the derived framework as an adaptive governance structure that emphasizes continuous monitoring, review cycles, and revision mechanisms rather than static compliance checklists. Still, the fast-moving landscape constrains the “timelessness” of specific implementation details or risk priorities.

4. Literature-only design and lack of primary validation: The thesis does not include primary data collection. While an SLR is appropriate for synthesizing and structuring the knowledge base, the resulting framework is not empirically validated within a specific organization. The framework should therefore be understood as an evidence-informed conceptual governance model, not as a proven intervention. This limitation is addressed explicitly by positioning empirical validation as a future research direction (e.g., case studies, design science evaluation, or longitudinal implementation studies).

These limitations imply that the framework proposed in *Chapter 5* should be interpreted as an executive-ready starting point that structures decision rights, accountability, controls, and monitoring in a coherent system while remaining adaptable to organizational context. The framework is designed to reduce the risk of governance blind spots by integrating value creation and risk mitigation into a single governance logic. However, executives and organizations adopting the framework should tailor it to their sector’s risk profile, regulatory environment, maturity level, and strategic

priorities and should treat implementation as an iterative process with built-in learning loops, auditability, and revision.

Overall, the quality assessment and limitations analysis strengthens the thesis by clarifying which conclusions are strongly supported across the evidence base and which remain indicative. This transparency supports responsible interpretation of the synthesis and provides a defensible foundation for deriving an executive governance framework that is both practically meaningful and academically rigorous.

3.4 Reliability and Validity Considerations

This systematic literature review is designed to produce a transparent and defensible synthesis of executive governance mechanisms for generative AI (GenAI) in knowledge-intensive organizations. Because the thesis derives a conceptual framework from heterogeneous literature, reliability and validity considerations are essential to clarify how methodological threats were managed and how the resulting insights should be interpreted.

Reliability in this review concerns whether the study identification, screening, and coding steps were applied consistently and in a way that another researcher could broadly replicate. First, procedural reliability was strengthened by applying predefined eligibility criteria and a two-stage screening process (title / abstract screening followed by full-text assessment). Reasons for full-text exclusions were recorded, enabling traceability from the initial record set to the final included sample. Second, because discovery engines can return variable results over time, reproducibility was supported through bounded screening (screening only the first 100 results per query, sorted by relevance) and by reporting the search period and core search terms. Third, interpretive reliability risks arise during qualitative synthesis, particularly when studies use different terminology for similar constructs (e.g., “governance”, “controls”, “oversight”, “assurance”). To mitigate this, data analysis was structured using three stable concept domains—value, risk, and governance—and coding followed an iterative approach. An initial open-coding pass captured relevant concepts, which were then consolidated into a stable codebook. A second coding pass was performed to reduce interpretive drift and improve consistency across the included records.

In the context of a literature review, internal validity is best understood as credibility: whether the synthesis accurately represents what the included evidence supports. Key threats include selective interpretation, overemphasizing single-study claims, and conflating normative recommendations with empirically supported mechanisms. These risks were addressed by applying a concept-centric synthesis logic: themes were derived from recurring patterns across studies rather than from isolated examples, and claims were framed in proportion to the strength and convergence of the underlying evidence. In addition, the synthesis explicitly distinguishes between (a) operational governance mechanisms (e.g., decision rights, accountability, control gates, monitoring and escalation practices) and (b) high-level principles that require translation into implementable organizational structures.

This distinction reduces the risk of treating aspirational guidance as if it were a tested governance design.

Construct validity is challenged by the definitional ambiguity of “AI governance” across disciplines. To strengthen construct validity, governance was operationalized in executive-relevant terms: decision rights, accountability structures, policies, controls, and monitoring / assurance mechanisms across the GenAI lifecycle. Risks were coded as multidimensional (strategic, operational, legal/regulatory, ethical, and security-related), and value was coded in terms of organisational value mechanisms rather than generic “benefits”. This operationalization served as a consistent coding anchor across studies, reducing the likelihood that different disciplinary meanings of governance would distort the synthesis.

External validity in an SLR is primarily the transferability of findings to comparable organizational settings. Because knowledge-intensive organisations differ substantially in sector constraints, risk appetite, regulatory exposure, and maturity, the framework derived in *Chapter 5* should not be interpreted as a universal “one-size-fits-all” model. Instead, it is designed as a modular governance architecture that can be tailored to context through use-case tiering, risk thresholds, and differentiated controls. Moreover, as the review is literature-only and the GenAI landscape evolves rapidly, the framework is conceptually grounded but requires empirical validation and refinement through real-world application.

Overall, reliability and validity were strengthened through transparent screening and exclusion documentation, bounded search procedures to improve reproducibility, iterative coding with a stable codebook to reduce interpretive drift, and explicit operationalization of governance constructs to improve conceptual alignment. Remaining limitations relate to platform and terminology variability, heterogeneity of study designs, and the literature-only nature of the work. These considerations frame the results in *Chapter 4* and support a cautious but actionable interpretation of the framework developed in *Chapter 5*.

4. Findings: Value and Risk of Generative AI

The discussion that follows delves into the tangible benefits that generative AI has to offer, as well as the risks that may hinder organizations from achieving success. This section explains the benefits organizations want and the risks that prevent them from getting them. This section in turn adds to the discussion of strategic governance in setting boundaries on how knowledge-intensive organizations may responsibly use generative AI.

Author (Year)	Findings
Bolden et al. (2024)	- Value creation & productivity outcomes across functions automation of knowledge tasks (e.g., RFP support) - Cost reduction across value chain, operational productivity uplift (e.g., scheduling/decision support), improved customer outcomes via AI-

	enabled interventions - Value constraints via data quality, workforce readiness, and governance maturity; value not realised without organisational capabilities - Need for executive oversight that aligns use cases to strategy; investment in data readiness, workforce enablement, and cross-functional governance to move beyond pilots; value tracking tied to operational adoption
Smith (2025)	- Scaling gap, adoption readiness, executive oversight highlights that only a small portion of organisations successfully scale GenAI; value comes when organisations invest deliberately in governance + skills + monitoring - Adoption failures driven by governance gaps, weak KPIs, insufficient internal competencies, and lack of risk readiness - Need for "technical and strategic oversight" structures (committee/operating model), KPI alignment across executive team, and governance that enables scaling beyond proof-of-concept
Waisam & Silver (2025)	- Success vs. failure patterns in GenAI projects - Value is contingent on governance integration; projects succeed when risk, talent, and KPIs are built early - Failure drivers: lack of internal talent, inadequate risk frameworks, poor KPI design, customer safety + data security failures; unmanaged risk destroys value - Executive governance must anchor GenAI to business outcomes; establish risk oversight, talent / upskilling strategy, and KPI architecture early; implement customer safety/data-security controls as "go/no-go" gates
Chandra & Rahman (2026)	- Customer-facing value mechanisms (service / experience) - Customer value via (1) mechanical AI (automation, speed), (2) thinking AI (personalisation, intent recognition), (3) feeling AI (rapport/empathy) and loyalty effects - Implicit risk exposure if customer interactions become unreliable, manipulative, or untrustworthy (trust as a fragile asset in customer use cases) - Governance must set boundaries for customer-facing GenAI, define quality thresholds, escalation paths, and monitoring for customer harm; treat customer trust as an executive KPI dimension
Shomali et al. (2025)	- Value in knowledge-intensive, high-stakes applied context (digital

	health example) - Value through personalised interventions and measurable outcome improvements (e.g., adherence/engagement and clinical outcome movement as reported in the study) - High-stakes context implies heightened requirements for accuracy, oversight, and risk controls (quality failures have direct harm potential) - Executive governance must define risk tiers and require stronger assurance for high-stakes contexts; monitoring must track both outcomes and safety/quality signals
Kaul et al. (2025)	- Security/misuse, cyber and operational risk expansion - Mitigating risk protects value; continuous monitoring and real-time response capability are prerequisites for sustainable GenAI use - Misuse by malicious actors, expanded attack surface, need for proactive defence; argues traditional reactive controls are insufficient - Executives need layered defence governance: monitoring, anomaly detection, incident response, scenario-based simulations and red-teaming integrated into governance; define accountability for security operations
Gehrmann et al. (2025)	- Reliability, testing limits, post-deployment risk control - Value depends on controlling model failure modes; without robust QA, value collapses Hallucinations/quality failures, unpredictability, limits of conventional testing; need for scenario testing/red teaming; requirement for audit trails and continuous improvement loops - Governance needs deployment-stage controls: human-in-the-loop QA for quality-sensitive work, stress testing, incident reporting loops, auditability/traceability, continuous monitoring linked to escalation protocols
Poon Affat (2025)	- Standards, compliance implementation gaps, fairness/privacy dominance - Standards can support trust and governance structure but may not reflect full operational reality - Implementation inconsistency; over-focus on fairness/privacy vs broader operational risks; capability/maturity gaps prevent effective adoption - Executives must adapt standards into an operational governance system: capability building, tailored controls beyond fairness/privacy, and a living risk register; standards should be translated into roles, processes,

	and metrics
Sinclair & Mehta (2023)	- Ethics, privacy, transparency and explainability as governance requirements - Ethics / privacy governance enables sustainable adoption by preserving trust and compliance exposure control - Privacy, fairness harms, inclusivity / security harms; lack of transparency increases regulatory scrutiny and stakeholder distrust - Governance must embed transparency/explainability, documentation, audit trails, and privacy-enhancing design; treat ethics / privacy compliance as executive-level accountability with measurable indicators
Strauss et al. (2025)	- Governance research gaps (deployment-stage and high-stakes focus) - Evidence mapping indicates that research attention is skewed toward pre-deployment concerns (e.g., bias checks) while deployment-stage behavioural and operational risks remain under-addressed; only a small minority of studies focus on high-stakes application domains where liability and societal harm potential is highest. => Creates an executive blind spot because many material risks emerge post-deployment through real-world use, drift, misuse, and organisational process failures rather than during model development alone - Executive governance must compensate through post-deployment controls: continuous monitoring and incident learning loops, clear escalation and accountability paths, and deployment governance that goes beyond static “pre-deploy compliance” routines
Batool et al. (2025)	- Organisational-level AI governance frameworks and implementability gaps - Review of organisational AI governance frameworks suggests that relatively few enterprise-level frameworks are available and that many lack actionable implementation guidance - Governance attention is often imbalanced toward fairness and privacy, while other executive-critical dimensions (e.g., accountability structures, reporting lines, operational controls, engagement mechanisms) are less developed - Limits practical adoption and scaling because executives lack a clear operating model for translating principles into organisational decision rights and control architectures - Executive implication: governance must be designed as an implementable enterprise systemroles, responsibilities, reporting,

	assurance, and operational processes-rather than as a principles-only framework
Attard-Frost & Lyons (2024)	- Ecosystem fragmentation and multi-actor governance complexity - Governance environment surrounding AI is highly fragmented, characterised by many actors producing overlapping guidance and policies, which weakens coherence and coordination - Fragmentation increases ambiguity for organisations about which standards and practices to follow and reduces the transferability of governance approaches across jurisdictions and sectors - Result is a practical executive challenge: aligning internal governance to an external environment that is dynamic, crowded, and often inconsistent - Executive implication: organisations require explicit "external alignment" governance capability-structured stakeholder engagement, regulatory/standards intelligence, and an internal translation layer that converts fragmented external guidance into one coherent governance operating model
Gianni et al. (2022)	- Limits of national / sector "soft policy" and non-binding guidance - National and sector AI strategies frequently rely on "soft" governance instruments (principles, voluntary guidelines, ethics statements), which creates variability and inconsistent implementation across organisations and industries - Non-binding guidance may raise awareness but does not reliably produce scalable, enforceable organisational controls, especially in cross-jurisdiction contexts - Amplifies executive burden: firms must operationalise governance internally even when external regimes are non-executable - Executive implication: governance must be implemented as enforceable organisational mechanisms (controls, accountability, monitoring, assurance) rather than treated as a compliance exercise against abstract principles
Madanchian & Taherdoost (2025)	- Cross-thematic patterns: scaling depends on capabilities, investment, and executive leadership - Synthesis indicates that organisations achieve higher GenAI implementation success when they invest deliberately in enabling infrastructure and workforce learning rather than treating GenAI as a purely technical deployment

	- Constraints emerge from resource limitations, maturity differences, and ethical divergence across contexts, which undermines “one-size-fits-all” governance models - Executive implication: governance should be modular and maturity-based (tiered controls, adaptable operating model), with executive leadership as the central lever to align investments, capability building, and risk thresholds to organisational strategy
Walz & Firth-Butterfield (2021)	- Workforce disruption and socio-psychological risk as governance domain - AI adoption creates material workforce risks, including job displacement pressures and negative psychological and social impacts that can reduce morale, engagement, and trust in leadership if unmanaged - While long-term effects may include new job creation, short- to mid-term transition costs and cultural impacts can directly undermine transformation outcomes - Executive implication: governance must include workforce transition mechanisms - reskilling and redeployment plans, transparency and employee involvement, and change-management indicators - so adoption does not erode organisational stability and performance
Sharma (2023)	- High-stakes model and data risk in critical domains High-stakes deployment contexts expose organisations to amplified model risk and data quality failure modes, where errors and low-quality outputs can translate into direct harm, liability, and reputational damage - Governance requirements therefore increase with criticality: assurance, oversight, and quality controls must be stronger than in low-stakes contexts - Executive implication: implement explicit risk tiering and stronger assurance for high stakes use cases (human control/oversight, stricter validation gates, monitoring of safety/quality signals, and accountability for quality outcomes).
Soroori Sarabi (2025)	- Foundational capability gaps as a root cause of value failure and risk exposure - Organisational vulnerabilities - insufficient data quality, scarcity of skilled personnel, and weak or uncoordinated operational management models - are identified as key drivers of both unrealised value and increased risk exposure in GenAI deployments - Without these foundations, organisations remain stuck at pilot stage,

	lack effective monitoring and control, and struggle to sustain safe scaling. - Executive implication: governance must prioritise foundational enablement (data readiness, talent strategy, operating model maturity, ownership clarity) as prerequisites for scaling and for maintaining risk controls over time.

Table 2: Summarised findings of included records (peer-reviewed + high-authority grey literature)

4.1 Value Creation Mechanisms

Generative AI technologies have emerged as crucial tools for enhancing organizational efficiency, with the automation of time-intensive tasks being a key driver of value creation. For example, procurement use cases include an RFP assistant that can save around five hours per draft and AI support that enables 50% faster tender drafting and offer comparison. It is also utilized for SEO content generation and results in a 95% reduction in costs and a fiftyfold decrease in the production timeline. The automation of repetitive tasks and the improvement of processes are valuable tools for organizations and result in cost reduction as well as allowing organizations to use resources for more strategic purposes (Bolden et al., 2024). However, the quality of data, workforce preparedness, and governance are essential for achieving the desired benefits from the use of AI technologies. Without these elements, AI may not deliver anticipated value for organizations.

Generative AI also adds value in the business function area by impacting cost transformation across the entire value chain. These implementations result in an average cost reduction of 8-12%, which equates to hundreds of millions of dollars in business value (Bolden et al., 2024). Organizations must strategically apply technology, which will affect the scalability and cost-saving strategies for AI. Poor management of governance and implementation of technologies may hinder cost reduction and the overall value creation of organizations.

The effectiveness of AI solutions has not been exclusive to white-collar applications; organizations can create value through generative AI in operational spaces as well. For example, an AI-driven scheduling and decision-support system can increase productivity 5-10% as well as decrease job duration by 15-20% and improve rework rates in frontline workers (Bolden et al., 2024). Organizations must empower their employees and in still confidence in them as they implement AI solutions to enhance operational practices.

This change will help support the strategic alignment of technology and maximize value. The effective integration of generative AI technologies in organizations also depends on data governance and workforce development.

Organizations that invest thoughtfully and purposefully through a technical and strategic oversight committee tend to have more consistently reported AI success (Smith, 2025). The use of AI applications may have a significant impact on cost reduction as well as customer experiences, but the quality of data and the workforce are crucial to generate optimal outputs.

One of the primary ways generative AI creates value for customers is by adding mechanical, thinking, and feeling functionalities to the interactions between organizations and customers. Mechanical AI enhances the functionality by automating routine customer service issues with more accuracy, less effort from the customer, and faster turnaround times. This increases customer satisfaction (Chandra & Rahman, 2026). Thinking AI delivers personalized, real-time interactions and dynamically recognizes customer intent and mood, which is compelling in creating loyalty and emotional bonds between the customers and the company.

Feeling AI builds customer trust and rapport by providing more human-like, emotionally rich responses and offers increased interactions with digital customers by adding empathy (Chandra & Rahman, 2026). When companies are looking for strategic value creation, customer interactions are an imperative dimension to address through technology.

Many areas within an organization, such as productivity and innovation, also improve with the successful use of generative AI. With the application of generative AI to customer interactions, for example, an organization can expect to receive twice the return on investment than it typically would. Moreover, AI-powered interventions such as digital health and asset maintenance are achieving higher engagement and improved outcomes for their customers as well as the organization (Bolden et al., 2024; Shomali et al., 2025). Understanding the necessary actions to optimize productivity and enhance innovation through AI technologies is crucial for reaping these benefits. A small portion of organizations (8%) are successful at strategically leveraging technology and, by design, creating value within the organization. The remainder are stuck at the proof of concept (pilot) stage, mainly because they don’t address challenges for adoption, which are AI risk, readiness, and scalability (Smith, 2025). For example, most organizations struggle with risk and governance issues with the utilization of AI and how to align strategic and financial KPIs across the executive team. Generative AI adds value in knowledge-intensive sectors through personalized interventions and data-driven solutions.

For instance, a specific case in digital health demonstrated its ability to enhance glycaemic control in patients with type 2 diabetes. Through the implementation of a digital health coaching platform, patients saw a 1.2% decline in A1C over the control groups as well as improved medication adherence, food consumption habits, and exercise habits (Shomali et al., 2025). With technological innovation and personalized healthcare for customers, organizations can create more engagement with their customers, ultimately resulting in improved value and positive outcomes for organizations. Generative AI value creation occurs in both the physical and operational worlds; however, its deployment is highly dependent on an organization’s ability to harness robust data, workforce, and governance. As research is developed, the society is beginning to see that organizations that are more successful in AI tend to purposefully use technical and strategic oversight (Smith, 2025; Bolden et al., 2024). 74% of organizations claim to implement generative AI solutions effectively, and they need a way to be more purposeful in their implementation and utilize a risk-aware, integrated, proactive executive governance framework to better address key challenges across a broad range

of functions (Smith, 2025). The proper workforce, strategy, data, and technology architecture must be in place for AI adoption and value creation to be successful. The executive governance framework should also create a strong partnership among finance, HR, IT, supply chain, and the board to create value for the company. The following chapters will help identify value creation strategies for generative AI within knowledge-intensive organizations and begin building a useful governance framework for AI at the executive level to realize value.

4.2 Risk Landscape

The exploitation of generative AI by malicious actors is an example of how the breadth and depth of cyber and operational risks are expanding. Cybercriminal group GTG 5004 leveraged Anthropic's Claude to automate attacks on sectors such as healthcare, emergency services, government, and religious organizations. Organizations that are knowledge-intensive need to be aware of the implications and utilize a layered strategy to address the risks effectively, as traditional controls are not sufficient to address these concerns. There must be continuous monitoring of generative AI outputs and early mitigation of potential harm. Such an approach requires the capability to respond effectively to actual and potential harm in real-time and identify anomalies (Kaul et al., 2025). With the continuously evolving nature of the technology, the attack surface for malicious actors expands as well. Technical and process vulnerabilities are leveraged to take advantage of this reality, thus rendering traditional response strategies, which are often reactive, unsuitable to address rapidly developing AI-related harm. To respond swiftly to unexpected events, organizations need to establish scenario-based simulations, red-teaming, and incident response with increased AI penetration.

Top executives will need to develop advanced defence capabilities and establish a governance framework that prioritizes vigilance and preparation. One method is the adoption of structured systems that will help organizations mitigate potential AI risks, as proposed by ISO/IEC 42001, which assists in governing AI use (Kaul et al., 2025; Poon Affat, 2025).

Harm associated with the misuse of generative AI comes in the form of loss to organizational resources, service disruption, and loss of customer trust. On top of this, there could be damage to regulatory and legal status and an organization's reputation. Structured executive governance, particularly the incorporation of AI risk indicators, can facilitate the establishment of trust and confidence in knowledge-intensive organizations characterized by elevated risk and compliance levels.

One method is the structured system implemented by the ISO/IEC 42001 framework, which establishes formal reviews of incident reports and defines best practices to foster reliability and trust among stakeholders (Kaul et al., 2025; Poon Affat, 2025).

A main concern in the field of generative AI is "model hallucinations," where the AI can produce incorrect or fabricated data and insights. For knowledge-intensive organizations, this could be a problem, as the propagation of false or misleading content may lead to regulatory enforcement and

reputational harm. When AI is used for high-risk processes, there needs to be "human-in-the-loop" quality assurance, meaning a human is present when the outputs are reviewed. MLOps dashboards must incorporate and continuously monitor AI bias and error detection tools. Without these interventions, organizations risk the credibility of their outputs (Kaul et al., 2025). This issue is made more complex as AI systems use proprietary models to produce outputs, which are also constantly in flux due to continuous AI evolution. Traditional quality testing methodologies do not apply to generative AI applications. AI applications cannot be adequately tested via regular means to detect issues such as hallucination or any other form of low-quality, biased, or discriminatory output.

One way to effectively test generative AI quality is through red-teaming and scenario-based simulations to pinpoint AI failure or harm. Governance systems need to implement accountability and traceability that incorporate audit trails. Furthermore, there needs to be a system for continual improvement, or monitoring incident reports, where the AI can learn from past errors and then improve. Such an approach contributes to an environment that adheres to regulatory expectations and leads to risk and compliance resilience (Gehrmann et al., 2025; Kaul et al., 2025). One of the prevalent risks in the world of AI is bias and discrimination, which occurs when generative AI models enhance and reinforce unfair practices already in place in society. In this case, AI can replicate and perpetuate harmful behaviours if not properly addressed. Regulatory bodies are becoming increasingly aware of these concerns, as indicated in the NAIC Model Bulletin on AI Use in Insurance. The issue is that, when organizations adopt this bulletin, there are inconsistencies in implementation, making adherence to the standards complicated. To mitigate bias and discrimination, companies can develop representative datasets, conduct fair audits to address the problem, and use AI tools to detect potential biases and automatically resolve the issues. This issue is especially critical for knowledge-intensive organizations, as these systems are often utilized in high-stakes decision-making situations with great effect on individuals and society (Poon Affat, 2025; Sinclair & Mehta, 2023).

The strategic and regulatory challenges in mitigating bias and discrimination risk are exacerbated as businesses and institutions operate globally across different geographical boundaries. Organizations need to know how each jurisdiction adopts regulations and how strictly they are enforced. For instance, organizations may see variability in whether the NAIC Model Bulletin on AI Use in Insurance is adopted, as each state has different guidelines, policies, and enforcement measures. This poses a compliance problem and can increase the risk of punitive consequences and reputational damages. Agile executive governance is needed, where the organization can easily adapt to changing laws and ensure compliance in all geographies. Fairness should be a core governance principle and integrated into the AI strategy, especially in high-risk processes. There must be defined key indicators to measure AI fairness, and performance should be tracked on an ongoing basis to guarantee compliance with AI regulations and mitigate regulatory fines (Poon Affat, 2025).

The use of AI also brings with it increased risks around ethical and privacy aspects. Ethical and privacy compliance needs to be a top priority, as research shows that building design decisions around privacy-enhancing technologies, as well as privacy and ethical guidelines, has made measurable progress in mitigating fairness, inclusivity, and security harms with AI. Knowledge-intensive organizations must rigorously address ethical risks related to the high volumes of sensitive personal and private data they collect, store, and transfer. Failures in this space will damage stakeholder confidence, lead to punitive actions from regulators, and ultimately have an adverse impact on the organizational reputation. Furthermore, failure to address these risks can result in increased fines for non-compliance. Ethical failures in AI can have a direct and cascading negative impact on compliance requirements (Sinclair & Mehta, 2023).

Transparency and explainability of how an AI decision occurred are mandatory for the executive governance of ethics and privacy risks with generative AI. The lack of transparency can cause distrust among stakeholders, increase regulatory investigation, and result in AI corrective measures. Organizations must integrate transparency into the design of AI as well as develop formal documentation and audit trails of decision-making processes. To improve stakeholder confidence and address ethical risks in AI, the explanation mechanism needs to be addressed at the beginning of the model training process. This will ensure that AI is explainable and the organization is not exposed to liabilities (Sinclair & Mehta, 2023).

Executive governance for AI risk management is complex and fragmented. There are international AI standards that serve as frameworks, such as ISO/IEC 42001. Many models and approaches to AI risk governance are available to knowledge-intensive organizations, for example, the HUDERAF framework. All these frameworks provide steps for the organization to follow. The issue comes in with aligning each framework to the organization's strategy and values, as well as knowing where to put resources to implement and how to comply with them. There must be an adaptable system, with direct executive engagement. Furthermore, the risk registers need to be updated constantly to address and adapt to risk incidents as generative AI technologies continue to evolve. An agile and dynamic risk management framework is paramount for knowledge-intensive organizations to sustain their reputation (Kaul et al., 2025; Poon Affat, 2025).

As the generative AI ecosystem becomes ever more prevalent and evolves to become a mainstream offering, the executive governance of risk management will need to focus on developing controls to pre-emptively manage harm to sustained operations and stakeholder confidence. Initiatives such as scenario-based stress testing, real-time performance monitoring of controls, and iterative feedback loops will increase the executive's capability to manage risks posed by AI across an organization. These systems ensure that governance controls operate within thresholds aligned to the company's risk appetite and appetite to harm, thereby supporting the achievement of governance objectives (Gehrmann et al., 2025).

In closing, as the field of generative AI and large language models grows, there will be risks associated with it. Knowledge-intensive organizations need to address these risks early on and

address the following points: technical, compliance, ethics, and security. They should incorporate metrics to monitor for any potential harms with respect to fairness and transparency. They should be transparent about how the system is working and how the data is handled, including obtaining user consent for data collection. Organizations should also monitor AI performance to determine if the model is delivering effective solutions and they should adjust if needed to improve effectiveness. With AI playing an ever-increasing role in how organizations operate, it's crucial to put an effective risk management plan in place early in the process.

4.3 Comparative Analysis of Value–Risk Tensions

The comparative examination of value–risk tensions for generative AI shows that the potential rewards of this technology and the risks associated with its implementation are interlinked. Empirical findings indicate that organizations could achieve 30%-50% enhanced operational processes and cost savings of 8%-12% throughout the value chain through the application of generative AI (Smith, 2025; Waisam & Silver, 2025). However, they are contingent on an organization’s risk management capabilities. Organizations that have weak executive governance (that is, governance, KPIs, internal competencies) typically remain at the pilot phase for years, and only about 8% fully scale generative AI use cases (Smith, 2025; Waisam & Silver, 2025).

An organization's ability to effectively govern across value/risk dimensions determines the value it derives from generative AI technologies. Although generative AI technologies can dramatically improve productivity, the improvements may be undermined by risks such as biased outputs, data privacy failures, and strategic, operational, or regulatory misalignments, which can diminish any gains. In some situations, risk factors left unaddressed can lead to negative net organizational productivity.

As Smith (2025) and Waisam and Silver (2025) noted, the degree to which risk management is integrated within executive governance often determines the organization’s ability to harvest value from generative AI technologies.

A closer look at organizations that were able to proceed beyond the pilot stage of generative AI indicates they had an intention to internally govern (that is, monitoring, risk assessment, workforce upskilling) these technologies. The analysis provided by Smith (2025) further substantiated the hypothesis that sustaining value creation from generative AI technologies requires proactive action by executives appropriate to the context of their organizations. Generative AI deployments’ capability to deliver real value is not something achieved by merely installing and using them. Scaling AI solutions effectively is more likely to be tied to how well an organization aligns its goals with appropriate organizational oversight structures than to the technology's readiness. Without such operational readiness, most AI transformations remain at the pilot phase, providing very limited value.

There is also a clear misalignment between generative AI and holistic organizational oversight capabilities. Organizational inertia, incomplete risk controls, and a lack of capabilities to manage

emergent technology risks can all contribute to the absence of adequate organizational oversight. As Waisam and Silver (2025) indicated, these issues must be addressed to facilitate scaling. An absence of holistic organizational oversight prevents organizations from effectively scaling AI solutions and addressing inherent technological and human biases. For generative AI to create sustained value creation opportunities, executive governance must address organizational inertia and fragmented risk controls.

Model bias, customer safety, data breaches, and regulatory compliance have a direct impact on value realization in generative AI deployments. Numerous instances have shown that AI-generated misinformation, cybersecurity breaches, and ethical AI transgressions have directly impacted organizations' bottom lines. While generative AI has great potential to improve overall productivity, it introduces risk factors that could harm organizational stability. People often overlook these factors, which can lead to significant financial losses, reputational harm, and even legal consequences (Waisam & Silver, 2025; Gehrmann et al., 2025). For example, model hallucination and bias can result in loss of customers and can be legally challenged if regulatory guidelines are violated. Similarly, data breaches can quickly degrade the benefits of AI. All applications that use any type of AI may be open to data breaches and malicious attacks. Companies may face heightened compliance requirements and reputational risks if they fail to protect their underlying technology infrastructure. These incidents restrict the potential value-producing capabilities of generative AI. As Waisam and Silver (2025) showed, organizational data protection preparedness is key to the scalability of generative AI deployments.

Regulatory risks such as increasing compliance obligations and a lack of standardization across regions can hinder the operations and returns of an organization using AI. Legal uncertainty and different and often incongruous standards of operational conduct across industries and geographies can constrain the value potential of generative AI technologies. These factors heighten operational risk and make value creation in the generative AI space less predictable. Furthermore, in the absence of proper executive governance processes, organizations may have difficulty adapting quickly to evolving regulatory environments (Smith, 2025; Waisam & Silver, 2025). Unmanaged risks can result in loss of reputation, business interruptions, and damage to operations; therefore, they hinder generative AI-driven value generation.

For example, errors in AI models could be a sign of poor model integrity or insufficient data quality and input validation. Operational governance factors such as poor incident escalation processes can lead to unresolved risks. The organizations in Gehrmann et al.'s (2025) research made it clear that real-time performance monitoring, risk controls, and overall risk preparedness were critical to maintaining stable environments. Furthermore, organizations can face challenges that derail generative AI implementations (Waisam & Silver, 2025). Poor monitoring mechanisms or risk readiness can disrupt organizational processes and erode the value potential of generative AI. Knowledge-intensive businesses in the financial services, healthcare, and insurance sectors face higher-risk environments while striving to adopt AI to transform their processes. To achieve

operational efficiencies and effectiveness with AI, organizations in these sectors must demonstrate a high degree of compliance with regulatory guidelines. All of the risks associated with generative AI, such as model hallucinations, model bias, and failure to comply with regulatory statutes (for example, the EU AI Act) and industry-specific compliance policies (for example, the NAIC Model Bulletin), directly constrain the realization of value from these applications (Gehrmann et al., 2025; Smith, 2025).

Therefore, the combination of sector-specific regulatory pressures and uncertainty concerning these compliance requirements has created a tension, one that requires organizations to use governance practices in response rather than proactively innovate. For instance, if an organization cannot properly govern its data management system to prevent cybersecurity threats, it is vulnerable to regulatory and reputational damages. The pressure to adapt to changing regulations negatively influences an organization's potential to capitalize on its competitive advantage. As well, there are no singular and comprehensive governance solutions, which exacerbates organizational and sector-specific issues (Madanchian & Taherdoost, 2025).

Job Displacement. The adoption of AI at the organization level poses the risk of workforce disruption in the global economy. Research indicates that AI is putting roughly 40% of jobs in advanced countries at risk. Further, the expectation is that machines could take over routine-based work by 2030, negatively affecting employment levels (Walz & Firth-Butterfield, 2021). However, the threat is considered only a short-term risk. In the long run, the creation of new jobs could potentially offset the displacement of existing ones. It may also require organizations to commit time and resources to upskilling programs. If proper reskilling strategies and organizational practices are not developed, the fear that AI can eliminate workers may lead to a decline in the general perception of the technology. Workforce disruption can lead to low trust in the leadership or an overall negative culture, hindering organizational transformations.

Negative Psychological and Social Impacts. While AI can offer great organizational efficiencies, it carries the risk of impacting human beings in detrimental ways. If employees displaced by automated systems have a negative perception about such initiatives or are feeling isolated, their morale can diminish, and thus their overall organizational productivity and engagement (Walz & Firth-Butterfield, 2021). The company's leadership and external stakeholders must implement continuous ethical reviews in the organization to mitigate these risks. Involving employees in governance also facilitates a higher degree of transparency and promotes organizational stability.

While existing mitigating frameworks can assist in managing risk factors, they primarily focus on discrete scenarios. It is unlikely that all potential value versus risk trade-offs can be addressed by current mitigation strategies. This is because the dynamics around generative AI can be quickly changing. Mitigating strategies such as pilot testing, red-teaming models, etc. may offer great insights that address value versus risk trade-offs for individual use cases. However, governance policies must be adaptive and dynamic; tools such as pilot test plans are too static to address the evolving dynamics of AI.

Organizations must be prepared to adopt strategies that dynamically address the value versus risk trade-offs by conducting scenario-based stress tests, establishing continuous feedback loops with stakeholders, and forming stakeholder alliances to proactively address emerging challenges (Smith, 2025; Waisam & Silver, 2025).

To sustainably maximize the value-creating potential of generative AI technologies, organizations will need to embrace a framework of evidence-based executive governance that considers the interaction between ethical/regulatory guidelines and the technology. A framework containing risk management, continuous upskilling, and effective escalation processes can help organizations sustain benefits gained through generative AI technology (Gehrmann et al., 2025; Walz & Firth-Butterfield, 2021). These can address the complex dynamics of value–risk trade-offs, providing organizational structures for generative AI to make a net contribution.

4.4 Governance Gaps Identified in Literature

Corporate AI research's lack of attention to deployment-stage behavioural and operational risks of AI creates a problem for generative AI governance. In fact, research indicates that only 4% of studies focus on high-stakes application areas, such as misinformation and persuasion or financial and healthcare services (Strauss et al., 2025). This imbalance reflects a clear discrepancy between the academic emphasis on pre-deployment risks and the operational risks for organizations during deployment, thereby under-equipping companies with the tools they need to govern AI applications. The point is further emphasized when observing that the limited attention to operational risks of AI and framing of the deployment context as simply testing for bias, in fact, overlooks the operational stage where generative AI is deployed. Consequently, risks that emerge in the operational stage of AI implementation (e.g., unpredictable and problematic behaviours, new biases) are typically not recognized, addressed, or monitored, even though it is at this stage that risks impacting the organization (e.g., business liability, compliance, reputation) are likely to occur and evolve. With generative AI applications rapidly growing in complexity and global impact, the deficiency in post-deployment monitoring capabilities and the lack of tools for this stage create a gap in risk mitigation strategies that executives need to fill through novel approaches.

This is further aggravated by the fact that AI governance frameworks for organizations are limited, particularly for knowledge-intensive organizations. Current frameworks are deficient in their guidance for organizational-level implementation and have yet to address specific generative AI risks, leaving executives unsupported in this domain. One study indicated that only seven proposed organizational-level AI frameworks are available to organizations, and only four of them include an implementation guide (Batool et al., 2025). More troubling is the fact that among the existing AI frameworks, there is a disproportionate amount of focus on fairness and privacy to the detriment of other governance dimensions, such as transparency, accountability, and engagement. While fairness and privacy are certainly important, their dominance creates an imbalance for executives, as they fail to address and provide guidance to structure effective and comprehensive AI

governance, leaving other governance dimensions underrepresented. For instance, the lack of transparency and reporting structures leaves organizations blind to unforeseen errors, unethical outputs, and unanticipated behaviours that the AI models can exhibit. As well, a lack of feedback systems also restricts organizational responsiveness to operational risks, particularly in integrating AI ethical oversight into daily operations. In sum, the over-emphasis on fairness and privacy undermines organizations' scalability and ability to respond and adapt to change and emerging operational and societal AI risks.

Moreover, most sector-specific and national strategies focus largely on “soft” policies, such as guidelines and standards (Gianni et al., 2022). This leads to fragmented governance approaches that cannot scale, creating issues for generative AI governance. Many governments, for example, are pushing for ethical frameworks that provide recommendations for AI governance practices but are not legally enforceable, leading to inconsistency across organizations and industries. As such, a lack of alignment among national and sector strategies complicates generative AI management in organizations for executives dealing with cross-jurisdictional operations, as differing policy regimes and regulatory overlaps necessitate fragmented approaches, increasing organizational complexity and operational expenses.

More worrying is the fact that most national AI governance frameworks are non-binding and non-executable by organizations. Although countries such as Denmark and Singapore are pushing for human-centered approaches to AI development, there are still no frameworks designed to enable organizations to adopt governance measures, address risk scenarios, and ensure compliance at the enterprise level. Ultimately, regulatory fragmentation and lack of legally binding AI governance frameworks impede effective adoption of industry practices and limits organizations' ability to address high-stakes risks. There also exists a fragmentation of the AI governance ecosystem, which impedes the development of a scaled and transferable governance model for generative AI.

To demonstrate this fragmentation, Attard-Frost and Lyons (2024) identify 120 unique governance actors in the Canadian AI policy ecosystem, including 44 organizational actors, highlighting a complex multi-actor environment where coordination and harmonization remain difficult. Unfortunately, this complex and fragmented network remains ineffective at achieving a holistic and coordinated response to generative AI operational risks. Collaborative forms of governance, such as co-governance, consistent standards development, and effective feedback loops, remain relatively sporadic or weak. This aspect has been identified as problematic, since fragmented environments hinder knowledge-sharing, standardization, and effective use of resources. In effect, fragmented policies and governance actors limit the ability of organizations to have a consistent set of standards, creating ambiguity about which practices work best for governance. For instance, due to the lack of international standards for AI governance and the weak level of coordination, executives have an insufficient understanding of emerging global governance approaches. Weak feedback loops and collaborative policymaking also mean that there is a limited organizational ability to respond and adapt to a constantly evolving generative AI risk landscape. As such, executives must adopt

strategies for multi-stakeholder engagement and pursue proactive and systematic efforts to develop a coherent AI ecosystem.The high-stakes and business-critical areas of AI and organizational use remain under-addressed. Academic and corporate literature on AI governance inadequately discuss high-stakes applications and organizational risks.

Of the literature analysed, only 6% in academia and 4% of the corporate literature directly address areas of high liability and higher societal risk, such as medicine, finance, and other knowledge-intensive services (Strauss et al., 2025; Batool et al., 2025). As a result, AI governance initiatives typically leave undiscussed and unaddressed AI risk scenarios for organizations, such as misinformation, financial fraud, and large-scale operational disruptions, which have the highest probability of occurrence. High-liability domain organizations, therefore, are left unequipped, as they lack best-practice guidance that can be adapted to the specific business and ethical contexts in which they operate. Even AI frameworks successfully applied in academic contexts fail to provide practical regulatory guidance for organizations managing the emergent risks of generative AI because existing compliance or ethical guidance does not address them. Organizations are still trying to understand how to best build systems for generative AI applications due to a lack of practical and applicable strategies, even as regulations are advancing and organizations' exposure to AI risk scenarios are continuously and quickly increasing. Organizations in high-stakes and business-critical applications must address this research gap to mitigate the impact and likelihood of AI risks. In conclusion, AI governance gaps identified throughout the literature have important implications for the organizational use of generative AI. These gaps represent practical and operational considerations executives must address to achieve effective AI management.

4.5 Synthesis of Cross-Thematic Patterns

The literature on generative AI governance in knowledge-intensive organizations surfaces several cross-thematic patterns that clarify how value creation and governance structure mutually reinforce - or constrain - each other. Across empirical studies, organizations that realize the highest value from generative AI tend to be those that deliberately implement governance structures and invest materially in workforce adaptation. Despite generative AI’s transformative potential, only a small proportion of organizations have progressed from pilots to enterprise-wide deployment, with some studies reporting rates as low as 8% (Waisam & Silver, 2025; Madanchian & Taherdoost, 2025). This implementation gap is consistently associated with underdeveloped organizational capacity-building and an unclear linkage between governance design and operational or strategic outcomes. A recurring pattern is that sustainable value depends on a combined approach: executive leadership that establishes clear decision rights, proactive upskilling to increase organizational absorptive capacity, and robust data governance to enable reliable and compliant use of generative AI (Madanchian & Taherdoost, 2025). Where these foundations are present, organizations are more likely to achieve efficiency gains and meaningful personalization, and to build resilience as work practices evolve. However, the scalability of these approaches remains uneven. Decentralized

organizations and those with constrained resources face persistent challenges in standardizing practices, addressing disparities in technical capabilities across departments, and maintaining governance consistency as adoption expands.

The pilot-to-scale barrier is also strongly linked to measurement deficits. Several studies indicate that organizations remain stalled because KPIs are poorly defined, governance structures are weak or fragmented, and progress cannot be evaluated in a way that supports scaling decisions (Waisam & Silver, 2025). This pattern suggests that KPI design is not merely a performance-management issue but a governance capability. KPIs that focus narrowly on model performance or isolated use cases can miss operational dependencies, cross-functional impacts, and ecosystem-level risks. As a result, organizations may be unable to demonstrate business value credibly, diagnose failure modes, or justify the controls needed for broader rollouts.

Infrastructure readiness and learning investment emerge as practical differentiators of implementation success. Strategic investment in enabling technologies, data pipelines, secure environments, and structured learning activities is associated with significantly higher success rates (Madanchian & Taherdoost, 2025). In contrast, ad hoc governance - where training and infrastructure are treated as reactive “technical fixes” - tends to produce fragile implementations that struggle under scale and change. This pattern is particularly pronounced for SMEs, where limited budgets and constrained expertise reduce the feasibility of building scalable support structures and continuous training for large language model development and operations. Consequently, collaboration mechanisms - such as shared services, vendor partnerships, consortia, and sector-level initiatives - become governance-relevant strategies for building and sustaining foundational capability.

Another cross-cutting theme is the persistent disconnect between academic proposals and practical implementation. Multiple studies highlight that governance concepts are frequently discussed at a conceptual level without translating into operationally feasible instruments aligned with organizational constraints (Waisam & Silver, 2025; Madanchian & Taherdoost, 2025). This gap implies that parts of the research agenda may be insufficiently grounded in organizational realities, including budget constraints, decentralized structures, legacy systems, and competing business priorities. A promising pathway suggested by the literature is deeper co-production between academics and practitioners to develop usable governance frameworks, templates, and decision models that reflect actual constraints rather than idealized governance designs. Operationalizing selected research insights into accessible tools or platforms may help reduce translation costs and improve adoption fidelity.

Ethical and technical risk controls - particularly fairness, privacy, and explainability - are widely emphasized in the literature, yet implementation outcomes remain limited, and many initiatives terminate at proof-of-concept stage (Waisam & Silver, 2025). This pattern indicates that risk awareness does not reliably convert into embedded controls, especially when governance remains compliance-centric and disconnected from operating models. Where controls are implemented primarily as procedural add-ons, they can become costly, hard to sustain, and vulnerable to being

deprioritized under delivery pressure. More durable outcomes appear to depend on integrating ethical and technical controls into everyday workflows, product governance, procurement, and performance management - so that risk management is treated as a core operational capability rather than a periodic compliance exercise.

Executive leadership is repeatedly presented as a gating factor for governance maturity. Without clear accountability at senior levels, even theoretically robust frameworks tend to remain aspirational and under-enforced (Madanchian & Taherdoost, 2025). This reflects a broader pattern that organizational change is unlikely when governance lacks executive sponsorship, authority, and incentives aligned with responsible deployment. Leadership capability in generative AI governance therefore appears to be an enabling condition: it shapes resource allocation, enforcement of standards, escalation pathways, and the organizational willingness to trade short-term delivery speed for long-term reliability and risk control.

The uncertainty profile of generative AI further supports iterative, adaptive governance rather than linear, one-off compliance programs. AI security risks evolve rapidly due to shifting model behaviours, emerging threats, changing regulatory expectations, and dependencies on external ecosystems (Waisam & Silver, 2025). These conditions reinforce the need for continuous monitoring, periodic re-assessment of risk models, and governance processes that can update controls and thresholds as systems and contexts change. This has implications for organizations operating across jurisdictions, where governance must remain responsive to evolving legal requirements and diverging enforcement practices while maintaining internal consistency.

Ecosystem complexity is another persistent structural constraint. Research mapping AI governance ecosystems with large numbers of actors (e.g., 120 distinct stakeholders) highlights limited coherence and coordination, as well as duplication of efforts and information asymmetries (Attard-Frost & Lyons, 2024). This fragmentation increases transaction costs and makes it harder for organizations outside established networks to access and implement best practices. The literature consequently points to the need for stronger intra-organizational coordination (across risk, legal, IT, data, HR, and business units) and expanded co-governance with external stakeholders, including regulators, vendors, industry bodies, and civil society, to improve alignment and reduce redundant or conflicting initiatives.

Jurisdictional and sectoral variation further complicates governance standardization. Ethical expectations and governance priorities can differ substantially across countries and industries, creating challenges for designing models that are both generalizable and context-sensitive (Madanchian & Taherdoost, 2025; Attard-Frost & Lyons, 2024). This divergence can push organizations toward fragmented governance approaches, where separate frameworks emerge for different contexts, increasing cost and reducing operational efficiency. The literature therefore implicitly supports the value of baseline governance principles that can be consistently applied, complemented by context-specific extensions that address local regulatory requirements and domain-specific ethical risks.

High-stakes domains illustrate the strongest coupling between model risk, data quality, and oversight requirements. In areas such as healthcare, disaster management, and social welfare systems, generative AI introduces heightened governance challenges because consequences are severe and the tolerance for error is low (Sharma, 2023). These deployments often require explicit human oversight and stronger assurance mechanisms to manage model uncertainty, operational dependencies, and risk of harm. At the same time, sustained human-in-the-loop oversight can be resource-intensive, raising practical concerns about scalability and long-term viability in organizations with constrained capacity.

Finally, organizational vulnerabilities repeatedly emerge as the underlying driver of both unrealized value and unmanaged risk. Limited investment in data quality, insufficient skilled personnel, and poorly coordinated operational management models reduce the ability to deploy, scale, and govern generative AI effectively (Soroori Sarabi, 2025). This pattern reinforces that “foundational capability” is not a background condition but a central governance variable: organizations that underinvest in data and skills are structurally less able to implement governance controls, demonstrate value, and manage operational risk. In this context, public–private collaboration - such as capability-building programs, shared datasets, and sector-level standards - can be interpreted as governance-enabling infrastructure that supports broader diffusion and more responsible adoption.

Overall, these cross-thematic patterns indicate that effective generative AI governance in knowledge-intensive organizations depends on integrated, interdependent capabilities spanning executive leadership, workforce transformation, data and technology foundations, operational controls, and ethical risk management. Organizations that align these elements are better positioned to improve productivity and adaptability while reducing exposure to strategic, operational, and ecosystem-level risks.

5. Executive Governance Framework and Discussion

This section synthesizes various governance mechanisms and strategies that facilitate the oversight of generative AI in knowledge-intensive organizations. It outlines how specific, tailor-made strategies and the inclusion of stakeholders in the process can assist in bridging the current gap between policy and practical implementation. This discussion builds upon the groundwork laid so far to provide guidance for organizations aiming to adopt a more structured and responsible approach to generative AI deployment.

5.1 Governance Mechanisms Identified in Literature

The governance of generative AI in knowledge-intensive organizations remains largely untested, particularly the integration of organization-level tailored mechanisms. While there are plenty of available AI governance frameworks presented in academic and practical literature, only seven of these are organizational. In the same vein, Batool et al. (2025) argued that available frameworks

are insufficient in executive oversight during real-world deployments. In the literature, a recurring theme is the overrepresentation of certain principles, like fairness and privacy, at the expense of accountability, transparency, and risk management post-deployment. Batool et al. (2025) also argued that fairness is the most prevalent concept and is overrepresented, which skews governance priority as it leaves out many operational and regulatory challenges. It's important to govern principles in an equitable manner to provide a holistic solution and a robust AI governance regime that is scalable.

Another critical gap in the governance mechanisms identified in the literature is the insufficient documentation of AI lifecycle activities, including provenance tracking, licensing conditions, and updates to training data and model parameters. Such deficiencies can substantially increase legal and compliance risks for executives (Maryala, 2025). Organizations frequently operate across multiple jurisdictions within a rapidly evolving regulatory landscape for artificial intelligence, which requires comprehensive documentation to ensure legal defensibility and auditability. A structured and systematically implemented AI governance framework can therefore support regulatory compliance while strengthening executive oversight, accountability, and internal control mechanisms. Furthermore, the literature highlights a persistent separation between formal policy development and practical implementation, indicating a gap between governance design and operational execution. Madanchian and Taherdoost (2025) state that proposed solutions do not provide practical guidelines beyond high-level policies to support AI governance during and post-deployment. This leaves executives poorly equipped to govern critical business risks as AI models move from testing to deployment.

For an organization to extract all the potential benefits of AI, it must scale past pilot phases to production. Only 8% of organizations can achieve this milestone (Smith, 2025). Poor governance is often a primary culprit when solutions can't scale in knowledge-intensive organizations. Bolden et al. (2024) assert that one of the major failures of AI is the inability to scale it in production. There are a few common reasons why governance fails to scale, namely undefined accountability, poor risk frameworks, and a deficiency in AI expertise. Ultimately, poor governance results in stalled initiatives, which can lead to value stagnation and, potentially, decay in the organization.

Poor ecosystem-level governance presents significant challenges for organizations that aim to utilize and govern generative AI, especially those with international reach. Attard-Frost and Lyons (2024) map Canada’s AI policy network and identify 120 unique actors, including 44 organizational actors, within the AI governance and policy space. Importantly, these figures represent analytical categories derived from the policy-network analysis rather than a sample size of interview participants or organizations. While several organizations engage in AI governance, the study suggests a lack of harmonization and standardization for AI governance, development, and implementation. Additionally, they discovered that participants of Canada's ecosystem lacked adequate collaboration or feedback. For example, co-governance activities and feedback loops to the governance actors were rare. All of this can hinder an organization's decision-making because

there is a lack of consensus within the AI ecosystem. For example, Attard-Frost and Lyons (2024) mention that a lack of co-governance activities and feedback loops creates fragmented and inefficient governance systems. Without feedback loops and multi-actor collaboration, it may prove challenging for organizations, especially those operating in multiple jurisdictions, to learn from each other to collectively improve their own governance mechanisms. A common problem for knowledge-intensive organizations that use generative AI is that they don't have enough resources. This problem includes both internal organizational resources such as staff and expertise as well as external ecosystem resources (data, standards, and frameworks). This phenomenon is primarily seen through knowledge and policy instruments for governance. Attard-Frost and Lyons (2024) assert that while there are multiple actors and instruments within AI governance, there is significant centralization in the contribution of both resources and financial instruments. Certain actors in the ecosystem possess more of these resources than others. Canada's ecosystem is dominated by government resources and policy instruments that drive governance mechanisms (Attard-Frost and Lyons, 2024). The authors suggest that the AI ecosystem may distribute data and financial resources less than others. Such an arrangement creates a large gap in ecosystem participants for smaller firms or those who are lacking in specific resources or power. In addition to a small number of co-governance activities, there were no instances of the use of national-level governance mechanisms discovered in Attard-Frost and Lyons's (2024) findings. The lack of high-level, overarching policy mechanisms and co-governance can increase fragmentation and decrease the ease of scalability or organizational benefits. National-level governance can improve the scalability of organizations as well as ensure that regulations and mechanisms are coherent. Governance guidelines can come from various organizations (Attard-Frost & Lyons, 2024), and multiple organizations may propose similar governance mechanisms. International knowledge-intensive organizations using generative AI can employ various ways and strategies for governing these algorithms. While international and global governance standards and principles exist, some authors posit that national, rather than multinational, frameworks could be more useful for organizations (Attard-Frost & Lyons, 2024; Poon Affat, 2025).

As part of their policy-network evidence base, the authors catalogue 650 distinct guidance documents produced or referenced by governance actors, underscoring the breadth - but also the potential incoherence - of available governance guidance (Attard-Frost & Lyons, 2024). This lack of harmonization can confuse AI governance within an organization and potentially lead to misaligned and ineffective governance (Poon Affat, 2025). Poon Affat (2025) also said that there is no clear flow of information and responsibility in global AI ecosystems. This is a problem because AI and governance standards have not yet been adopted and unified around the world.

There has been growing interest by governments around the world toward ISO standards, and countries and organizations are pushing organizations toward adopting standardized protocols to help govern AI. Yet Poon Affat (2025) acknowledges that ISO standards are fragmented and non-hierarchical, which creates gaps in effective control. Furthermore, frameworks, such as the NAIC

Model Bulletin, have not been widely adopted in the insurance industry. To be well governed, an organization will have to follow best practices in a multitude of categories and use international and industry-wide standards for its own unique governance system. This is a key reason why executive management needs to understand the organization's risk architecture.

It is necessary for organizations to acknowledge the diversity and independence of the board as another key component to effectively govern AI and maximize performance. Boards can play a crucial role in risk management and AI governance when they properly control, monitor, and direct executives who use, create, and deploy AI systems.

A board's monitoring quality is better when the number of women serving on the board increases (Coulson-Thomas, 2023). Additionally, Coulson-Thomas (2023) argues that a balanced gender makeup on the board leads to greater emphasis on corporate social responsibility. Moreover, the presence of at least three independent directors can have the same impact, as a board will perform better in terms of controlling executive actions and behaviours and will monitor managers more effectively (Coulson-Thomas, 2023). According to Coulson-Thomas (2023), such an arrangement also improves decision quality on the board by lowering the risk of groupthink and challenged decisions that ultimately provide poor oversight for AI. For organizations that utilize a large volume of personal data, transparency in organizations can have a significant impact on whether or not a data breach is avoided and will provide increased ethical awareness within the organization. These behaviours lead to more effective governance and risk management (Coulson-Thomas, 2023).

Coulson-Thomas (2023) indicates that ethical awareness depends on individual beliefs in right or wrong. Additionally, organizational trust, when higher, is associated with organizational openness, adaptability, and a greater ability to improve existing governance processes, creating an environment more open to AI implementation and AI integration.

The success of organizational AI and integration also hinges on how charismatic the CEO is, which can increase employees' trust in the organization (Coulson-Thomas, 2023). When a company and leaders lead with charisma, that tends to evoke change management as employees feel secure and open to new policies and procedures (Coulson-Thomas, 2023). If trust is cultivated and executives are willing to share the risks and uncertainties with their employees, organizations become more open and can successfully adapt and transform business practices when integrating generative AI (Coulson-Thomas, 2023).

For the previous mechanism listed as executive leadership, leadership diversity can also lead to organizational resilience when faced with difficult situations like change management or integrating an uncertain technology. Leaders can enhance the resilience of a knowledge-intensive organization by cultivating trust and improving collaboration with and between the organization, stakeholders, and employees. Despite all the identified principles in the literature that are required for good AI governance, Madanchian and Taherdoost (2025) mention that there are fewer guidelines and practical examples available for high-level ethical guidelines compared to

governance concepts. While ethical guidelines can be a beneficial framework to start from, they must be integrated and operationalized into each individual organization in ways that work specifically for that company and provide value (Madanchian & Taherdoost, 2025). Additionally, organizations should consider all principles within their internal system and not prioritize the easier-to-adopt ethical principles. The lack of implementation guidelines from academia can leave management in high-risk organizations at a disadvantage.

5.2 Integrative Conceptual Derivation

The harmonization of regulatory standards on generative AI within knowledge-intensive organizations is difficult due to different ethical standards and regulatory requirements across national and sector contexts. For example, based on an analysis of 200 regulations from Madanchian and Taherdoost (2025), Poon Affat (2025), and Perry and Uuk (2019), regulations varied in their ethics and operations considerations. The result indicates that implementing a single effective governance model is improbable since organizations need to establish governance models that are context-specific and acceptable at the local levels in which they are operating. This requirement is especially pertinent for multinational organizations that face EU transparency standards versus U.S. sector flexibility. This suggests the need for strong executive leadership to manage such diverse regulations, such as in the NAIC Model Bulletin for the insurance sector in the U.S., to ensure context-appropriate governance rather than strictly one-size-fits-all compliance (Poon Affat, 2025). However, whether context-specific models enhance or inhibit overall operational efficiency remains a critical inquiry since such outcomes will depend on their cost in each context. For example, in an organization operating in 20 different contexts that have vastly different ethical and regulatory conditions, a context-specific AI governance model can be quite efficient and effective by maximizing local acceptability and compliance while being highly inefficient due to operating 20 separate governance models. If, instead, these 20 governance models could be synthesized, the overall effectiveness of the AI governance model would be substantially increased; however, the local effectiveness of each model would be lower due to limited contextualization of the context-independent part of the synthesized AI governance model.

The relationship between the intensity of AI usage and performance on macro-level indicators suggests the problem is at the organizational level, given the overall macroeconomic value generated versus the low number of generative AI projects operationalized. Madanchian and Taherdoost (2025) estimate that a 10% increase in AI use in a high-income, knowledge-intensive sector could increase sector GDP by 0.3%. However, many AI initiatives still stall at experimentation rather than sustained production (Dasgupta & Wendler, 2019). This scenario represents a failed executive commitment for governing AI beyond adoption since AI-based performance and risk measurement systems are largely not at the organizational level. It is essential to implement appropriate performance measures, such as workforce impact measurements, in addition to technology-based AI performance measures. Researchers should emphasize post-adoption risk

management, compliance, and other governance measures, in addition to measuring these performance metrics. However, while these performance measurements are necessary, in practice, their feasibility must be considered, and implementation steps should be included. For instance, most small and medium-sized enterprises will lack the capacity and resources to measure and implement the governance constructs from the frameworks discussed previously, and thus, the governance strategies are practically infeasible. Such an outcome calls for collaborations between organizations, governance solutions, and external resources that can make AI governance feasible and available for more organizations.

Fractured industry standards, slow hierarchical processes, and executive execution of these principles may hinder the integration of organizational, regulatory, and ecosystem governance mechanisms. Although frameworks, such as ISO 42001 and the NAIC Model Bulletin, are foundational in their focus on compliance, risk, and fairness, the operational scalability, applicability, and overall effectiveness of these standards are challenged by fragmentation across industries. The abundance of industry standards also requires executives to determine the appropriate methods and standards for the context of their organization, limiting scalability by increasing operational and regulatory burdens. Although governmental regulators prioritize organizations that adhere to industry-recognized standards, it seems that these are often used to avoid regulatory penalties and prevent innovative activities. These governance models are effective in that their adherence is a competitive advantage through reduced regulatory risk, but the strategic response to government regulations is to avoid high-value creative endeavours, since their existence opens the door for missteps that can be penalized by these standard-setting authorities. Such an approach implies that organizations may need to develop standards-appropriate models that can rapidly scale and adjust, but at a cost in terms of scalability across the organization. Moreover, having numerous fragmented standards can become burdensome since the interpretation and integration of these many can cause executives to divert their attention to governance compliance that could otherwise be used for more creative endeavours. Research on the ecosystem fragmentation of AI governance identified the root cause as multiple governance actors with weak linkages. In the Canadian context, over 120 unique AI governance actors were found, and such fragmentation resulted in standards being inconsistent and poorly applied (Attard-Frost & Lyons, 2024). To ensure regulatory compliance and scalability of governance across the AI ecosystem, more mechanisms need to be in place to ensure consistent governance models. Furthermore, organizations that have robust internal governance strategies in place can benefit organizations that are unable to adopt these practices within the organization by enhancing governance at the ecosystem level through feedback and improvement mechanisms. The inefficiency and lack of scalability inherent within a multi-actor regulatory ecosystem indicate the necessity of improved collaboration, coordination, feedback, and joint policymaking between these actors and within organizations. A challenge here is how to make these concepts actionable and implementable, since simply suggesting increased communication and coordination does not solve the issue if, in practice, stakeholders have differing priorities or if the transaction costs are too high

to make collaboration sustainable. These costs include monetary or time-related factors, as well as non-monetary factors, such as an executive's time being occupied by these cross-stakeholder governance processes, which may result in an opportunity cost in their time being spent on something else.

Most existing AI risk management theories are focused heavily on technical aspects of fairness and privacy, which largely avoid areas of broader governance issues regarding risk in general, compliance, and organization or operational liability, as well as fairness in high-stakes deployments (Batool et al., 2025). Corporate-led studies consider sector governance in knowledge-intensive sectors (e.g., healthcare or finance) only 4% of the time, compared to 6% in research-led studies (Batool et al., 2025; Strauss et al., 2025; Madanchian & Taherdoost, 2025). This gap suggests a need to further incorporate risk monitoring and post-deployment risk management into the governance models. Moreover, there is also a need for continuous evaluation of new risks during and after generative AI implementations and deployments. Given that most theoretical AI risk management research occurs at a technical level, most recommendations for governing the implementation and deployment of AI are regarding the technology; however, this suggests the need for incorporating strategic governance models that include risk identification and mitigation from a holistic viewpoint, not just a technical risk viewpoint. The challenge lies in developing operationalized solutions that can be practically implemented, taking into account organizational resources and scale.

In summary, the integrative conceptual derivation of governance models and strategies for generative AI implies that several major gaps exist within knowledge-intensive organizations that must be addressed, which largely relate to fragmentation across all governance factors and levels. The fragmented issues included regulatory standards, insufficient operational and strategic scalability, and a lack of post-deployment performance measurement and risk management strategies. These fragmented standards need to be streamlined across the entire organization so generative AI adoption can effectively be accelerated and standardized. Scaling organizational standards may require additional efforts to improve scalability between and across firms. To effectively measure performance in high-stakes environments, additional post-deployment measurement and risk mitigation must occur that are embedded into governance, which allows for a quicker and more agile responsiveness to dynamic environments and organizational demands.

5.3 Development of an Executive Governance Framework

Drawing on the gaps identified in *Section 5.1* and the integrative tensions synthesized in *Section 5.2*, the framework is designed to operationalize five recurring governance requirements: (1) explicit accountability and decision rights to address diffusion of responsibility; (2) lifecycle documentation and auditable evidence to mitigate legal and regulatory exposure; (3) post-deployment monitoring and incident governance to close the policy–action gap; (4) proportional controls through risk-tiering to balance innovation speed with enterprise assurance; and (5) capability enablement to reduce

shadow use and behavioural risk during scaling. These design requirements directly reflect the literature's emphasis on insufficient executive oversight, weak scalability into production, and fragmented governance instruments.

This section (5.3) presents an executive-operational governance framework for generative AI in knowledge-intensive organizations that translates the SLR synthesis into a practical, implementable governance model. It positions GenAI governance as a portfolio-level, risk-tiered executive capability that aligns use-case value mechanisms with multidimensional risk mitigation through clear decision rights, lifecycle integration, and measurable oversight. The framework operationalizes the literature-identified tensions between speed and control, innovation and accountability, and decentralized experimentation and enterprise assurance by establishing (i) a small set of interdependent governance domains, (ii) a lifecycle with embedded executive stage gates, (iii) a tiered control logic that enables proportionality, (iv) explicit accountability for risk acceptance and operational execution, and (v) a KPI architecture that integrates value delivery with control effectiveness and risk exposure. The result is a governance design that executives can deploy as an operating model for disciplined GenAI value creation under bounded and auditable risk.

Existing AI governance guidance, as synthesized in the SLR, is frequently insufficiently operational for executive decision-making in knowledge-intensive contexts because it tends to be either principle-heavy (articulating desired ethical or responsible outcomes without specifying executable mechanisms) or technically oriented (focusing on model development practices rather than executive oversight designs). The SLR further indicates a persistent principles-to-practice gap in which governance artifacts exist but do not reliably translate into enforceable decision rights, consistent deployment-stage controls, or measurable accountability. In such environments, executives face a dual exposure: they are accountable for value realization from GenAI adoption while simultaneously bearing enterprise risk for misuse, regulatory breach, reputational damage, and operational failure. Without structural clarity, organizations default to fragmented governance - either permissive experimentation that scales risk or excessive restriction that suppresses value mechanisms identified in the SLR (such as productivity gains, knowledge augmentation, accelerated content generation, and improved decision support).

The Executive GenAI Governance Framework is derived conceptually from the SLR's synthesis of value–risk tensions. It treats governance as an executive design problem in which value alignment and risk control must be structurally integrated rather than managed as sequential or competing agendas. Accordingly, the framework addresses executive-level governance of GenAI use and deployment—portfolio shaping, risk appetite setting, control allocation, accountability design, and oversight - rather than technical AI model design. Technical safeguards remain necessary, but they are positioned as implementable controls within an executive system of decision-making, assurance, and continuous improvement.

The framework is organized into five interdependent governance domains that together convert literature-identified governance mechanisms into an executive-operational architecture. The first

domain, *Strategic and Portfolio Steering*, defines the organisation's GenAI intent, portfolio boundaries, and value thesis by translating enterprise strategy into an "approved problem space" for GenAI. This domain ensures GenAI is adopted for defensible, priority outcomes rather than opportunistic tool use, and it establishes portfolio-level value criteria (for example, knowledge work acceleration, service quality uplift, or improved decision consistency) alongside explicit boundary conditions, including prohibited or restricted contexts.

The second domain, *Risk Appetite and Policy Control*, embeds risk appetite into actionable policy by defining what levels of risk exposure the executive is willing to accept for different categories of GenAI use. Risk appetite is operationalized through a tiered classification logic, minimum control requirements per tier, and clear rules for risk acceptance. In practice, this domain anchors proportionality: controls are not uniform but scaled to the multidimensional risk profile and exposure of each use case.

The third domain, *Operational Enablement and Capability*, ensures that adoption is feasible and disciplined by establishing the organizational capability required to execute governance consistently. This includes workforce enablement, minimum competency expectations, approved tooling pathways, and change leadership that addresses behavioural and cultural risks highlighted in the SLR (such as shadow use, over-reliance, and inappropriate automation of judgement). This domain makes governance implementable by ensuring that policy and controls are usable within the rhythms of knowledge work.

The fourth area, *Assurance, Monitoring, and Incident Governance*, turns governance promises into proof that can be checked and effective management. It defines monitoring expectations, KPIs and reporting structures, assurance activities, and incident handling workflows. Critically, this domain integrates value and risk monitoring so executives can manage GenAI as a performance-and-exposure portfolio rather than as isolated technical implementations.

The fifth domain, *Accountability and Decision Rights*, defines who decides, who executes, and who accepts residual risk. It specifies executive ownership for the GenAI portfolio, assigns accountable business ownership for individual use cases, and formalizes the roles of compliance, IT/security, and assurance functions. This domain addresses the SLR-identified governance gap in executive-operational design by ensuring that decision rights are not implicit or distributed by default but explicitly allocated.

These domains are mutually reinforcing. Strategic choices determine which use cases enter the pipeline; risk appetite governs permissible exposure and control depth; enablement ensures adoption occurs through approved pathways rather than informal workarounds; assurance exposes whether value is being realized and whether risks are controlled; and decision rights determine how trade-offs are resolved and escalated.

The framework is operationalized through a governance lifecycle that treats each GenAI use case as a managed asset from identification to decommissioning, with executive oversight embedded through risk-tiered stage gates. The lifecycle begins with use-case identification, where business

owners articulate the intended value mechanism, target users, expected scale, and decision criticality. This stage requires explicit framing of how GenAI will be used (assistive, advisory, or automating), because the SLR suggests that risk exposure materially shifts when GenAI output influences consequential decisions or external stakeholders.

Next, risk classification applies the tiered logic by evaluating multidimensional risk factors and exposure drivers, including the sensitivity of inputs and outputs, externality (internal-only versus customer-facing), degree of autonomy, regulatory and contractual constraints, and potential harm severity. Tiering functions as the governance “router” that determines the depth of required controls and the level of executive involvement. Control allocation then translates the risk tier into a tailored control set, combining minimum mandatory controls with context-specific additions. This step integrates value and risk by ensuring that controls maintain the intended value of the use case while limiting risk; if the controls necessary for ensuring the safety of a use case undermine its business rationale, the lifecycle is structured to either halt or redesign the use case instead of continuing without consideration.

Executive approval occurs at a tier-appropriate gate. Low-risk use cases may be approved through delegated authority under defined policy constraints, whereas higher-risk deployments require executive sponsor sign-off and risk acceptance, and the highest tiers require committee-level or executive committee endorsement. Deployment oversight ensures that the approved control set is implemented before scaling, including evidence of control completion and readiness. Monitoring then operates as a continuous governance function rather than a post hoc audit: performance indicators track value realization and adoption behaviour, while risk indicators track incidents, near misses, policy breaches, and control effectiveness.

Incident handling is integrated as a formal pathway with predefined thresholds for escalation, containment, and communication. This enables executives to manage GenAI incidents with speed and legitimacy, addressing the SLR’s emphasis on accountability and trust. Review and adaptation occur on a cadence aligned to risk tier and organizational learning needs; monitoring signals and incident learnings are explicitly used to refine tier criteria, update controls, and adjust portfolio decisions. Finally, decommissioning is treated as a governed outcome, not a failure, ensuring that use cases are retired when value decays, risk exposure increases, or organizational priorities shift. Across the lifecycle, governance is designed to operationalize value–risk management by making trade-offs explicit, staged, evidence-based, and traceable to accountable decision rights.

Risk-tiering operationalizes proportional governance by aligning the intensity of controls and executive attention to the risk profile and exposure of GenAI use cases. Tiering acknowledges that GenAI risks are not uniform; they fluctuate based on context, scale, and the ramifications of error or misuse. The framework therefore classifies use cases into tiers using a consistent logic that considers both likelihood and impact across multiple risk dimensions (for example, privacy, security, legal/IP, reliability, bias, reputational harm, and operational disruption), while also accounting for exposure amplifiers such as external deployment, automation of judgement, and integration into

critical workflows. Tiering is important because it helps executives create value by allowing low-risk innovation to move quickly while making sure that high-risk use cases get more attention, stronger controls, and clear acceptance of risk.

The table below summarizes the tiered control logic and its executive implications.

Tier	Risk Profile	Required Controls	Executive Involvement	Monitoring Intensity
Tier 1	Internal, low-consequence use; non-sensitive inputs/outputs; assistive use with human judgement retained	Approved tools only; data classification checks; user guidance and training; basic logging; human review of outputs before use	Delegated approval under policy; periodic portfolio visibility	Baseline, trend-based (e.g., quarterly)
Tier 2	Customer-facing or business-material use; limited sensitive data; advisory outputs influencing decisions but not determinative	Tier 1 controls plus documented use-case assessment; control evidence; content quality checks; privacy and security review; model / vendor due diligence proportionate to exposure	Executive sponsor visibility; steering committee approval for scaling	Regular (e.g., monthly) with exception reporting
Tier 3	Regulated, high-impact, or sensitive-data use; integration into core processes; heightened legal, reputational, or harm potential	Tier 2 controls plus formal risk assessment and sign-off; enhanced testing and assurance; stricter access control; audit-ready documentation; incident playbooks; change control	Executive sponsor approval and explicit risk acceptance; committee endorsement prior to launch	High frequency (e.g., weekly) and event-driven
Tier 4	High-stakes, severe-harm potential, or prohibited-by-policy contexts; autonomy or	Presumptive prohibition or exceptional-case governance;	Executive committee and/or board approval only;	Continuous, real-time signals and immediate escalation

	externality creates unacceptable exposure without exceptional justification	independent assurance; executive committee/board gate; stringent containment and kill-switch capability; continuous monitoring requirements	documented rationale and accountability	

Table 3: Decision Rights Allocation Matrix

Interpreting the table, the intent is not to create bureaucratic friction but to embed proportionality and defensibility. Tier 1 accelerates value capture for low-risk productivity use, while Tier 3 and Tier 4 ensure that higher-risk initiatives are governed as enterprise commitments with explicit accountability, auditable controls, and heightened oversight. The tiering system also enables a disciplined “stop or redesign” decision when a use case cannot be made acceptable within the organisation’s risk appetite without undermining its value proposition.

The framework allocates accountability to ensure that GenAI outcomes are owned, risks are accepted transparently, and operational execution is enforceable. Executive ownership at the portfolio level is designated to an executive sponsor or accountable group, ensuring alignment with enterprise strategy and accountability for balancing value delivery with risk. Risk acceptance is treated as an explicit decision right rather than an implied consequence of deployment; for higher tiers, the executive sponsor (and, where required, an executive committee) formally accepts residual risk after confirming that mandated controls are implemented and evidence is available.

Compliance and legal functions operate as control authorities and advisors on obligations, but not as default owners of business value or operational delivery. Their role is to translate external requirements and internal policy into executable constraints, validate that risk assessments and approvals meet governance standards, and support escalation when non-compliance or unacceptable risk exposure is detected. Business owners are accountable for defining use-case value, ensuring appropriate use within business processes, and maintaining operational controls in day-to-day practice, particularly around human oversight and decision use. IT and security functions are accountable for enabling approved technical pathways (access, identity, tooling, integration, and security controls) and for monitoring technical risk signals, while assurance (internal audit or a comparable function) validates whether governance operates as designed and whether control evidence is reliable.

Escalation pathways are defined to prevent ambiguity in time-critical situations. Operational risks or breaches escalate from use-case owners to the GenAI governance hub, then to the executive sponsor or committee depending on tier and severity, with board-level visibility for material incidents or Tier 4 cases. Escalation is triggered by predefined thresholds (for example, severe incidents,

repeated policy breaches, or sustained value underperformance), ensuring executives receive timely signals rather than retrospective reporting.

A concise allocation of decision rights is summarized below.

Legend (RACI): A = Accountable (final decision/ownership), **R = Responsible** (executes the work), **C = Consulted** (provides input), **I = Informed** (kept updated).

Governance Decision	Executive Sponsor	GenAI Steering Committee	Business Use-case Owner	Compliance / Legal	IT/Security
Portfolio prioritisation and boundaries	A	R	R	C	C
Risk tier assignment and approval pathway	A	R	R	C	C
Control tailoring and implementation sign-off	A (Tier 3–4)	R	R	C	R
Residual risk acceptance	A	C	C	C	C
Incident severity classification and external notification	A (material)	R	R	R	R

Table 4: Decision Rights & Accountability Matrix (RACI) for GenAI Governance Decisions

In this table, accountability (A) anchors executive ownership of trade-offs and risk acceptance, responsibility (R) ensures operational execution and governance processing occurs reliably, and consultation (C) prevents siloed decisions by integrating relevant expertise without dissolving ownership.

Monitoring must integrate value and risk because executive governance is fundamentally a performance-and-exposure responsibility: GenAI that delivers efficiency without controls creates latent liability, while stringent controls without value realization erode strategic legitimacy and investment discipline. A KPI architecture therefore functions as an executive instrument panel that consolidates leading and lagging indicators across value delivery, risk exposure, control effectiveness, and capability maturity. This enables executives to manage a GenAI portfolio with the

same discipline applied to other enterprise assets—tracking whether expected benefits are realized, whether risk appetite remains respected, and whether governance mechanisms are functioning.

The KPI system is set up to allow for tiered oversight. High-risk tiers require more frequent review and tighter escalation thresholds, while low-risk tiers can be monitored through sampled assurance and trend signals. Escalation thresholds operationalize accountability by converting KPI deviations into decision triggers: sustained under-delivery of value prompts reassessment of use-case viability; control failures and rising incidents prompt containment, remediation, or suspension; and behavioural indicators (such as shadow tool usage) prompt targeted enablement and enforcement. Continuous improvement loops are embedded by making KPI outcomes and incident learnings direct inputs into governance adjustment, including refinement of tier criteria, control catalogues, enablement content, and portfolio prioritization.

The table below provides a compact KPI architecture suitable for executive oversight.

Governance Objective	KPI Category	Example KPI	Review Cadence	Escalation Trigger
Realise risk-adjusted value	Value realisation	Benefit delivery vs. business case; cycle-time reduction in target knowledge tasks	Monthly (portfolio), quarterly (exec)	Material variance vs plan over two cycles; value < control cost threshold
Maintain acceptable risk exposure	Risk events and near misses	Number and severity of incidents; privacy or security events linked to GenAI use	Weekly (Tier 3–4), monthly (Tier 1–2)	Any severe incident; upward trend in high-severity events
Ensure control effectiveness	Assurance	Control completion rate; open audit exceptions; overdue risk assessments	Monthly	Control compliance below threshold; repeated exceptions in same control area
Sustain compliant adoption behaviour	Workforce and usage	Percentage trained; approved-tool usage rate vs unapproved / shadow use	Monthly	Shadow use exceeds threshold; training completion below target in impacted units
Improve responsiveness and learning	Operational resilience	Mean time to detect / respond; time to remediate;	Monthly; immediate for major incidents	Response time breaches; recurrence indicates ineffective remediation

		recurrence rate post-fix		

Table 5: Executive KPI Architecture for Generative AI Governance

Interpreting this table, the key design choice is to treat governance success as a balanced outcome. Executives should expect that KPIs will sometimes trade off (for example, faster adoption may temporarily increase incidents), but the architecture ensures those tensions are visible, discussed, and managed through deliberate decisions rather than unmanaged drift.

Implementation is most effective when staged as an executive change program rather than a one-off policy release. In the Foundation phase, executives establish the governance backbone: define strategic intent and risk appetite for GenAI, create the steering structure and decision rights, stand up a use-case registry, and operationalize initial tier criteria with minimum viable controls and approval gates. This phase prioritizes clarity, legitimacy, and enforceability over completeness, ensuring that early adoption occurs through governed pathways.

In the institutionalization phase, governance is embedded into core organizational systems, so it scales. Tiering and approvals are integrated into investment and delivery processes; procurement and vendor management align to GenAI risk expectations; monitoring and KPI reporting become a routine executive cadence; and enablement is expanded to address behavioural and cultural risks.

During the optimization phase, organizations grow and learn at the same time. Controls are improved based on incidents and performance, monitoring is automated as much as possible, assurance is more focused on risk, and the portfolio is actively rebalanced to focus on risk-adjusted value. Across all phases, the emphasis remains on executive ownership of trade-offs, evidence-based oversight, and continuous adaptation.

The framework is conceptually derived from the SLR synthesis and is therefore not presented as empirically validated in specific organizational settings. Its operational mechanisms require contextual adaptation to organizational size, risk profile, regulatory environment, and maturity, particularly in defining tier thresholds, approval gates, and KPI targets. GenAI governance also remains dynamic: both value mechanisms and risk landscapes evolve as models, tooling ecosystems, and regulatory expectations change. The framework should therefore be treated as an adaptive executive operating model that demands ongoing review and recalibration rather than a static compliance artifact.

5.4 KPI and Monitoring Architecture

To effectively govern generative AI at the executive level, there needs to be a monitoring system that sees GenAI as a performance-and-exposure portfolio instead of a group of separate technical deployments. Consistent with the framework in *Section 5.3*, KPI design must simultaneously (i) evidence value realization, (ii) detect and manage risk exposure, and (iii) demonstrate control effectiveness and accountability in auditable form. Standard AI performance metrics are therefore

insufficient on their own; instead, KPIs must be anchored in business objectives, risk factors, and compliance requirements, while remaining operationally actionable for scaling decisions (Milkau, 2018; Kaul et al., 2025).

A multi-level KPI architecture should begin with value-aligned outcome metrics that reflect the intended value mechanism of each use case (e.g., cycle-time reduction in a knowledge task, quality uplift in service responses, or cost reduction through automation). Such KPIs move beyond generic productivity tracking and capture whether GenAI produces new or amplified outcomes rather than merely increasing throughput (Milkau, 2018; Kaul et al., 2025). Operational examples include reductions in documentation effort or certification preparation cost where GenAI is deployed to augment routine knowledge work. However, just looking at value metrics isn't enough for proper oversight, because even successful cases can still have hidden risks like privacy issues, weak tracking of data sources, or over-dependence. For this reason, portfolio reporting should connect value outcomes to the cost and strength of controls, enabling executives to evaluate risk-adjusted value rather than gross benefit.

Second, the monitoring architecture should embed accountability-oriented metrics that translate governance responsibilities into measurable expectations. The Australian Institute of Company Directors and the UTS Human Technology Institute (2024) emphasize that accountability improves when AI decision-making roles are mapped explicitly to board and management responsibilities and then monitored through indicators such as board engagement cadence, decision traceability, and escalation effectiveness. In governance practice, accountability KPIs should measure both capability and behaviour (e.g., role-based training completion, adoption of approved tools, compliance with review requirements, and timeliness of risk acceptance decisions). When companies want to improve their governance by changing the makeup of their boards and how they are overseen, diversity and independence can be seen as conditions that make governance easier and support better monitoring and challenge, even if they need more resources to put into place (Coulson-Thomas, 2023).

Third, because GenAI risk is dynamic and often emergent during real-world use, the KPI set must include adaptive risk indicators that function as executive early-warning signals rather than retrospective audit artifacts. Kaul et al. (2025) recommend operational indicators such as human-in-the-loop intervention rates for high-risk use, remediation cycle times for detected issues (e.g., bias or harmful output patterns), and monitoring of MLOps or usage dashboards. Governance should also keep an eye on the number and severity of incidents, near-miss reports, policy violations, and the number of times something happens again after it has been fixed. These indicators support faster containment and learning and are especially critical when GenAI outputs influence consequential decisions or external stakeholders, where exposure escalates rapidly. Freeman et al. (2025) further argue that KPI systems must avoid becoming over-engineered and obsolete; therefore, indicator sets should be periodically reviewed and streamlined, prioritizing measurable signals that reliably trigger operational decisions and escalation.

Fourth, the monitoring architecture must provide compliance and audit readiness metrics that evidence alignment with external frameworks and shifting regulatory expectations. The dominance of ISO / IEC 42001 and the broader fragmentation of standards create a governance requirement to monitor not only technical conformance but also the organisation's ability to demonstrate compliance under audit conditions (Poon Affat, 2025). Practical compliance KPIs include the completion rate of mandatory risk assessments, evidence availability for lifecycle documentation (e.g., provenance, licensing, and change logs), closure time for audit exceptions, and the number and severity of compliance-related incidents. However, an exclusively compliance-driven KPI regime risks incentivizing a defensive posture that deprioritizes value creation and context-specific harm management. Executive monitoring therefore needs to balance external alignment with internal fit-for-purpose governance-ensuring that compliance metrics reinforce, rather than replace, substantive risk control and value delivery (Poon Affat, 2025).

Finally, continuous improvement should be built into the KPI architecture through feedback loops that operationalize governance learning. Freeman et al. (2025) emphasize that the implementation, assessment, and iterative revision of governance frameworks fosters trust and robustness. In practical terms, KPI outcomes, incidents, and audit findings should feed back into governance calibration: adjusting thresholds, refining review cadences, updating enablement content, and modifying tier criteria or control requirements where evidence indicates either over-control (value suppression) or under-control (excess exposure). Pilot-based validation - especially in higher-risk functions - can be used to test whether indicators are both reliable and decision-relevant before scaling organization-wide. While iterative improvement is resource-intensive, it is necessary to maintain alignment between monitoring signals, evolving organizational use patterns, and the changing GenAI risk landscape.

In summary, a KPI and monitoring architecture that supports executive governance of generative AI must integrate risk-adjusted value measurement, accountability metrics, adaptive risk indicators, and compliance evidence with structured feedback loops for continuous recalibration. This integrated design ensures that GenAI scaling decisions remain defensible, auditable, and aligned with organizational strategy while maintaining responsiveness to emergent risks (Kaul et al., 2025; Freeman et al., 2025; Australian Institute of Company Directors & UTS Human Technology Institute, 2024).

5.5 Managerial Implications and Theoretical Contribution

Managerial accountability and the development of skills help ensure effective generative AI governance by bridging gaps in organizational readiness for emerging technologies and new rules. By assigning clear decision-making responsibilities at both management and board levels, the Australian Institute of Company Directors and UTS Human Technology Institute (2024) highlight the importance of accountability in generative AI compliance and traceability. However, this recommendation raises the question of how smaller organizations with limited management capacity

can effectively implement these processes. Assigning roles or committees to ensure accountability is recommended but needs to be balanced with the added resource pressures imposed on smaller companies (Australian Institute of Company Directors & UTS Human Technology Institute, 2024). Furthermore, executives need to be continually trained in generative AI-related skills to adapt to the fast-changing technology environment. Through providing educational programs and training opportunities, executives will be able to improve AI governance culture to ensure continuous improvements, according to the Australian Institute of Company Directors and UTS Human Technology Institute (2024). However, providing these initiatives may further constrain small organizations with tight budgets. Collaborative initiatives or government support may be required to improve AI governance capacity development in resource-constrained organizations.
Performance metrics and monitoring systems for generative AI governance enable sustainable organizational growth and effective executive oversight. While KPIs have been used to track AI performance, most of them have remained generic and failed to meet the needs of generative AI governance (Kaul et al., 2025). A multi-level KPI architecture measures performance beyond operational efficiency and allows organizations to track levels of compliance, workforce development, and value delivery, using an international framework like ISO / IEC 42001 (Kaul et al., 2025). By tracking the levels of these KPIs, executive governance can monitor for continual improvements in generative AI governance, allowing them to align AI strategies to business strategies. KPI architecture design for generative AI must address the tension between effective compliance management and innovation. Having defined KPIs related to the boardroom, such as ‘frequency and outcome of board-level AI engagements’ and ‘the number of employees whose compensation is linked to the responsible management of generative AI’, can further facilitate an informed executive decision-making process (Kaul et al., 2025). Also, adaptive KPIs, such as ‘number of times a human-in-the-loop check led to intervention for generative AI systems in a high-risk area in a 3-month period’, further mitigate risks and improve governance in real time (Kaul et al., 2025). However, these KPI designs can lead to added financial and administrative burdens for those organizations with low budgets. For organizations to achieve sustainable KPI systems, they should not only design comprehensive sets of relevant KPIs, but they also need to be co-developed in iterative processes, which requires further organizational learning and engagement from the management team (Freeman et al., 2025).
Diversity and multi-stakeholder involvement in AI governance structures improve adaptability, mitigate groupthink, and enhance decision quality for generative AI. Smith (2025) has found that greater diversity within senior executive teams and the board has shown improvements to corporate social responsibility, stronger risk monitoring, and increased innovation rates. Incorporating diversity and including a variety of expertise can better identify emerging risks and respond to public sentiment. However, the implementation of these elements also needs to overcome practical barriers. In organizational structures where hierarchies exist, employees and departments can be highly reluctant to engage in organizational change. Further to this, implementing diversity on the

board can pose a resource and logistical burden. Furthermore, the participation of stakeholders from various sectors can lead to the addition of new perspectives.

Involving multi-stakeholder groups such as regulators and domain experts or external representatives in internal governance processes enhances the legitimacy of internal AI deployments and fosters trust in organizations (Rozenblit et al., 2025). These strategies ensure legitimacy in AI deployments. For instance, Rozenblit et al. (2025) noted that the health AI consortium in Canada brings together health-sector stakeholders from both public and private backgrounds to incentivize collaboration and legitimize AI systems. This type of participatory governance aligns incentives across various parties and stakeholders. However, these practices are underexploited because cross-sectoral and interdisciplinary governance can be costly, difficult to implement, and hard to scale. Generative AI governance needs to be adaptable to operate within fragmented regulatory and ethical environments for knowledge-intensive organizations. Standards such as ISO 42001 differ worldwide and from country to country (Poon Affat, 2025).

Organizations with headquarters and branches scattered across the globe need to adapt their governance approaches across multiple countries. This requires executives to not just tailor their approach to regulatory inconsistencies but ethical inconsistencies as well. A common challenge that has been identified with current executive governance approaches to AI is the ‘one-size-fits-all’ governance framework that is often described at a high level, lacking practical examples. Further to this, the dynamic nature of rules and regulations can make it challenging to adapt to the static implementation of many of them. According to evidence identified by Madanchian and Taherdoost (2025), adaptive governance is considered an optimal approach to facilitating ethical standards across a knowledge-intensive organization by continually updating the business with its ever-changing environments. However, implementing adaptive governance is a costly commitment and places burdens on organizations. Organizations are required to monitor frequently, invest in learning programs, and engage in dynamic interactions between their organizational members. As a next phase for organizational adaptation, scholars need to create hybrid governance models that integrate the adaptable aspects of governance into established standards and frameworks that can address the variations between jurisdictions and sectors.

Transparency and collaborative governance processes help foster trust and legitimacy in the deployment of generative AI for knowledge-intensive organizations. Organizations can improve trust through transparency in governance, providing transparency in documentation, audits, communication, and feedback. Through transparency, companies are more open and honest about their decisions and practices, diminishing information asymmetry (Rozenblit et al., 2025). For example, in health AI consortia, participatory governance methods align incentives across various stakeholder groups, legitimizing AI systems for broader healthcare applications, making them more trusted in society (Rozenblit et al., 2025). This, as a result, reduces the number of challenges that arise regarding public support. Such an approach in turn encourages AI adoption and acceptance. Transparency often needs to be managed effectively, or it can result in overly complex and lengthy

audit processes and communication chains, hindering performance (Rozenblit et al., 2025). This creates difficulties for organizations to move forward efficiently, as any time and effort wasted can be the determining factor for business failure. But, by gathering and incorporating external stakeholder feedback to governance processes, such as advisory panels and public consultation, organizations can gain legitimacy by ensuring decisions are fairer and more accountable, enhancing effectiveness. Organizations also need to make AI adoption more transparent to build legitimacy in society (Rozenblit et al., 2025). Stakeholders require a company's business processes to reflect values, such as ethical considerations. Incorporating this element enhances the perceived fairness and legitimacy of the knowledge-intensive business. Additionally, incorporating stakeholder and customer feedback will make processes more justifiable by clearly explaining the reasons behind organizational decisions and demonstrating how AI improves people's lives (Rozenblit et al., 2025). However, if these improvements are ignored, generative AI can create social harm to the customers of knowledge-intensive organizations, resulting in illegitimacy and decreased trust.
In conclusion, the managerial implications and theoretical contribution of executive governance for generative AI include innovative and adaptable KPI architectures, increased diversity and multi-stakeholder involvement, and dynamic systems that improve levels of trust and legitimacy.

6. Conclusion

The central objective of this thesis was to propose an integrated executive governance framework that enables knowledge-intensive organizations to realize generative AI's value-creation potential while systematically mitigating strategic, operational, and societal risks. While generative AI can drive operational optimization, business model innovation, enhanced decision-making, new customer experiences, and accelerated creative output, the literature and practice indicate a recurring execution gap: many organizations initiate pilots, but only a minority institutionalize governance mechanisms that allow scalable, repeatable, and auditable value realization. In line with the research question - How can executives design governance structures for generative AI in knowledge-intensive organizations to maximize organizational value while minimizing strategic and operational risks? - This capstone thesis conceptualizes governance as a deliberate, ongoing organizational capability that connects performance management, risk management, accountability, and control assurance across the generative AI lifecycle.

The findings synthesized through the systematic literature review (PRISMA-guided) indicate that generative AI in knowledge-intensive settings primarily creates value through (1) productivity and operational efficiency gains, (2) accelerated content and knowledge artifact creation, (3) improved customer engagement and personalization, and (4) enhanced problem-solving and decision support. However, value capture is frequently constrained by governance-related inhibitors: unclear decision rights and accountabilities, fragmented KPI architectures, insufficient capabilities and talent, weak risk evaluation routines, inconsistent controls across use cases, and disjointed technology and data architectures. In other words, organizations often fail to transition from experimentation to

operationalization not because of the absence of use cases, but because governance is under-specified at the executive level and under-instrumented at the operational level.

In parallel, the review identified a broad set of generative-AI risk domains with direct implications for value realization and organizational legitimacy, including cybersecurity vulnerabilities, data privacy breaches, misinformation and manipulation, regulatory uncertainty, bias and discrimination, ethical misuse, intellectual property conflicts, and workforce impacts. Importantly, the evidence also highlights “human and organizational” risk drivers - such as employee resistance, privacy concerns, low confidence in outputs, and negative psychological impacts - that can undermine adoption even when technical controls exist. This reinforces the need for governance that is both control-oriented and participatory: proactive testing and scenario exercises, feedback loops, adaptive KPI management, stakeholder engagement, and transparent communication are not “nice-to-haves,” but foundational for sustainable scaling.

Against this background, the capstone thesis addresses a core gap in prior work: while many studies emphasize regulation and compliance, the literature remains fragmented in its operational guidance for multi-level governance that reconciles regulatory dispersion with fast-moving technological change. Building on these observations, this capstone thesis contributes an integrated executive governance framework that explicitly links value drivers, risk tiers, decision rights (RACI), escalation pathways, and executive-grade KPI architecture. The framework is designed to function as an end-to-end governance “operating system” rather than a purely regulatory overlay: it specifies what must be governed, who decides, which controls are required at different risk levels, how performance is measured, and how incidents and underperformance are escalated to ensure timely executive oversight.

Nevertheless, several limitations affect the generalizability of the findings and the proposed framework. First, the evidence base in the reviewed literature is uneven across sectors; many detailed governance cases originate from heavily regulated domains (e.g., healthcare), and governance patterns may not transfer directly to less regulated industries with different risk appetites and control expectations (Brandtzaeg & Følstad, 2017). Second, universally validated metrics for value creation - particularly for intangible outcomes such as innovation, creativity, knowledge quality, and reputational effects - remain underdeveloped, which can complicate KPI standardization across organizations. Third, the rapid evolution of generative AI capabilities and ongoing regulatory changes imply that governance mechanisms require continuous calibration; static frameworks risk becoming obsolete or misaligned with emerging threat vectors and compliance obligations. Finally, behavioural, psychological, and ecosystem-level risks are still comparatively underexplored in the risk management literature, limiting the availability of validated constructs and measurement approaches for these dimensions.

Future research should therefore prioritize empirical evaluation and refinement of the proposed framework in real organizational settings and across multiple industries. Quantitative and mixed-methods studies could test causal pathways between governance mechanisms (e.g., decision rights

clarity, escalation design, KPI architectures, control maturity) and outcomes (e.g., adoption scale, incident frequency, realized value). Sector-specific standards and reference implementations could translate ethics and transparency requirements into operationally actionable control catalogues. Moreover, longitudinal research can illuminate how AI governance capabilities evolve over time, including the effectiveness of executive upskilling and role-based accountability programs in sustaining governance maturity. Finally, co-governance approaches spanning vendors, regulators, and industry alliances merit deeper examination, particularly for multinational organizations facing regulatory fragmentation. Behavioural risk should be expanded into measurable governance outcomes related to organizational justice, trust, well-being, perceived support, employee engagement, and stakeholder satisfaction in the context of generative AI.

Overall, this thesis underscores a practical conclusion: generative AI delivers durable value in knowledge-intensive organizations only when executive governance deliberately integrates performance management and risk management, aligns decision rights with accountability, and institutionalizes control assurance across the lifecycle. In an era of accelerating technological change and evolving regulation, integrated governance is not a constraint on innovation - it is the enabling architecture that allows organizations to scale generative AI responsibly, measurably, and sustainably.

Bibliography

Attard-Frost, B., & Lyons, K. (2024). AI governance systems: a multi-scale analysis framework, empirical findings, and future directions. AI and Ethics, 1–17. https://doi.org/10.1007/s43681-024-00569-5

Australian Institute of Company Directors & UTS Human Technology Institute. (2024). Eight elements of effective, safe, and responsible AI governance. Australian Institute of Company Directors and Human Technology Institute. https://www.aicd.com.au/content/dam/aicd/pdf/tools-resources/director-resources/concise-snapshot-web.pdf

Batool, A., Zowghi, D., & Bano, M. (2025). AI governance: a systematic literature review. AI and Ethics, 5, 3265-3279. https://doi.org/10.1007/s43681-024-00653-w

Baidya Nath Mukherjee (2025). Navigating AI governance: National and international legal and regulatory frameworks. In AURO University (Ed.), IGI Global (pp. 24). IGI Global. https://doi.org/10.4018/979-8-3373-1210-1.ch008

Bolden, D., Lukic, V., Martin, D., Luther, A., Kropp, M., Iyer, S., Bedard, J., Viner, B., Kleine, D., Mills, S., Martines, D., Ghai, J., McBride, J., Demyttenaere, M., Palumbo, S., Burke, D., Lin, A., Charanya, T., De Bellefonds, N., ... Meier, C. (2024). AI Unlocked: Value Creation with AI (10th ed.). Boston Consulting Group. https://www.bcg.com/assets/2024/executive-perspectives-value-creation-with-ai-17dec.pdf

Brandtzaeg, P. B., & Følstad, A. (2017). Why People Use Chatbots. In I. Kompatsiaris, J. Cave, A. Satsiou, G. Carle, A. Passani, E. Kontopoulos, S. Diplaris, & D. McMillan (Eds.), *Internet Science* (Vol. 10673, pp. 377–392). Springer International Publishing. https://doi.org/10.1007/978-3-319-70284-1_30

Chandra, B., & Rahman, Z. (2026). A framework for AI-driven value co-creation across customer journey stages. The Service Industries Journal, https://doi.org/10.1080/02642069.2026.2612702

Chompunuch, S., & Lubart, T. (2025). AI as a Helper: Leveraging generative AI tools across common parts of the creative process. J. Intell., 13(5), 57. https://doi.org/10.3390/jintelligence13050057

Coulson-Thomas, C. (2023). AI, Executive and Board Leadership, and Our Collective Future. Effective Executive, 26(3), 5–29.

https://gala.gre.ac.uk/id/eprint/44611/7/44611_COULSON%20THOMAS_AI_executive_and_board_leadership_and_our_collective_future.pdf

Dasgupta, A., & Wendler, S. (2019). AI adoption strategies (Working Paper Series – No. 9). Centre for Technology and Global Affairs, University of Oxford. https://www.politics.ox.ac.uk/sites/default/files/2022-03/201903-CTGA-Dasgupta%20A-Wendler%20S-aiadoptionstrategies.pdf

Eling, M., Nuessle, D., & Staubli, J. (2021). The impact of artificial intelligence along the insurance value chain and on the insurability of risks. The Geneva Papers on Risk and Insurance - Issues and Practice, 47, 205-241. https://doi.org/10.1057/s41288-020-00201-7

Freeman, S., Wang, A., Saraf, S., Potts, E., McKimm, A., Coiera, E., & Magrabi, F. (2025). Developing an AI governance framework for safe and responsible AI in health care organizations: Protocol for a multimethod study. JMIR Research Protocols, 14, e75702. https://doi.org/10.2196/75702

Gehrmann, S., Huang, C., Teng, X., Yurovski, S., Shode, I., Patel, C. S., Bhorkar, A., Thomas, N., Doucette, J., Rosenberg, D., Dredze, M., & Rabinowitz, D. (2025). Understanding and mitigating risks of generative AI in financial services (Preprint). Bloomberg. https://assets.bbhub.io/company/sites/51/2025/04/arXiv-Understanding-and-Mitigating-Risks-of-Generative-AI-in-Financial-Services-FINAL-4-25-25.pdf

Gianni, R., Lehtinen, S., & Nieminen, M. (2022). Governance of responsible AI: From ethical guidelines to cooperative policies. Frontiers in Computer Science, 4, 873437. https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2022.873437/pdf

Google. (2018). Perspectives on issues in AI governance. Google. https://ai.google/static/documents/perspectives-on-issues-in-ai-governance.pdf

Ghosh, A., Saini, A. & Barad, H. Artificial intelligence in governance: recent trends, risks, challenges, innovative frameworks and future directions. *AI & Soc* **40**, 5685–5707 (2025). https://doi.org/10.1007/s00146-025-02312-y

He, Y., Li, B., Cao, Y., Xu, H., Chen, Q., Chen, S., & Ren, S. (2026). Seeking human security consensus: A unified value scale for generative AI value safety. Research Institute for Data Management & Innovation, Nanjing University. https://www.arxiv.org/pdf/2601.09112

Hilb, M. (2020). Toward artificial governance? The role of artificial intelligence in shaping the future of corporate governance. Journal of Management and Governance. https://doi.org/10.1007/S10997-020-09519-9

Hogenhout, L. (2021). A framework for ethical AI at the United Nations (Unite Paper 2021). United Nations. https://arxiv.org/pdf/2104.12547

Kaul, A., Hneini, N., & Rahman, R. (2025). AI Risk Mitigation Framework. Roland Berger GmbH. https://www.rolandberger.com/publications/publication_pdf/AI_Risk_Mitigation_Framework.pdf?v=1594751

Kolt, N., Shur-Ofry, M., & Cohen, R. (2025). Lessons from complexity theory for AI governance. https://arxiv.org/pdf/2502.00012

Lee, C. H., Wang, Z., Wang, D., Lyu, S., & Chen, C. H. (2025). Artificial-intelligence-driven governance: addressing emerging risks with a comprehensive risk-prevention-centered model for public health crisis management. Health Research Policy and Systems, 23(115), 1-14. https://doi.org/10.1186/s12961-025-01390-0

Madanchian, M., & Taherdoost, H. (2025). Ethical theories, governance models, and strategic frameworks for responsible AI adoption and organizational success. Frontiers in Artificial Intelligence, 8, 1619029. https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2025.1619029/pdf

Maryala, B. T. R. (2025). Data governance in generative AI: A framework for transparency, compliance, and ethical practice. Journal of Computer Science and Technology Studies, 7(3), 964–971. https://doi.org/10.32996/jcsts.2025.7.3.108

Milkau, U. (2018). Value creation within AI-enabled data platforms. Journal of Creating Value, 5(1), 25-39. https://doi.org/10.1177/2394964318803244

NTT DATA. (2025). Mastering AI governance. NTT DATA, Inc. https://www.nttdata.com/global/en/-/media/nttdataglobal/1_files/insights/reports/agentic-ai/mastering-ai-governance-ntt-data.pdf

Nwashili, O. G. (2025). Scaling AI features in large organizations: A product management perspective. IRASS Journal of Economics and Business Management, 2(12), 23-30. https://irasspublisher.com/assets/articles/1765375019.pdf

Perry, B., & Uuk, R. (2019). AI governance and the policymaking process: Key considerations for reducing AI risk. Big Data Cogn. Comput., 3(2), 26. https://doi.org/10.3390/bdcc3020026

Poon Affat, R. L. (2025). All risk management frameworks: An expert panel discussion. Society of Actuaries Research Institute. https://www.soa.org/497d6f/globalassets/assets/files/resources/research-report/2025/ai-risk-management-frameworks-report-2025.pdf

Rao, A., Veillet, A., Kuperholz, M., Labovich, M., Cameron, E., & Ghosh, S. (2021). Responsible AI – Maturing from theory to practice. PwC. https://www.pwc.com/gx/en/issues/data-and-analytics/artificial-intelligence/what-is-responsible-ai/pwc-responsible-ai-maturing-from-theory-to-practice.pdf

Rozenblit, L., Price, A., Solomonides, A., Joseph, A. L., Srivastava, G., Labkoff, S., deBronkart, D., Singh, R., Dattani, K., Lopez-Gonzalez, M., Barr, P. J., Koski, E., Lin, B., Cheung, E., Weiner, M. G., Williams, T., Bui, T. T. T., & Quintana, Y. (2025). Towards a multi-stakeholder process for developing responsible AI governance in consumer health. International Journal of Medical Informatics, 195, 105713. https://doi.org/10.1016/j.ijmedinf.2024.105713

Santos, M. R. C., Carvalho, L. C., & Francisco, E. (2025). A capability-based framework for knowledge-driven AI innovation and sustainability. Information, 16(11), 1-18. https://doi.org/10.3390/info16110987

Schneider, J., Kuss, P., Abraham, R., & Meske, C. (2024). Governance of generative artificial intelligence for companies. https://arxiv.org/pdf/2403.08802

Sezgin, E. (2024). Redefining virtual assistants in health care: The future with large language models. J Med Internet Res, 26, e53225. https://doi.org/10.2196/53225

Sharma, Y. S. (2023). Generating wildfire risk maps for critical infrastructure systems using integrated generative AI and simulation techniques under information uncertainty [Doctoral dissertation, State University of New York at Buffalo]. ProQuest LLC. https://search.proquest.com/openview/165af4cf9402b51b98604abbc1e459ba/1?pq-origsite=gscholar&cbl=18750&diss=y

Shomali, M., Kumbara, A., MacLeod, J., & Iyer, A. (2025). Personalized cardiometabolic care powered by artificial intelligence. Frontiers in Endocrinology, 16, 1593321. https://www.frontiersin.org/journals/endocrinology/articles/10.3389/fendo.2025.1593321/pdf

Sinclair, K., & Mehta, H. (2023). Ethical considerations in AI governance: Towards responsible AI development. ITSI Transactions on Electrical and Electronics Engineering, 12(01), 1–17. https://journals.mriindia.com/index.php/itsiteee/article/download/143/130

Smith, G. K. (2025). Strategic integration of generative AI: Opportunities, challenges, and organizational impacts. Law, Economics and Society, 1(1), 156–168. https://doi.org/10.30560/les.v1n1p156

Soroori Sarabi, A. (2025). AI, global governance, and the need for an integrated disaster risk management system. Journal of World Sociopolitical Studies, 9(4), 815-851. https://doi.org/10.22059/wsps.2025.396240.1527

Strauss, I., Moure, I., O'Reilly, T., & Rosenblat, S. (2025). Real-World Gaps in AI Governance Research. arXiv. https://arxiv.org/pdf/2505.00174

Waisam, M., & Silver, N. (2025). From proof of concept to production: A guide to de-risking generative AI for real-world impact. Publicis Sapient. https://www.publicissapient.com/content/dam/ps-reinvent/us/en/2025/06/insights-lp/ai-risk-management-playbook/docs/A-Guide-to-Risking-Generative-AI.pdf

Walz, A., & Firth-Butterfield, K. (2021). Implementing ethics into artificial intelligence: A contribution, from a legal perspective to the development of an AI governance regime. Duke Law & Technology Review, 18, 180–231. https://scholarship.law.duke.edu/cgi/viewcontent.cgi?article=1352&context=dltr

Declaration of Authenticity

I hereby declare that I have completed this capstone thesis on my own and without any additional external assistance. I have made use of only those sources and aids specified, and I have listed all the sources from which I have extracted text and content. This capstone thesis, or parts thereof, has never been presented to another examination board. I agree to a plagiarism check of my capstone thesis via a plagiarism detection service.

Brisbane, 13.03.2026

Place, Date

Student signature