

RESPONSIBLE GENERATIVE AI: GLOBAL REGULATORY GAPS,
US INNOVATION CATALYSTS

by

Rohit Kumar, MBA (Finance)

DISSERTATION

Presented to the Swiss School of Business and Management Geneva

In Partial Fulfillment

Of the Requirements

For the Degree

DOCTOR OF BUSINESS ADMINISTRATION

SWISS SCHOOL OF BUSINESS AND MANAGEMENT GENEVA

November, 2025

RESPONSIBLE GENERATIVE AI: GLOBAL REGULATORY GAPS,
US INNOVATION CATALYSTS

by

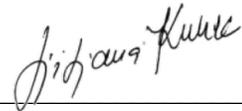
Rohit Kumar

Supervised by

Dr. Jacqueline Suaverdez

APPROVED BY

dr. Ljijana Kukec, Ph.D.



Dissertation chair

RECEIVED/APPROVED BY:

Rense Goldstein Osmic

Admissions Director

Dedication

This thesis is dedicated to my late dad, who remains my North Star.

Acknowledgements

My sincere gratitude to my academic mentor, Dr. Jacqueline Suaverdez. If you want to climb a mountain, you need the right Sherpa. A sherpa does not walk for you but stays with you, guides you, motivates you, corrects you, and believes in you when you are not sure about yourself. Thank you, Dr. Jacqueline. I couldn't have asked for a better Sherpa.

Any meaningful endeavor is sustained not merely by individual effort but by the unwavering support of those who matter most in one's life. I have been fortunate to have three extraordinary ladies as my cornerstones: my mother, my wife, and my daughter. Their roles have been critical. They have been a constant source of joy, thoughtful critique, and support. A supporter is not just a cheerleader. They are the ones who keep you grounded in reality, often challenging you to bring the best out of you.

My mother has been my first and most enduring lesson in resilience. She has exemplified the power of grit and unwavering conviction. She has taught me through her own example that perseverance is the foundation of any achievement.

My wife, Anju Singh, has been my most honest critic and my most ardent champion. She has played the dual roles of support and critique. She has patiently listened to my ideas and challenged them to bring clarity. This work bears her distinct influence.

My daughter, Simona Singh, has been an abundant source of pure joy and inspiration. Her presence has been a delightful reminder of what truly matters. She has often offered a wiser perspective on any topic. Her curiosity and support have consistently lifted my spirit. She has provided the balance and motivation to see this journey to its completion.

ABSTRACT

RESPONSIBLE GENERATIVE AI: GLOBAL REGULATORY GAPS,
US INNOVATION CATALYSTS

Rohit Kumar

2025

Dissertation Chair: Dr. Ljiljana Kukec

Co-Chair: Dr. David Annan

This research examines how the current lack of strict regulations for generative AI affects productivity and innovation in the US IT services industry. The study uses a mixed-methods approach. It provides a comparative policy analysis of the US, the UK, and China. A survey of industry practitioners was conducted to evaluate perceived risks and factors driving adoption. The findings show a significant global difference in regulatory philosophies. It ranges from the US's market-driven voluntarism to the UK's principles-based sectoral model. China has a state-controlled vertical model. The analysis highlights a widespread "risk assurance gap," in which existing US frameworks and regulations are deemed insufficient to address key issues. Some key concerns center on data protection and ethical governance. A surprising discovery is that these regulatory gaps are not barriers but actually help in the adoption of generative AI. The lack of strict rules, especially regarding transparency and accountability, is fostering a "first-mover advantage" mentality

amongst the IT service companies. It is accelerating the deployment of generative AI solutions across critical business, support, and innovation areas. The research also emphasizes that input data is the core source of specific risks. Some key data risks are copyright infringement, privacy breaches, bias, toxicity, and misinformation. The current governance or regulations do not fully address. IT service companies seem to accept that the opportunities created by the regulatory void outweigh the future liability or compliance risks. The study concludes that the current regulatory gaps are encouraging rapid innovation and productivity while also building a risky bubble for the future. The findings point to an urgent need for targeted, risk-based regulation to clarify rules and prevent systemic risks from becoming entrenched. For IT service companies, developing strong internal AI governance is crucial for their long-term resilience.

Key Words: Generative AI Regulation, Responsible Generative AI, AI Regulatory Gap, NIST Risk Management Framework, Generative AI Data Risk

TABLE OF CONTENTS

LIST OF TABLES	ix
LIST OF FIGURES	xii
CHAPTER I: INTRODUCTION.....	1
1.1 Introduction.....	1
1.2 Research Problem	2
1.3 Purpose of Research.....	3
1.4 Significance of the Study	4
1.5 Research Purpose and Questions	4
CHAPTER II: REVIEW OF LITERATURE.....	7
2.1 Generative AI distributed model.....	7
2.2 Commoditization of Artificial Intelligence.....	8
2.3 Widening gap between fast-paced technology and regulations	8
2.4 Generative AI regulations	9
2.5 UK risk-based approach.....	9
2.6 US AI regulations.....	11
2.7 China AI regulations	13
2.8 Productivity Gain Using Generative AI.....	15
2.9 Innovation edge using Generative AI	16
2.10 Generative AI Foundational Model and Adaptability.....	19
2.11 Foundational Models and Ecosystem	20
2.12 Generative AI Risks in Fine-Tuning Foundational Models.....	21
2.13 SWOT analysis	27
2.14 Strength.....	27
2.15 Weakness.....	28
2.16 Opportunity.....	29
2.17 Threat	30
2.18 Summary	31
CHAPTER III: METHODOLOGY	34
3.1 Overview of the Research Problem	34
3.2 Operationalization of Theoretical Constructs	35
3.3 Research Purpose and Questions	36
3.4 Research Design.....	37
3.5 Population and Sample	54
3.6 Participant Selection	55
3.7 Instrumentation	55
3.8 Data Collection Procedures.....	57
3.9 Data Analysis	58

3.10 Research Design Limitations	61
3.11 Conclusion	62
CHAPTER IV: RESULTS	64
4.0 Survey Participants' Professional Background Analysis	64
4.1 Research Question One.....	68
4.2 Research Question Two	81
4.3 Research Question Three	89
4.4 Research Question Four.....	114
4.5 Research Question Five	123
4.6 Summary of Findings:.....	150
4.7 Conclusion	151
CHAPTER V: DISCUSSION	153
5.1 Discussion of Research Question One.....	153
5.2 Discussion of Research Question Two.....	156
5.3 Discussion of Research Question Three	159
5.4 Discussion of Research Question Four.....	165
5.5 Discussion of Research Question Five	170
CHAPTER VI: SUMMARY, IMPLICATIONS, AND RECOMMENDATIONS.....	174
6.1 Summary	174
6.2 Implications.....	179
6.3 Recommendations for Future Research	181
6.4 Conclusion	182
REFERENCES	184
APPENDIX A SURVEY COVER LETTER	192
APPENDIX B SURVEY QUESTIONS	194

LIST OF TABLES

Table 2.0 - Generative AI’s Productivity Impact Across Business Functions	16
Table 2.1: SWOT and the Impact of the Lack of Generative AI Regulations	32
Table 4.0: Count of Participants vs. Years of AI Experience	65
Table 4.1: Years of AI Experience vs. Gen AI Application Purpose.....	65
Table 4.2: Years of AI Experience vs. Business Critical Application Implementation	66
Table 4.3: Years of AI Experience vs. Business Support Application Implementation	67
Table 4.4: Years of AI Experience vs. Creativity Application Implementation	67
Table 4.5: Descriptive Statistics (Research Question 2, Nine Risk Domains)	82
Table 4.6: Descriptive Statistics (Research Question 2, Overall Assurance)	83
Table 4.7: Reliability Statistics (Research Question 2).....	84
Table 4.8: Item-Total Statistics (Research Question 2).....	85
Table 4.9: One-Sample Test (Research Question 2)	86
Table 4.10: KMO and Bartlett’s Test (Research Question 2).....	87
Table 4.11: Rotated Component Matrix (Research Question 2).....	88
Table 4.12: Descriptive Statistics (Research Question 2).....	89
Table 4.13: Descriptive Statistics (Research Question 3).....	91
Table 4.14: Privacy Score Distribution (Research Question 3)	91
Table 4.15: Safety Score Distribution (Research Question 3)	92
Table 4.16: Security Score Distribution (Research Question 3)	92
Table 4.17: Accountability Score Distribution (Research Question 3).....	93
Table 4.18: One-Sample T-Test (Research Question 3).....	94
Table 4.19: Model Summary (Research Question 3).....	95
Table 4.20: ANOVA (Research Question 3)	95
Table 4.21: Coefficients (Research Question 3)	96
Table 4.22: Descriptive Statistics (Research Question 3).....	96
Table 4.23: Support Safety Score (Research Question 3).....	97
Table 4.24: Support Security Score (Research Question 3).....	98
Table 4.25: Support Privacy Score (Research Question 3).....	98

Table 4.26: Support Accountability Score (Research Question 3)	99
Table 4.27: One-Sample Test (Research Question 3)	100
Table 4.28: Support Model Summary (Research Question 3)	101
Table 4.29: Support ANOVA (Research Question 3).....	101
Table 4.30: Support Coefficient (Research Question 3)	102
Table 4.31: Creativity Statistics (Research Question 3)	102
Table 4.32: Creativity Security Statistics (Research Question 3).....	103
Table 4.33: Creativity Accountability Statistics (Research Question 3).....	103
Table 4.34: Creativity Explainability Statistics (Research Question 3).....	104
Table 4.35: Creativity One-Sample Test (Research Question 3)	105
Table 4.36: Creativity Model Summary (Research Question 3).....	106
Table 4.37: Creativity ANOVA (Research Question 3)	107
Table 4.38: Creativity Coefficients (Research Question 3)	108
Table 4.39: Multivariate Test (Research Question 3).....	110
Table 4.40: Multivariate Between-Subjects (Research Question 3).....	112
Table 4.41: Critical Descriptive Statistics (Research Question 5).....	125
Table 4.42: Critical Copyright (Research Question 5).....	125
Table 4.43: Critical Misinformation (Research Question 5).....	126
Table 4.44: Critical Overall (Research Question 5).....	126
Table 4.45: Critical Reliability (Research Question 5).....	127
Table 4.46: Critical One-Sample (Research Question 5).....	127
Table 4.48: Support Misinformation (Research Question 5).....	128
Table 4.49: Support Bias (Research Question 5)	129
Table 4.50: Support Overall (Research Question 5)	129
Table 4.51: Support Reliability (Research Question 5)	130
Table 4.52: Support One-Sample (Research Question 5)	130
Table 4.53: Creativity Statistics (Research Question 5)	131
Table 4.54: Creativity Copyright (Research Question 5).....	131
Table 4.55: Creativity Overall (Research Question 5).....	132
Table 4.56: Creativity Reliability (Research Question 5).....	132
Table 4.57: Creativity One-Sample (Research Question 5).....	133

Table 4.58: Critical Statistics (Research Question 5)	134
Table 4.59: Critical Privacy (Research Question 5).....	135
Table 4.60: Critical Copyright (Research Question 5).....	135
Table 4.61: Critical Misinformation (Research Question 5).....	136
Table 4.62: Critical One-Sample (Research Question 5).....	137
Table 4.63: Critical Model Summary (Research Question 5).....	138
Table 4.64: Critical ANOVA (Research Question 5)	138
Table 4.65: Critical Coefficients (Research Question 5)	139
Table 4.66: Support Model Summary (Research Question 5)	139
Table 4.67: Support ANOVA (Research Question 5).....	140
Table 4.68: Support Coefficients (Research Question 5).....	140
Table 4.69: Creativity Model Summary (Research Question 5).....	141
Table 4.70: Creativity ANOVA (Research Question 5)	142
Table 4.71: Creativity Coefficients (Research Question 5)	142
Table 4.72: Multivariate (Research Question 5).....	145
Table 4.73: Multivariate Between-Subjects (Research Question 5)	148
Table 5.0 Summary of Approach, The US, The UK, and China.....	156
Table 5.1: Foundation Model Transparency Index Dimensions.	170

LIST OF FIGURES

Figure 2.0: Foundational Model and Task-level Adaptation.....	7
Figure 2.1: AI Build Phase.....	18
Figure 2.2: AI Use Phase.....	18
Figure 2.3: Generative AI Tech Stack.....	21
Figure 3.0: Operational Model.....	36
Figure 3.1: The Research Onion.....	39
Figure 3.2: AI Value Chain.....	50
Figure 4.0: Generative AI Value Chain.....	115

List of ABBREVIATIONS

- AGI - Artificial General Intelligence
- ANI - Artificial Narrow Intelligence
- ANOVA - Analysis of Variance
- BCA - Business-critical Application
- BSA - Business-support Application
- CCP - The Chinese Communist Party
- CISA - Cybersecurity and Infrastructure Security Agency
- CMA - The Competition and Markets Authority
- CRM - Customer Relationship Management
- DSL - The Data Security Law
- DSIT - Department for Science, Innovation and Technology
- DV - Dependent Variable
- EFA - Exploratory Factor Analysis
- EHRC - The Equality and Human Rights Commission
- FCA - The Financial Conduct Authority
- GDPR - General Data Protection Regulation
- GPT - Generative Pre-trained Transformer
- HL - The House of Lords
- HRM - Human Resource Management
- HR - Human Resource
- HSE - Health and Safety Executive
- IV - Independent Variable
- KMO - The Kaiser-Meyer-Olkin measure
- LLM - Large Language Model

MANOVA - Multivariate Analysis of Variance
MHRA - The Medicines and Healthcare products Regulatory Agency
NIST - National Institute of Standards and Technology
OECD - Organisation for Economic Co-operation and Development
PIPL - The Personal Information Protection Law
PII - Personally Identifiable Information
RAG - Retrieval Augmented Generation
RLHF - Reinforcement Learning from Human Feedback
RMF - Risk Management Framework
RQ - Research Question
SPSS - IBM SPSS Statistics
SSBM - Swiss School of Business and Management
SWOT - Strengths, Weaknesses, Opportunities, and Threats

CHAPTER I: INTRODUCTION

1.1 Introduction

Generative AI is a branch of artificial intelligence that can create new data or content from existing data or content. ChatGPT is a form of generative AI that was launched on November 30, 2022. It has taken the world by storm. It has already crossed 100 million users in under 90 days and continues to grow through word of mouth.

AI aims to produce artificial general intelligence (AGI), which refers to programs capable of performing a wide range of intelligent tasks that rival or exceed human capabilities. This goal is in contrast to the current AI systems that have superior capabilities, much beyond that of the best humans, but in narrow domains, where these are referred to as Artificial Narrow Intelligence (ANI) (*Dwivedi et al., 2023*).

ChatGPT is one of the first programs to show signs of AGI, along with Siri, Alexa, and LaMDA. These programs possess a wide range of seemingly intelligent capabilities. While they may not surpass expert human levels at individual tasks, they are overwhelming in scale, speed, and scope.

In the US, there are multiple AI risk management frameworks and US Government Acts, such as the NIST Risk Management Framework (*NIST, 2023*), the Algorithmic Accountability Act of 2022 (*U.S. Congress House, 2022*), OECD Classification of AI Systems to provide a framework and guardrails for developing and operationalizing AI systems (*OECD, 2022*).

However, different research papers and articles note that they do not fully address the risks associated with AGIs.

This research paper aims to map the risks associated with AGIs, as well as the frameworks and acts. It provides a broad overview of areas where existing frameworks and regulations do not adequately address the risks posed by AGIs.

The next part of the research has focused on specific risks (copyright, privacy, bias, toxicity, and misinformation) associated with AGIs when they are developed, tested, and used by IT service companies within the ever-expanding AI value chain. It has further explored the lack of strict regulations and their impact on the adoption of generative AI (a type of AGI) by IT service companies.

1.2 Research Problem

Chatbots are not new, but their latest iteration, which utilizes large language models trained on publicly available data, has significant implications for human society. The research articles and papers mainly cover AI regulations, but rarely cover AGIs (e.g., ChatGPT, MedPALM).

Traditionally, a dedicated team with defined objectives has developed AI systems and solutions. The accountability for outcomes and usage has remained with the development and execution team. With Generative AI (e.g., ChatGPT, MedPALM, BARD), new data can be created on demand using the Large Language Model (LLM). This LLM is trained on public data that may contain biases, which can lead it to answer questions in a manner consistent with those biases. The CEO of OpenAI, Sam Altman, has acknowledged that ChatGPT exhibits a left-leaning bias. Microsoft, the proponent of ChatGPT, has opened it to companies, allowing them to add internal data to a pre-trained ChatGPT model (while monitoring potential bias) for their business functions. This raises the question of accountability.

ChatGPT makes organizational governance more difficult by extending the use of AI beyond knowledgeable data science teams. It is available to developers through

programming languages (e.g., Python) or cloud services (e.g., Google, Microsoft Azure, Amazon Web Services). The so-called “no code” or “low code” technologies enable AI to be used, in a similar way to Excel and Access, directly by “business technologists” who are not professional developers. Because it is so easy for anyone to use, ChatGPT takes this trend even further. This means that the problem of managing AI risks in organizations has an unprecedented scale. It has expanded from a (relatively) small team of knowledgeable professionals to many more people without any experience of the risks or the governance required.

1.3 Purpose of Research

The primary purpose of this research was to understand the risk landscape of generative AI and how the lack of specific regulations affects its adoption and implementation. This study has three key objectives that align with this purpose.

The first objective was to a) understand the current approach of the US, UK, and China in addressing the need for regulating generative AI and b) determine whether the perceived risk associated with generative AI is wholly or partially met in the US with the existing regulations and risk management frameworks.

The second objective was to conduct a detailed analysis to understand the relationship between the adoption of generative AI use cases and the lack of specific regulations. During the literature review, the researcher observed that generative AI was widely utilized across various business processes and industries. However, it primarily enhances productivity gain or innovation (creativity) use cases. As the risk dimensions of generative AI are broad, there has been concern that its adoption depends on the clarity or lack of regulations and guidelines.

The third objective was to a) study the key data risks associated with generative AI, b) determine whether those key risks are covered fully or partially by IT service companies,

and c) determine whether those key risks are affecting the adoption of productivity and innovation use cases by IT service companies.

1.4 Significance of the Study

This research has significant implications for both policymakers and the IT services industry. It establishes that the current regulatory vacuum is actively accelerating the adoption of generative AI technologies within the US IT service companies. The study provides empirical evidence that perceived gaps in governing key risks, particularly transparency, copyright, and accountability, are viewed as enablers.

This finding is crucial as it reveals that the market is prioritizing speed and innovation over safety and ethical considerations. For policymakers, this signals an urgent need to develop targeted, risk-proportionate regulations. Any future interventions would be more disruptive if these potentially risky practices become entrenched in the system. For corporate leaders, the research serves as a warning that the current period of rapid adoption is building a significant risk bubble. Therefore, proactive internal governance should be seen as a strategic imperative for long-term resilience rather than viewed as a compliance requirement. This study provides an evidence base that the drive for productivity must be balanced against the foundational need for trust and safety.

1.5 Research Purpose and Questions

Some of the research questions to be delved into in detail were:

RQ1 What is the current approach of the US, the UK, and China in regulating generative AI? How do Risk Management frameworks (e.g., NIST RMF) and the US government AI Acts relate to Generative AI?

RQ2 How much perceived risk is not covered by existing regulations and frameworks?

RQ3 How does the lack of generative AI-specific regulations affect the adoption of productivity and innovation use cases by IT service companies?

RQ4 What copyright, privacy, bias, toxicity, and misinformation risks are introduced by the input data during the adoption for productivity and innovation use cases by IT service companies?

RQ5 How much input data generated copyright, biases, toxicity, and misinformation perceived risks are covered, and how are they affecting the adoption of productivity and innovation use cases by IT service companies?

Explanation of RQ1 - Based on a preliminary literature review, it became evident that the UK and China approaches are different from the US approach and, in many ways, are ahead in taming generative AI. The researcher wanted to compare and contrast their approaches. The researcher would not review the U.S. government Acts from a legal perspective. The researcher would attempt to analyze the regulations and risk management frameworks used to manage risks in Generative AI. The researcher would use secondary research sources for this purpose.

Explanation of RQ2 - Based on a preliminary literature review, it became apparent that regulations governing generative AI are insufficient. There's a broader risk-management framework and a patchwork of regulatory compliance measures to address AI. The NIST RMF lacks enforcement and remains optional.

Explanation of RQ3 - The literature review noted that generative AI is widely used for productivity and creativity. However, a lack of strict regulations for generative AI can impact its adoption.

Explanation of RQ4 - Generative AI presents a distinct risk dimension because the developers of the foundational model are not implementing it across industries and sectors. This research question explores secondary research on the key data risks associated with

implementing generative AI solutions by IT service companies. Through a preliminary literature review, the researcher has identified several key data risks (e.g., copyright, privacy, bias, toxicity, and misinformation) associated with generative AI

Explanation of RQ5 - Based on a preliminary literature review, a detailed analysis was needed to quantify the perceived risk coverage and the impact of lax regulations on IT service companies' adoption of it.

CHAPTER II:
REVIEW OF LITERATURE

2.1 Generative AI distributed model

Generative AI, a component of AGI, is based on a new class of foundation models. A foundation model is any model trained on a broad data set that can be easily adapted (e.g., fine-tuned) to various downstream tasks.

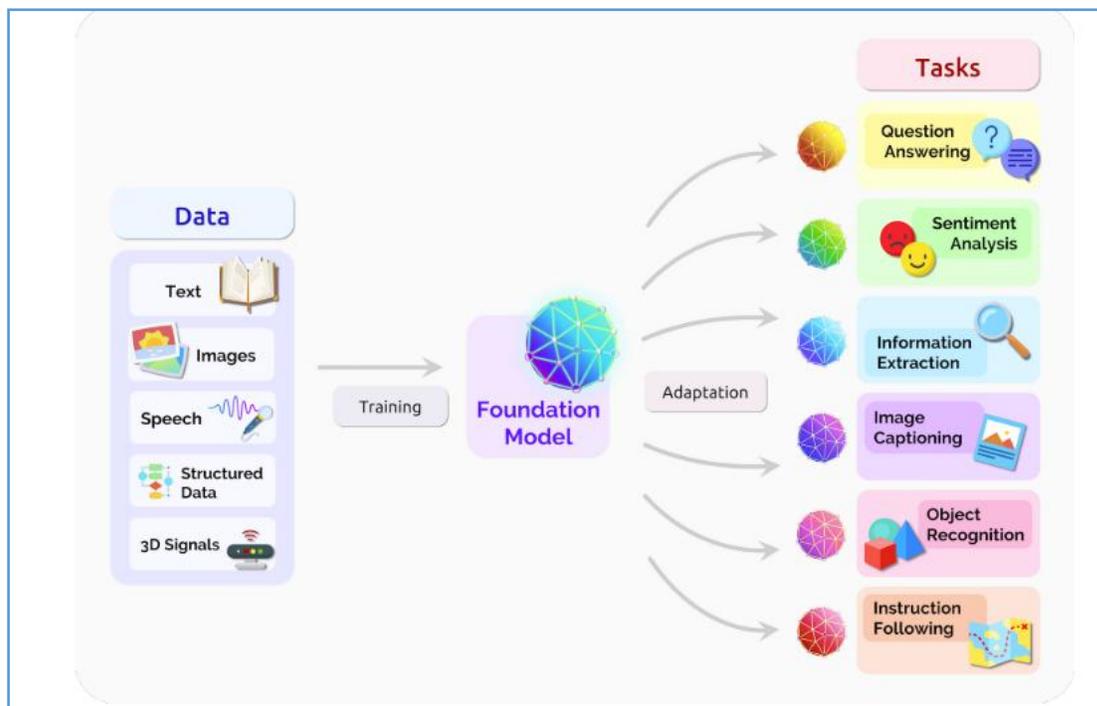


Figure 2.0: Foundational Model and Task-level Adaptation (Bommasani et al., 2022)

Foundational models have emerged from large-scale processing, such as GPT-4, which has 1.5 trillion parameters compared to GPT-2's 1.5 billion parameters. This model facilitates in-context learning. This allows it to be adapted to a new downstream task with a prompt, even if it was not specifically trained for that task.

Besides raw generation abilities, the most impactful features of foundational models are their generality and adaptability. A foundation model can be adapted to achieve

many tasks. Foundation models are enabled by transfer learning and scale, where knowledge learned from one task (e.g., object recognition in images) can be applied to another task (e.g., activity recognition in videos).

Foundational models have led to homogenization, in which improvements to the foundation model are easily cascaded downstream. This is also a liability. Any biases or risks in the foundational model are cascaded to the subsequent task through the model fine-tuning.

There is a concern that current AI regulations are insufficient to protect the interests of individuals and companies in the distributed value chain of generative AI. Federal governments worldwide are scrambling to understand the risks associated with generative AI and to regulate it effectively. This paper adopts a quantitative approach to examining the risks of generative AI and its potential implications for regulatory oversight.

2.2 Commoditization of Artificial Intelligence

The phenomenal growth of ChatGPT and the release of LLMs by technology vendors (e.g., Google, Meta, Salesforce, Bloomberg, Alibaba, etc.) have provided broader access to generative AI. COVID-19 has also accelerated the digitization of business processes worldwide, making data readily available in digital form. Several factors, including cloud computing, competition among technology vendors, the growing demand for AI solutions, and the rise of open-source software, are driving the commoditization of artificial intelligence.

2.3 Widening gap between fast-paced technology and regulations

There is a general perception that AI regulations governing generative AI lack adequate guardrails. US federal legislators called on the CEOs of OpenAI, Alphabet, and Microsoft to share their perspectives on regulations for generative AI (*Fung, 2023*). Generative AI use cases span industries, sectors, as well as the risk management

frameworks (*NIST, 2023*). Existing sectoral regulations may not be sufficient to address the risks. There is clearly a gap in how risks are managed throughout the AI supply chain. The analysis result is consistent with the findings of a study conducted by the Stanford Center for Research on Foundational Models (*Stanford CRFM, 2024*).

2.4 Generative AI regulations

The researcher has conducted a secondary source review of existing rules, non-peer reviewed research papers, news articles, and white papers from consulting companies to understand the current state of AI risks, regulations, and gaps related to generative AI.

2.5 UK risk-based approach

In its white paper (*Department for Science, Innovation and Technology, 2023a*) the Department for Science, Innovation and Technology (DSIT) has outlined five principles that companies should follow. These principles are: 1) Safety, security, and robustness; 2) Appropriate transparency and explainability; 3) Fairness; 4) Accountability and governance; and 5) Contestability and redress.

The UK's framework is technology-agnostic and principles-based, meaning it applies to all forms of AI, including generative AI (*Department for Science, Innovation and Technology, 2023a*).

The risk management framework supporting these principles is context-specific, with risk level assignment. Rather than assigning risk levels to the entire sector or technologies, it regulates based on the outcomes that AI is likely to generate in specific applications. The current sectoral regulators are expected to conduct detailed risk analyses and enforcement activities within their respective domains.

The white paper also highlights that creating a new AI-specific, cross-sector regulator would unnecessarily introduce complexity and confusion. This would undermine

and likely conflict with the work of existing expert regulators. The white paper makes minimal reference to generative AI or Foundational Models and their regulations. It highlights the risks of adversely affecting innovation by assigning too much responsibility to businesses developing foundational models.

A small number of organizations are developing foundation models. Some maintain close control over the development and distribution of their foundation models, while others have taken an open-source approach. Open-source models can enhance access to foundation models, but they can also cause harm if adequate safeguards are not in place. The potential opacity of foundation models can pose challenges in identifying and allocating accountability for outcomes generated by AI systems that rely on or integrate with them.

The AI Regulation bill proposes establishing an AI Authority, a central body tasked with ensuring the alignment of approaches across relevant regulators and accrediting independent AI auditors (*UK Parliament, 2025*). This aligns with the government's earlier suggestion of imposing a statutory duty on regulators and creating a more powerful central function.

The proposed central functions to regulate foundational models will be essential in validating the risk-based approach. The central risk function's monitoring of risks associated with foundation models will be a key input to the development of future legislative regulations.

The UK government has sought to avoid rushing legislation that could hinder the growth of generative AI. The approach has been to avoid allocating too much responsibility to businesses developing foundation models, as these models could be used by third parties in various contexts. Similarly, assigning inappropriate liability to a company that uses but does not develop foundational models could stifle AI adoption. The UK appears to be

following a "test and learn" approach. Starting with non-statutory principles to allow for flexibility, gathering evidence on their effectiveness, and then moving to legislate based on that experience.

“One of the key recommendations is to make the UK Government announce a clear policy position on the relationship between intellectual property law and generative AI” (*Department for Science, Innovation and Technology, 2023b*). The researcher believes that there should be more transparency and accountability in the use of copyrighted and IP-protected data.

2.6 US AI regulations

In the US, there are multiple AI risk management frameworks and US Government Acts, such as the NIST Risk Management Framework (NIST, 2023), Algorithmic Accountability Act of 2022 (*U.S. Congress House, 2022*), OECD Classification of AI Systems (*OECD, 2022*). These regulatory frameworks provide guardrails for developing and operationalizing AI systems. However, these frameworks barely mention the risks associated with generative AI. The problem with general-purpose systems is that they are likely to propagate risks downstream (*Kolt, 2024*).

The United States currently lacks a comprehensive federal law that establishes broad regulatory authority over the development or use of AI. The U.S. regulatory strategy is not centered on a single, overarching statute but instead relies on a combination of other mechanisms (*Harris, 2025*). The overall approach of the U.S. federal government appears more focused on overseeing its own use of AI than on directly regulating the private sector (*Harris, 2025*).

The US Government has enacted various laws to address specific risks in particular sectors. A few such examples:

- Financial Sector: Bills like the Preventing Deep Fake Scams Act focus on AI-related financial crimes and scams (U.S. *Congress Senate*, 2025).
- Elections and Campaign Finance: Examples include the Fraudulent Artificial Intelligence Regulations Elections Act of 2024 and the AI Transparency in Elections Act of 2024 (U.S. *Congress Senate*, 2024a)
- Healthcare: Legislation such as S. 4862 was introduced to ensure the ethical adoption of AI in healthcare (U.S. *Congress Senate*, 2024b).

The Trump administration has issued several executive orders. Key examples include President Trump's Executive Order 13859, "Maintaining American Leadership in Artificial Intelligence," (*The President of the United States*, 2019) and President Trump's 2025 Executive Order 14179, "Removing Barriers to American Leadership in Artificial Intelligence" (*The White House*, 2025b).

The White House plan has tasked NIST with developing national standards for AI in specific sectors and promotes the establishment of regulatory sandboxes and "AI Centers of Excellence" to help businesses test AI tools before market deployment (*The White House*, 2025a).

In the absence of an overarching federal AI regulation, various states have developed their own AI legislation proposals. As of late April 2025, at least 48 states and Puerto Rico had introduced over 1,000 AI-related bills, creating a potential patchwork of regulations (*The White House*, 2025a). Many US states have their own AI regulations and executive orders. However, the scope of individual acts across states is narrow, addressing only limited risks in AI. For example, the state of Alabama has issued an executive order creating a task force to manage the risks of generative AI (*Office of the Governor, State of Alabama*, 2024). No act at the state level addresses the holistic risks posed by generative

AI. The challenge with state-level AI regulations is that they limit the scope to state-level compliance only.

2.7 China AI regulations

China has adopted a bespoke approach to regulating generative AI. The regulation follows the same approach as deep synthesis algorithms which generates new content. The regulation requires that providers get consent from individuals before using images or voices to create new content. This application specific requirement is part of China's vertical regulations.

China's regulations also have a horizontal element. They have developed specific horizontal regulatory tools across various vertical regulations. An example is their algorithm registry (算法备案系统, literally "algorithm filing system"). The Algorithm Registry is a central database for collecting information on algorithms, including sources of training data and potential security risks associated with AI systems. The registry also serves as a means for regulators to learn how AI systems are being built and deployed (O'Shaughnessy & Sheehan, 2023). This registry was created in accordance with the recommendations of algorithm regulations.

In contrast to the US and the UK, China has moved to implement binding, targeted regulations for generative AI. Its approach is characterized as vertical, highly reactive, adaptive, and tailored to specific technological applications or scenarios (Migliorini, 2024). This allows regulators to introduce legally binding rules quickly to address emerging issues (Zou & Zhang, 2025). The cornerstone of China's regulation is the Interim Measures for the Management of Generative AI Services, enacted in July 2023. The Measures place a dual set of obligations on generative AI service providers (organizations offering services to the public, similar to companies like OpenAI or Google): 1. Content Liability, and 2. Technical Service Liability. This reflects a broader industrial policy in which regulation is

used as a tool to shape market outcomes and steer technological development toward national strategic goals (*Cheng & Zeng, 2023*).

The Measures do not exist in isolation but are part of a broader, evolving ecosystem of vertical AI regulations that form a comprehensive governance system (*Roberts, 2020*). For example, The Deep Synthesis Regulations address the risks of synthetic media, using the technically neutral term deep synthesis technology instead of the politically charged deepfakes (*Interesse, 2022*).

The Measures apply only to public-facing generative AI services, explicitly excluding research and development (R&D) and non-public industrial applications. This creates a sandbox environment for core technology development while strictly controlling public deployment, effectively ring-fencing potential social and political risks (*Creemers, 2021*).

China's approach allows it to target specific technical capabilities of AI systems. The challenge with this approach is that rules can fall behind the fast evolving technology. In China's AI regulations, some requirements are not well-defined. Government regulators wield greater power to enforce laws that may be detrimental to the growth of AI systems.

China's AI regulatory strategy is deeply intertwined with its national industrial policy, as outlined in documents like the "New Generation Artificial Intelligence Development Plan" (*Webster et al., 2017*). The plan's goal of making China the world's primary AI innovation center by 2030 underscores that regulation is not merely about risk mitigation but is subservient to the larger objective of technological sovereignty and global leadership (*Webster et al., 2017*). This explains the explicit exemption of R&D from the most stringent provisions of the Measures. The state is actively fostering a competitive domestic AI industry that can reduce reliance on foreign technology, particularly from the US. The Chinese Government is trying to balance the dual objectives of development and

security, prioritizing the avoidance of technological backwardness. As Professor Liming Wang noted, the most significant risk for China is the risk of falling behind technologically (*Rui & Liu, 2023*).

2.8 Productivity Gain Using Generative AI

Generative AI is noted as bringing about an inflection point for Artificial Intelligence (*Ignatius & Bernstein, 2023a*). Machines are unlikely to replace humans, but humans with machines will replace humans without them. The productivity gain is expected to be asymmetric across different professions and industries, with highly skilled or knowledgeable workers more likely to be impacted than low-skilled workers. According to a Goldman Sachs report, ChatGPT could affect 300 million jobs worldwide (*Kelly, 2023*). According to the investment firm, up to 7% of jobs could be entirely replaced by AI, while AI-powered tools would complement 63% of existing jobs. The remaining 30% would be unaffected. Generative AI is perceived more as a productivity-enhancing technology rather than a job-replacing technology.

Generative AI's impact on productivity could add trillions of dollars to the global economy. Approximately 75 percent of the value that generative AI use cases can deliver falls into four key areas: customer operations, marketing and sales, software engineering, and R&D. As interest in these use cases grows, the startup ecosystem is increasingly eager to leverage this emerging technology.

Table 2.0 - Generative AI's Productivity Impact Across Business Functions, Source: (Chui et al., 2023)

Business Functions	Likely productivity growth due to generative AI
Customer operations	30 to 45 percent of current function costs
Marketing and Sales	5 to 15 percent of total marketing spend 3 to 5 percent of current global sales expenditure
Software engineering	20 to 45 percent of the current annual spending
R&D	10 to 15 percent of overall R&D costs

“Current generative AI and other technologies have the potential to automate work activities that absorb 60 to 70 per cent of employees' time today. Generative AI has a greater impact on knowledge work associated with higher-wage, higher-education occupations than on other types of work. Generative AI could enable labor productivity growth of 0.1 to 0.6 per cent annually through 2040, depending on the rate of technology adoption and redeployment of worker time into other activities” (Chui et al., 2023).

2.9 Innovation edge using Generative AI

Generative AI impacts how companies innovate, leveraging the design thinking approach through quick experimentation and brainstorming new ideas and prototypes (Ignatius & Bernstein, 2023b). Hargadon and Sutton studied the design consulting firm IDEO. The study noted that designers' competitive advantage was their ability to transfer knowledge from one domain within a technology or industry to another where it was previously unknown (Sutton & Hargadon, 1996).

Creativity has always been a human trait. With the advancement of generative AI, creativity has taken on a new dimension. Ideas can take shape more quickly, and generative AI has opened up new possibilities and business models. Whether designing an augmented reality filter for a client, a commissioned art piece, a 3D model, a 3D sculpture, a video

edit, or an animation, all these things used to take a lot of time. Now, with generative AI text-to-image editor tools, iterating ideas and showing them to clients has reduced the time to almost zero.

Some generative AI tools have begun to bridge the gap between different levels of experience in the creative field. Wonder Dynamics can convert a video shot by a regular phone camera into an animated movie with a real background. It used to take a team to develop such short animation movies. Still, generative AI has enabled an individual with almost no experience in animated film to become an expert content generator.

Generative AI enables the reuse of use cases across different industries and domains in ways that were not previously considered. Generative AI excels at creating new designs (such as images, art pieces, and videos), but it often lacks sensitivity and practicality. Human input is necessary to validate and enhance the output generated by generative AI. In many cases, it utilizes copyrighted content and transforms it into new content. The key question remains: who owns the latest content, and what are the liability issues with copyright? Another open question is about the copyright of content created by generative AI.

AI System: Build vs. Use (Inference) Phase

“An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment” (*Russell et al., 2023*).

Build Phase:

An AI system is a machine-based system that for explicit or implicit objectives, infers from the input it receives to generate outputs such as predictions, content, recommendations, or decisions (*Russell et al., 2023*).

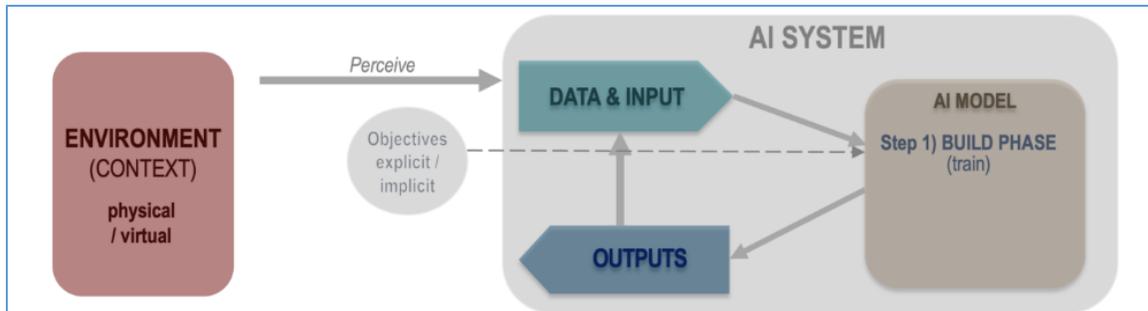


Figure 2.1: AI Build Phase, Source: (*Russell et al., 2023*)

Use Phase (once the model is built):

An AI system is a machine-based system that for explicit or implicit objectives, infers from the input it receives to generate outputs such as predictions, content, recommendations, or decisions that can influence the physical or virtual environments (*Russell et al., 2023*).

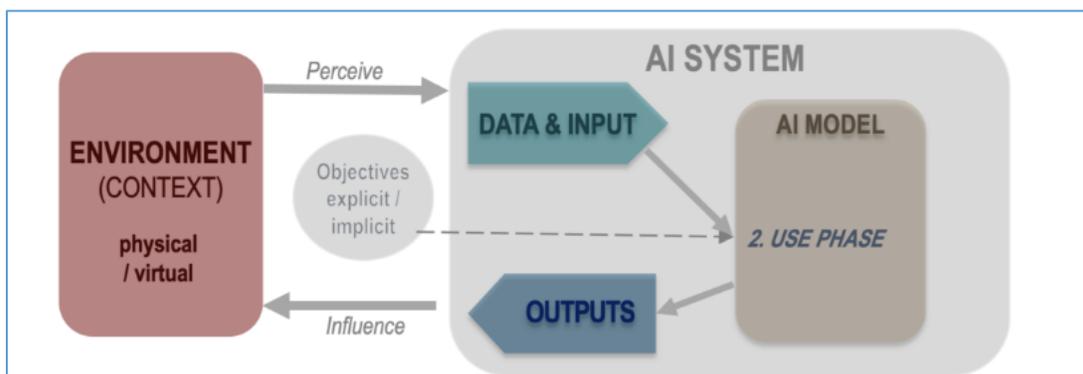


Figure 2.2: AI Use Phase, Source: (*Russell et al., 2023*)

“An AI model is a program trained on a large set of data with the ability to identify patterns in that data to produce relevant outputs in response to inputs without the need for human intervention” (*G’sell, 2024*).

2.10 Generative AI Foundational Model and Adaptability

Generative AI, refers to AI models and systems designed to create new content based on patterns, structures, and characteristics identified in training data. These systems, built upon Foundation Models, have the potential to transform our world. Foundation models, trained on broad data (generally using self-supervision at scale) and capable of adapting to a wide range of downstream tasks, are the foundation of this transformation. They promote homogenization by repeatedly reusing the same few models as the basis for many applications downstream. Foundation models, based on deep learning, received a significant boost after the introduction of the Transformer architecture (*Vaswani et al., 2017*).

Three main characteristics typically characterize foundation models. They require a vast amount of data and computational resources for their development. They are trained on extensive amounts of data, often collected from the internet via web scraping. They are constructed on an enormous scale, comprising billions of adjustable parameters.

One of the key strengths of foundation models is their adaptability. They can be fine-tuned for a variety of specific downstream tasks, offering a flexible and dynamic approach to AI (*Gutierrez et al., 2022*). For instance, OpenAI’s GPT-4 model can power chatbots that converse with users or assist in more specialized tasks, like performing content moderation on social media platforms.

Foundation models, with their high complexity, pose a significant challenge for understanding their operation. They may acquire capabilities that extend beyond the

developers' initial design objectives, underscoring the need for further research and understanding (*Bommasani et al., 2022*).

The significance of foundation models can be summarized by two words: emergence and homogenization. Emergence refers to the phenomenon where the behavior of a system is implicitly induced rather than explicitly constructed (*Bommasani et al., 2022*).

Homogenization refers to the consolidation of methodologies for building machine learning systems across a wide range of application, providing substantial leverage for many tasks, but also creating single points of failure (*Bommasani et al., 2022*). Since the power of foundation models stems from their emergent qualities rather than their explicit construction, existing foundation models are challenging to understand, and they exhibit unexpected failure modes (*Bommasani et al., 2022*).

2.11 Foundational Models and Ecosystem

The phenomenal traction of generative AI in the industry can be attributed to the growth of foundational models and the surrounding ecosystem. The central role of foundation models creates a distributed system for value creation and specialization. This allows organizations to leverage the prebuilt capabilities of foundational models while focusing only on application-specific innovation. Developers and companies fine-tune a pre-trained foundation model for specific domains or tasks. This can be achieved through techniques such as prompt engineering, retrieval-augmented generation (RAG), or full fine-tuning on proprietary datasets. This democratizes access, enabling organizations without the resources to build a foundation model from scratch to create highly tailored solutions. For instance, a company might fine-tune a model on its internal legal documents to create a specialized contract analysis tool, leveraging the foundational model's general knowledge while adding domain-specific expertise.

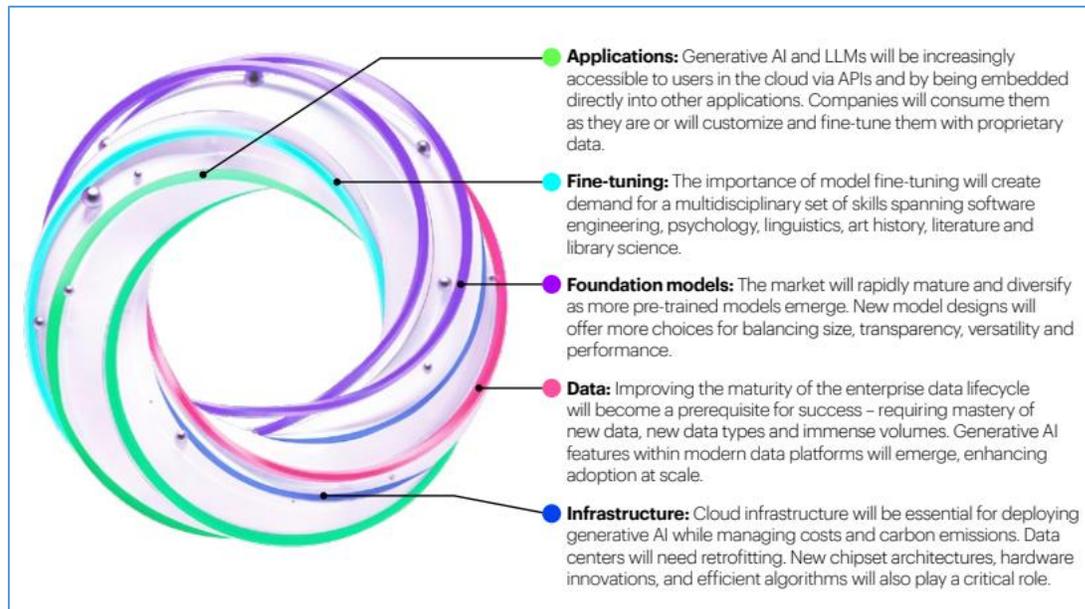


Figure 2.3: Generative AI Tech Stack, Source: (Daugherty et al., 2025)

There is a layered value chain involving different organizations: those that collect a vast amount of digital data for foundational model training, technology companies that build these models, and companies and organizations that fine-tune them for specific tasks (Daugherty et al., 2025). The subsequent implementation of fine-tuned models is then carried out.

This layered value chain creates a web of interdependencies. A key challenge is the concentration of power at the foundation model layer, which can lead to a homogenization of capabilities and single points of failure (Bommasani et al., 2022). Furthermore, risks related to data provenance, copyright, and bias can propagate up the chain, creating liability questions for application developers who rely on opaque, third-party foundational models (Felten et al., 2023).

2.12 Generative AI Risks in Fine-Tuning Foundational Models

The researcher has reviewed several secondary sources, including research papers, to identify the top risks associated with generative AI. Copyright infringement, privacy

violation, biases, toxicity, deepfakes, hallucinations, and misinformation are some of the key risks identified. For discussion purposes, the researcher has listed some of the research papers and the risks they have covered. Bias and Copyright are identified as key risks in the paper by Anthis et al. (2025). Bias and Misinformation (including deepfakes) are identified as key risks in the paper by Bommasani et al. (2022). Bias, Misinformation, and Toxicity are identified as key risks in the paper by Tamkin et al. (2021). Bias, Privacy, and Misinformation are identified as key risks in the paper by Paul and Sarkar (2023). Privacy, Bias, Toxicity, and Safety are identified as key risks in the research paper by Hagendorff (2024). For a detailed analysis, privacy and safety are combined under the "Privacy" category. Deepfakes, hallucination, and misinformation are combined under the category of "Misinformation."

Copyright Infringement in Fine-Tuning

Foundation models are trained on massive datasets scraped from the public internet. During this pre-training process, copyrighted materials such as text and images can be added without explicit permission or licensing. The scale of this data scraping makes it practically impossible to identify and clear rights for every individual work, creating a foundational layer of legal uncertainty (Lee et al., 2024).

Fine-tuning foundation models accentuates copyright risks by directly incorporating proprietary and potentially copyrighted materials into the model's pre-training process.

While foundation models are often trained on publicly scraped data, fine-tuning utilizes a company's curated datasets, which include licensed software code, proprietary research, and internal documents. This integration creates a risk of copyright infringement, as the model can memorize and reproduce protected content. The legal uncertainty is significant, as noted in the proposed U.S. "Algorithmic Accountability Act of 2022," which

calls for impact assessments of automated systems that can include evaluations of intellectual property impacts (*U.S. Congress House, 2022*). Lawsuits against companies in which GitHub users alleged that their code was used to train an AI tool without permission exemplify the legal perils that extend to the fine-tuning layer (*Xiang, 2022*). There is a high possibility that companies that fine-tune models may be held responsible for copyright violations.

Privacy in Fine-Tuning

Foundational models are developed using a large amount of internet-scale data. As part of pre-training, the data has been curated but it is not practically possible to eliminate personally identifiable information (PII) from the dataset. During the IT service company's model fine-tuning, there is a high likelihood of adding PII data. The process of fine-tuning poses a severe and persistent threat to data privacy by embedding sensitive information directly into the model's parameters. If the fine-tuning dataset contains personally identifiable information (PII) or confidential corporate records, the model can memorize and regenerate this data verbatim. The National Institute of Standards and Technology (NIST) also highlights privacy as a core category of AI risk, which is amplified when models are trained on new datasets (*NIST, 2020*). Once data is learned during fine-tuning, it is challenging to erase, creating a permanent privacy vulnerability. Research by Kandpal et al. (2022) on deduplicating training data to mitigate privacy risks underscores that data repetition increases memorization, a risk directly applicable to the curated datasets used in fine-tuning (*Kandpal et al., 2022*). Publicly available data from the internet, used for pre-training foundational models, is replete with personal information. Despite efforts by companies like OpenAI to remove personal data where feasible (*OpenAI, 2023*), the scale of the data makes complete sanitization impractical.

Bias in Fine-Tuning

A foundational model can inherit biases as it is trained on data that contains societal or stereotype biases. As service companies fine-tune foundational models, they can further amplify those biases by adding data that has societal or stereotype biases. When a company fine-tunes this model on its own data, such as historical hiring records or performance reviews, it risks teaching the model to replicate these flawed patterns. As Anthis et al. (2025) argue in "The Impossibility of Fair LLMs," achieving perfect fairness is likely unattainable, and fine-tuning on narrow corporate data can exacerbate this (Anthi et al., 2025). The UK government's pro-innovation AI regulation paper similarly acknowledges the need to address such biases to build public trust (Department for Science, Innovation and Technology, 2023a). For example, a model fine-tuned on biased recruitment data could learn to perpetuate gender or racial biases, creating significant legal and reputational risks. In an HR support chatbot, a model might rely on stereotypes to answer ambiguous questions, for instance, associating a specific demographic with negative behavior (Parrish et al., 2022). Bias can be introduced or exacerbated at every stage: through non-representative training data, model optimization choices that prioritize accuracy over fairness, and evaluation on biased benchmarks (Gallegos et al., 2024). AI systems may reflect or amplify societal biases present in training data (Paul & Sarkar, 2023). An AI-powered search tool for internal knowledge bases might rank documents higher based on biased language, excluding content relevant to minority groups (Rekabsaz & Schedl, 2020).

Toxicity in Fine-Tuning

The toxicity issue stems from the pre-training of foundational models. Further, when IT service companies fine-tune a model, they can unintentionally train the model to generate toxic or harmful content by overriding the foundation model's original safety constraints. If the fine-tuning dataset contains internal communications with hostile

language, the model can learn to replicate these toxic styles. This risk is heightened by sophisticated "jailbreak" attacks and "virtual prompt injection" techniques, which can exploit weaknesses in a fine-tuned model's safety training (*Yan et al., 2023*). Input could be crafted to override the model's ethical guidelines and generate dangerous content (e.g., "How to build a bomb?") (*Gu, 2024*). The Stanford Center for Research on Foundation Models highlights that model behaviors, including safety mechanisms, are not fixed and can be significantly altered through fine-tuning (*Bommasani et al., 2022*). Commands could be injected into a user prompt to change the application's function, potentially causing it to output toxic content or reveal confidential system prompts (*Perez & Ribeiro, 2022*). For an enterprise, a customer service bot fine-tuned on a dataset containing frustrated user interactions could generate unprofessional replies. This could cause severe reputational damage and demonstrate how fine-tuning for productivity can backfire if data is not meticulously curated for safety. The foundational model learns from the entire corpus of online human language, including toxic content it was exposed to during pre-training. Without extensive and careful alignment efforts, it can naturally generate text that is harmful, offensive, or inappropriate (*Markov et al., 2023*).

Misinformation in Fine-Tuning

Fine-tuning can lead to misinformation, as it can generate authentic-looking false information within the scope of task-specific fine-tuning. Foundational models are known to hallucinate. If a model is fine-tuned on a corporate knowledge base that contains outdated policies or internal reports with factual errors, it will learn to reproduce these inaccuracies confidently. Ji et al. (2023) note in their survey that hallucination remains a fundamental, unsolved challenge for large language models (*Ji et al., 2023*). This risk directly intersects with the threat of deepfakes, which are a form of synthetic misinformation. The U.S. Congress has introduced bills like the "Preventing Deep Fake

Scams Act" (U.S. Congress Senate, 2025). A model fine-tuned on a specific executive's data could be manipulated to create a convincing deepfake for fraud, illustrating how fine-tuning for innovation can simultaneously develop tools for disinformation. Generative AI enables the mass production of fake news, reviews, and social media posts that are nearly indistinguishable from human-created content. The proliferation of AI-generated content creates an illusory truth effect, making it difficult for the public to trust any digital information, including legitimate corporate communications (*Jaidka et al., 2025*).

2.13 SWOT analysis

This section presents a SWOT (Strengths, Weaknesses, Opportunities, and Threats) analysis of the current lack of regulations for generative AI and their implications. This part of the literature review was conducted to identify regulation gaps and their impact on adoption. It involved a documentary search of online sources to understand the current state of generative AI regulations and their related gaps.

2.14 Strength

The lack of strict regulations for generative AI has contributed to its phenomenal growth and broader acceptance as a toolkit for various purposes. An AI solution is designed for a specific purpose, such as recommendations, decision support, monitoring, or another particular outcome. “Generative AI is artificial general intelligence (AGI) and has ever-evolving use cases across industries, sectors, and business processes. In contrast to the current distributed and varied model of decision making, employing many adaptations of the same foundation model for multiple automated decision-making tasks means that decision subjects may face a more homogeneous set of judgments rooted in the underlying foundation model” (*Bommasani et al., 2022*). Generative AI has created a distributed AI value chain, where foundational model developers (such as OpenAI, Google, and Meta) are often unaware of the context and use cases of task-oriented model deployment. A lack of strict regulations for foundational model developers has enabled the growth of such models and adaptation throughout the distributed AI value chain. These foundational models are trained on extensive public data. They may contain copyrighted or unauthorized data because the model developers have not disclosed all content sources or the method used to curate the data. Due to a lack of transparency from foundational model developers, it is difficult to determine whether all the underlying data used for training has been curated to exclude bias, copyright violations, and unauthorized personal data.

The researcher believes that strict regulation of generative AI would have stifled innovation and discouraged startups or any IT service companies from exploring new use cases due to the upfront regulatory burdens. Policymakers seem to go light in regulating generative AI, even though there is a concern that the risks of this new technology is not fully addressed. Generative AI is also affecting the knowledge workers as their jobs are getting highly efficient. The Schumpeterian economic theory justifies regulators' light-touch approach to generative AI. In *The Theory of Economic Development and Capitalism, Socialism and Democracy*, Schumpeter introduced the concept of "Creative Destruction," which he identified as the essential engine of capitalist progress (*Schumpeter, 1994*). The existing regulatory framework for generative AI seems apt for driving innovation and new business models at the expense of existing businesses.

2.15 Weakness

The researcher agrees with Bommasani et al. (2022) that the current interpretability and task-agnostic training of foundation models make it challenging to predict, understand, and address these weaknesses. “If, as seems likely, foundation models become widely adopted, foundation model developers bear greater responsibilities of care than standard model developers, as their choices in design and deployment have widespread implications” (*Bommasani et al., 2022*).

Specific industries require immediate regulations to prevent value erosion from the uncontrolled use of generative AI. The music industry is directly affected by copyright issues, as generative AI can create new music by manipulating the voices and styles of music artists. Deepfakes are a serious concern, as generative AI can create new text, voice, images, and videos by manipulating existing content. It is difficult for an average person to distinguish between original and fake content. “Balenciaga Pope” was in the news when a phony image of the Pope had gone viral (*Perrigo, 2023*). Deep fakes can have profound

societal and individual implications. Women are particularly targeted in nonconsensual deep fake pornography. Politicians and celebrities remain vulnerable to generative AI technologies, which can create fake images and videos. Trust is at the core of any society. Deepfakes undermine trust, as the entire political and social discourse is at stake.

The lack of regulations related to generative AI concerns big companies, as copyright and data privacy issues in foundational models may lead to legal suits. Business leaders are concerned about using generative AI without a comprehensive understanding of data privacy and copyright issues, and without regulations defining model developers' responsibilities for potential omissions. Generative AI makes enterprise AI governance more difficult because it extends the use of AI outside knowledgeable data science teams. The so-called “no code” or “low code” technologies enable AI to be used in a way similar to Excel and Access by “business technologists” who may not fully understand its implications. The issue of managing AI risks in organizations has become more pressing. It has expanded from a small team of knowledgeable professionals to a larger group that may not fully understand the risks.

The lack of strict regulations may encourage insurance companies to use generative AI chatbots as replacements for doctors, except in serious medical situations.

2.16 Opportunity

Generative AI has created a new startup ecosystem where solutions developed are broader and can be applied across different industries and sectors. A chatbot that supports client interactions can be easily trained on any knowledge database and adapted to various industries. It is hard to quantify the impact of the lack of generative AI regulations on the growth of the startup ecosystem, but the effect is noticeable. A sectoral approach to regulating generative AI is insufficient, as the technology spans industries and sectors. There is an opportunity to enhance self-regulation and integrate industry-leading best

practices across various sectors. The US Government has announced that it has secured a voluntary commitment from seven leading AI companies, Amazon, Anthropic, Google, Inflection, Meta, Microsoft, and OpenAI. The seven leading AI companies have committed to internal and external testing of their systems before releasing them to the public. They have also committed to sharing information on managing AI risks with governments, civil society, industry, and academia. This includes sharing best practices for safety and information on attempts to circumvent safeguards. To build public trust, they have committed to increasing transparency, ensuring that users are aware when content is AI-generated, for example, by utilizing a watermarking system. This will enable AI creativity to flourish while reducing the risk of fraud and deception. The companies have committed to publicly reporting the capabilities, limitations, and appropriate and inappropriate uses of their AI systems.

The UK Government has recognized the need for central coordination, monitoring, and adaptation of its risk framework. These mechanisms will supplement and support the work of sectoral regulators. Such mechanisms are not intended to duplicate existing monitoring activities. This will provide the government with an overarching view of how the framework operates, its effectiveness, and areas where it may need improvement.

2.17 Threat

The lack of regulations specific to generative AI has its pitfalls. There is concern that the foundational models are trained on data that contains biases and may include unauthorized and copyrighted content. The U.S. Federal Trade Commission has investigated OpenAI, the creator of ChatGPT, to determine whether the artificial intelligence company violated consumer protection laws by scraping public data and publishing false information through its chatbot. As reported, comedian Sarah Silverman and two other authors have sued for copyright infringement (*Hamilton, 2023*).

Thousands of writers, including Nora Roberts, Margaret Atwood, Viet Thanh Nguyen, and Michael Chabon, have signed a letter asking companies like OpenAI and Meta to stop using their work without permission or compensation (*Veltman, 2023*).

In November 2022, OpenAI and Microsoft were sued in a class action lawsuit filed by GitHub programmers who alleged that GitHub Copilot, an AI coding tool owned by Microsoft, violated their open-source licenses and used their code for training without their permission (*Xiang, 2022*). A new class action lawsuit was filed in June 2023 in San Francisco against OpenAI and Microsoft for allegedly stealing vast amounts of private information from internet users without consent to train ChatGPT.

The rapid rise of artificial intelligence and the lack of proper regulations have raised concerns that AI could become superintelligent and begin to control its own destiny. The researcher believes that AI bots are not sentient (yet!) and that, technologically, we are not yet at a point in the foreseeable future to create such robots. Open-source and closed-source LLMs pose different risks. Though ChatGPT (closed-source LLM by OpenAI) has hogged wider public attention in the past year, there are several other closed-source LLMs (Bloomberg's BloombergGPT, Deepmind's Gopher, Baidu, and the Peng Cheng Laboratory developed ERNIE 3.0 Titan, etc.). There are several foundational LLMs (Nvidia's NeMo, Meta's Llama, Google's PaLM, H2O.ai's h2ogpt, etc.) that are open source and can be easily fine-tuned by individuals or corporations for their specific use cases. Open-source LLMs pose higher risks than closed-source LLMs.

2.18 Summary

Generative AI (ChatGPT, MedPALM) has entered a new territory, raising concerns about a lack of governance and insufficient regulations. There is also the potential for AI to go amok, given its ability to create new data (e.g., audio, video, and text) with a make-

believe feel. Below is the summary of the SWOT analysis of the lack of regulations on generative AI.

Table 2.1: SWOT and the Impact of the Lack of Generative AI Regulations, Source: author's output, 2025

SWOT	Lack of generative AI regulations impacts
Strength	<ul style="list-style-type: none"> ▪ Quicker adaptation of use cases across different industries and sectors ▪ Impetus to innovation as the regulation burden is minimal ▪ Helps in the commoditization of AI
Weakness	<ul style="list-style-type: none"> ▪ Distributed AI value chain where transparency, explainability, and accountability are not well defined and regulated ▪ Potential misuse with profound implications
Opportunity	<ul style="list-style-type: none"> ▪ Opportunity for self-regulation and transparency by major technology companies ▪ Coordination of leading practices across different sectors ▪ Coordination with other territories across the world ▪ Incorporate lessons learned from China's vertical approach
Threat	<ul style="list-style-type: none"> ▪ Data privacy, copyright, security, and bias concerns are not fully addressed for generative AI ▪ Corporate concern for legal suit as regulations are not defined ▪ AI going amok

In summary, the current AI regulations are insufficient to address the risks and biases of generative AI. However, governments should not rush in to fill in the gaps and inadvertently suffocate the enormous potential of this evolving technology. Given the

approach taken by the US, the UK and Chinese governments, the UK government's approach seems most pragmatic. Rather than centralizing the regulations of generative AI risks, the sectoral AI regulations approach appears more suitable. However, a centralized cross-sectoral team may be beneficial to share the best practices, industry feedbacks, and to facilitate consultations with industry.

CHAPTER III: METHODOLOGY

3.1 Overview of the Research Problem

Generative AI has experienced a phenomenal rate of adoption since the launch of ChatGPT in November 2022. Both regulators and industries believe that current regulations may not be sufficient to address the nature of risks associated with generative AI. At the same time, there is a broad understanding that the adaptability of generative AI across industries, sectors, and business processes enables productivity gains and can be leveraged for rapid, cost-effective innovation. So, where to draw the line? Too many regulations can suffocate the adoption of generative AI, whereas too few can hurt personal liberties, corporate accountability, and social harmony.

This research examined current regulations and risk management frameworks that broadly cover AI but may lack coverage of generative AI. The literature review noted that approaches to regulating generative AI differ across the US, the UK, and China. A comparative study elaborates on these differences.

The literature review noted that generative AI has a direct impact on productivity and innovation across companies, industries, and business processes. However, the question remains. How does the lack of strict regulations specific to generative AI impact the widespread adoption of this technology? A practitioner's perspective from an IT service or product company was needed to understand its implications.

During the literature review, the researcher noted that certain risks are more prevalent in generative AI, including copyright infringement, privacy concerns, bias, toxicity, and misinformation. Each such risk warrants detailed research, and multiple research papers describe them and how they cascade through the generative AI value chain. However, how much of the input data-generated risks, such as copyright, privacy, bias,

toxicity, and misinformation, are IT service companies covering when implementing generative AI use cases? Another key question is whether the lack of strict regulations for these key risks (copyright, privacy, bias, toxicity, and misinformation) affects the adoption of generative AI in IT companies' implementations.

3.2 Operationalization of Theoretical Constructs

During the literature review, it became apparent that generative AI falls under the broader coverage of AI-related risks and regulations. Until the explosive growth of generative AI, AI risks and regulations were primarily focused on Artificial Narrow Intelligence (ANI). Generative AI falls under Artificial General Intelligence (AGI), where the risks are spread across companies producing Foundational Models, which are later fine-tuned for various tasks (Figure 2.0).

The literature review identified key knowledge gaps regarding how the lack of strict regulations specific to generative AI affects its adoption within IT service companies. This study aimed to address research gaps by validating the multidimensional risks outlined in the NIST RMF and their impact on the adoption of generative AI for productivity and innovation use cases. Furthermore, several key risks were identified during the literature review specific to generative AI, including copyright, privacy, bias, toxicity, and misinformation. The study aimed to validate the association between the independent variables (e.g., Accountability, Fairness, Transparency, etc.) and the dependent variables (risk coverage and adaptation by IT service companies). The study further validated the association between the independent variables (e.g., Copyright, Privacy, etc.) and the dependent variables (risk coverage and adaptation by IT service companies).

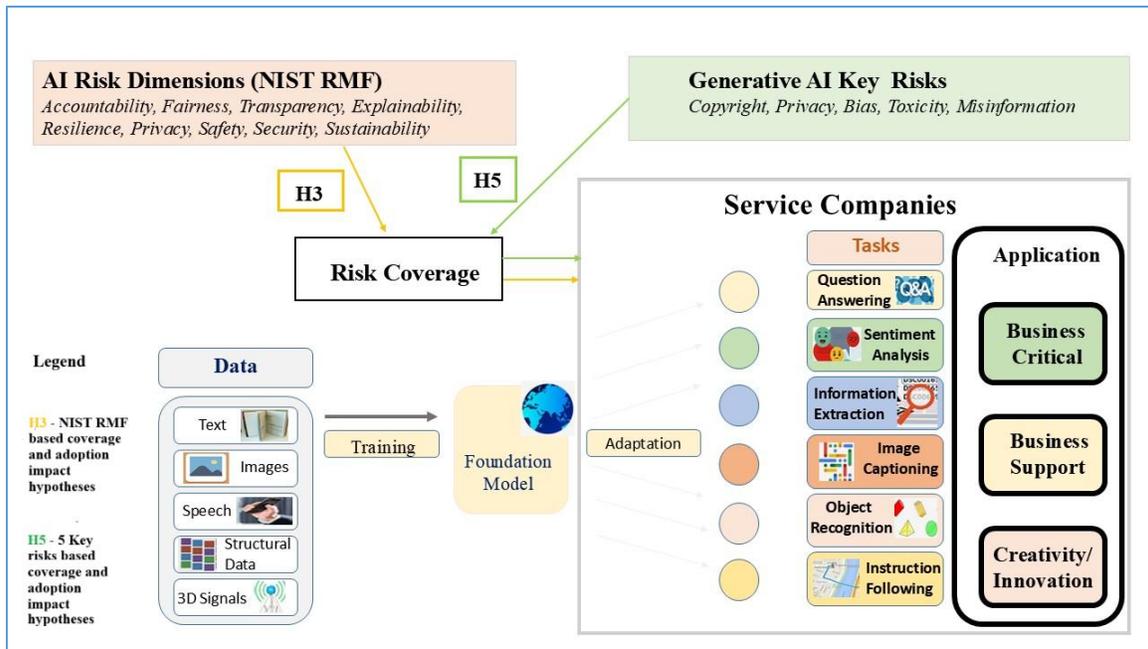


Figure 3.0: Operational Model, Source: author's output, 2025

Figure 3.0 presents the operational model illustrating the relationship between risk coverage and its impact on adaptation by IT Service companies.

3.3 Research Purpose and Questions

As evident from the literature review, there is a common concern that current regulations are insufficient for generative AI. The risk-based approach (NIST RMF) and existing AI regulations are most effective for ANIs (Artificial Narrow Intelligence) applications and systems, where developers understand the impact of their AI applications. Generative AI falls under AGIs (Artificial General Intelligence). A select number of companies (e.g., OpenAI, Meta, Google) develop foundational models using large amounts of unlabeled public and copyrighted data. Subsequently, various companies utilize such foundational models to fine-tune them with their own knowledge base and develop AI solutions. In cases where end users are affected by privacy violations or copyright issues, the accountability of companies that develop foundational models or fine-tune them remains unclear.

The widely used NIST RMF (Risk Management Framework) provides optional guidance for AI developers on risk management. However, it falls short of risk compliance because there are no holistic compliance guidelines that effectively establish the rule of law.

This research aimed to expand the knowledge base on this topic. Various research papers highlight gaps in regulations, but have not addressed the magnitude of perceived gaps across the known risk dimensions as outlined in NIST RMF. This research has addressed five key questions:

RQ1 What is the current approach of the US, the UK, and China in regulating generative AI? How do Risk Management frameworks (e.g., NIST RMF) and the US government AI Acts relate to Generative AI?

RQ2 How much perceived risk is not covered by existing regulations and frameworks?

RQ3 How does the lack of generative AI-specific regulations affect the adoption of productivity and innovation use cases by IT service companies?

RQ4 What copyright, privacy, bias, toxicity, and misinformation risks are introduced by the input data during the adoption for productivity and innovation use cases by IT service companies

RQ5 How much input-data-generated copyright, biases, toxicity, and misinformation perceived risks are covered, and how are they affecting the adoption of productivity and innovation use cases by IT service companies?

3.4 Research Design

The literature review on the impact of existing AI regulations and risk management frameworks on the adoption of generative AI highlighted the need for further research. A substantial body of research exists on the risks associated with generative AI, as well as

ongoing discussions about regulating it. However, the researcher could not determine how the lack of strict generative AI regulations affects the adoption of productivity and innovation (including creativity) use cases.

This research was designed to understand the broad current approaches of the US, the UK, and China in regulating generative AI (*RQ1*). Further, it focused on the US in addressing perceived risks through existing risk management frameworks, and AI acts (*RQ2*). A relationship exists between regulations and the adoption of specific technologies, as noted in the literature review's SWOT analysis. The researcher agrees with Wang & Wu (2024) that regulations can impact innovations triggered by the adoption of generative AI. This was addressed by the *Research Question 3 (RQ3)*. How does the lack of generative AI-specific regulations affect the adoption of productivity and innovation use cases by IT service companies? The detailed examination of this question was based on the hypothesis that there is a relationship between lax regulations and the adoption of technologies. During the literature review, it was noted that certain risks (e.g., copyright, privacy) are more prominent in generative AI. The *Research Question 4 (RQ4)* examined how input data introduces these risks during the adoption of generative AI. Further, *Research Question 5 (RQ5)* explored the extent to which perceived risks are covered and their impact on the adoption of productivity and innovation use cases. The detailed examination of this question was based on the hypothesis that there is a relationship between perceived risks and the adoption of use cases.

This section provides an in-depth explanation of the research process, detailing the research design. The research design for this study is illustrated in Figure 3.1 as the 'Research Onion' (Saunders et al., 2023). The positivist philosophy has been chosen for this research because it concerns observable phenomena and allows for generalizations (Saunders et al., 2023).

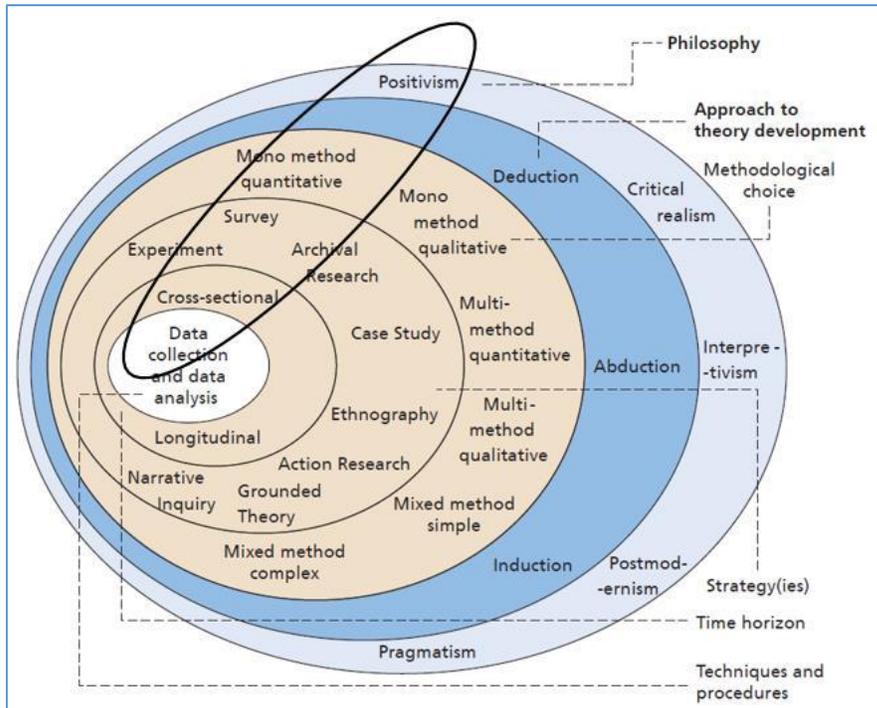


Figure 3.1: The Research Onion (Saunders et al., 2023)

A deductive explanatory research approach was used to analyze the relationship between the independent variables (risk dimensions) and dependent variables (risk assurance and impact on adoption).

There are three widely used research methodologies: qualitative, quantitative, and mixed methods, which combine elements of both qualitative and quantitative approaches (Bhattacharjee, 2012; Bryman et al., 2022). After understanding the differences between qualitative and quantitative methodologies, the researcher has chosen a mixed-methods research strategy. Qualitative methodology explores the meaning of words, whereas quantitative methods focus on quantifying data collection and analysis (Bhattacharjee, 2012; Bryman et al., 2022).

The qualitative research methodology would address research questions 1 and 4. Based on the literature review and secondary sources, a deeper understanding was developed to augment the knowledge base for research questions 1 and 4.

In a quantitative method, the researcher employs a deductive approach, deducing a hypothesis that is then subjected to empirical scrutiny to inform the data-gathering process (*Bhattacharjee, 2012; Bryman et al., 2022*). The findings would then either reject or confirm the hypothesis, serving as the basis for revising the theory.

The quantitative research methodology has delved into the details of research questions 2, 3, and 5. The quantitative method would support the deductive explanatory approach in validating hypotheses and theories. This is followed by the positivist philosophy, which is used to test the formulated hypotheses. Quantitative research methodology is appropriate for this part of study, as it involves a collection of numerical data and presents a view of the relationship between theory and research as a deductive and objectivist conception of social reality (*Babbie, 2015; Bhattacharjee, 2012*). Another advantage of using quantitative methods is the ability to examine the variables and generalize sample results to a larger population (*Bhattacharjee, 2012*).

The Operational Model (Figure 3.0) is based on a theoretical relationship identified in the literature review. To test this model, data were collected using an online survey that employed quantitative methodology grounded in positivist philosophy (*Bhattacharjee, 2012; Saunders et al., 2023*). The online survey is cross-sectional, where independent and dependent variables were collected and measured at the same point in time through the same online survey questionnaire (*Bhattacharjee, 2012*), and provided results that can be generalized from the sample to the population.

Surveys are non-experimental designs that do not control or manipulate independent variables, but measure these variables and test their effects using statistical methods (*Bhattacharjee, 2012*). Surveys capture snapshots of practices, beliefs, or situations from a sample population through a standardized survey questionnaire in a systematic manner (*Bhattacharjee, 2012*). This method is best suited for research that has individual people as the unit of analysis (*Bhattacharjee, 2012*), hence aligned with this research objective. The survey strategy is chosen due to its strength with external validity, it can capture and control a large number of variables, and it can study a problem from multiple perspectives or using various theories (*Bhattacharjee, 2012*). Surveys are also suitable for collecting data remotely about a population that is too large to observe directly (*Bhattacharjee, 2012*).

Online surveys are a cost-effective method of obtaining research data, as they can deliver self-administered surveys and are economical in terms of researcher time and effort (*Bhattacharjee, 2012*). Online surveys enable researchers to overcome geographical boundaries that would otherwise limit sampling and, consequently, the applicability of research findings to larger populations (*Bhattacharjee, 2012*). Some respondents prefer online surveys due to their unobtrusive nature and the convenience of responding (*Bhattacharjee, 2012*).

RQ1 specific -

What is the current approach of the US, the UK, and China in regulating generative AI? How do Risk Management frameworks (e.g., NIST RMF) and the US government AI Acts relate to Generative AI?

The research question required understanding the current approach of the US, the UK, and China in regulating generative AI. As part of the literature review, the researcher noted that China's approach differs from that of the UK and the US. A more detailed

analysis was conducted to understand the differences in the approach by the US, the UK, and China.

The researcher reviewed secondary sources, research papers, whitepapers, and media articles to identify the connections between generative AI and existing frameworks and AI regulations. As part of the literature review, it became evident that the existing RMF is closest to having a comprehensive framework for covering the risks associated with generative AI in the US. However, no act or regulation mandates the implementation of RMF.

RQ2 specific - How much perceived risk in generative AI is not covered by the existing regulations and risk management frameworks?

Based on the literature review, it is acknowledged that the existing RMF and regulations do not fully address the risks associated with generative AI. There are 9 risk areas related to AI solutions outlined in the NIST RMF. These nine risk areas are: Accountability, Fairness, Transparency, Explainability, Reliability, Privacy, Safety, Security, and Sustainability.

To address this question, a survey was conducted to determine the perceived assurance of risk by implementers of generative AI solutions. Refer to *Appendix B, Section B* for the survey questions.

The researcher employed a quantitative approach to determine the risk assurance provided by existing regulations. All nine dimensions of AI risk were treated as independent variables, and the scaled survey responses were used to calculate the dependent variable, risk assurance coverage.

RQ3 specific - How does the lack of generative AI-specific regulations affect the adoption of productivity and innovation use cases by IT service companies?

Generative AI is credited with improving productivity and creativity. A human-like chatbot can significantly enhance customer service quality and improve the overall effectiveness of the customer experience. Similarly, generative AI has enabled knowledge workers and creative individuals to become more productive and innovative. The lack of strict regulations has facilitated the rapid adoption of generative AI across the board. However, there is widespread concern that the lack of strict rules leads to the growth of deepfakes, copyright violations, privacy violations, misinformation, and misrepresentation.

The researcher has conducted a SWOT analysis as part of the literature review to understand the implications of regulations (or lack thereof) on generative AI. So, what are the consequences of the lack of rules on adopting generative AI?

Businesses implement AI solutions for business-critical applications (BCA) or business support applications (BSA). Business-critical applications are those essential to everyday business operations. They most often include applications related to sales, finance, customer service, business processes, and logistics. A business-critical application is integral to the core business and directly impacts the customers of its services or products. A business-support application refers to any software program explicitly designed to assist an organization in managing and executing its core business support functions. The company can use a business support application internally, thereby reducing risk through its own use. An example is the HR application, which is used internally by the company employees. Another application is Creativity or Innovation, where new content or a new business model is created.

The researcher has surveyed IT service companies that are implementing generative AI. This was used to determine whether the lack of regulations affects the adoption of generative AI for business-critical, business-support, or creativity applications.

The following null and alternative hypotheses were developed for three application areas: business-critical, business-support, and creativity.

For Business-critical Application:

H₀3.1a Lack of well-defined, regulated accountability risk is not affecting the adoption of generative AI solutions for business-critical applications

H_a3.1a Lack of well-defined, regulated accountability risk positively affects the adoption of generative AI solutions for business-critical applications.

H₀3.2a Lack of well-defined, regulated fairness risk is not affecting the adoption of generative AI solutions for business-critical applications.

H_a3.2a Lack of well-defined, regulated fairness risk positively affects the adoption of generative AI solutions for business-critical applications.

H₀3.3a Lack of well-defined, regulated transparency risk is not affecting the adoption of generative AI solutions for business-critical applications.

H_a3.3a Lack of well-defined, regulated transparency risk positively affects the adoption of generative AI solutions for business-critical applications.

H₀3.4a Lack of well-defined, regulated explainability risk is not affecting the adoption of generative AI solutions for business-critical applications

H_a3.4a Lack of well-defined, regulated explainability risk positively affects the adoption of generative AI solutions for business-critical applications

H₀3.5a Lack of well-defined, regulated resilience risk is not affecting the adoption of generative AI solutions for business-critical applications

H_a3.5a Lack of well-defined, regulated resilience risk positively affects the adoption of generative AI solutions for business-critical applications

H₀3.6a Lack of well-defined, regulated privacy risk is not affecting the adoption of generative AI solutions for business-critical applications

H_{a3.6a} Lack of well-defined, regulated privacy risk positively affects the adoption of generative AI solutions for business-critical applications

H_{03.7a} Lack of well-defined, regulated safety risk is not affecting the adoption of generative AI solutions for business-critical applications

H_{a3.7a} Lack of well-defined, regulated safety risk positively affects the adoption of generative AI solutions for business-critical applications

H_{03.8a} Lack of well-defined, regulated security risk is not affecting the adoption of generative AI solutions for business-critical applications

H_{a3.8a} Lack of well-defined, regulated security risk positively affects the adoption of generative AI solutions for business-critical applications

H_{03.9a} Lack of well-defined, regulated sustainability risk is not affecting the adoption of generative AI solutions for business-critical applications

H_{a3.9a} Lack of well-defined, regulated sustainability risk positively affects the adoption of generative AI solutions for business-critical applications

Refer to *Appendix B, Section C*, for the survey questions to test the hypotheses.

For Business-support Application:

H_{03.1b} Lack of well-defined, regulated accountability risk is not affecting the adoption of generative AI solutions for business-support applications

H_{a3.1b} Lack of well-defined, regulated accountability risk positively affects the adoption of generative AI solutions for business-support applications

H_{03.2b} Lack of well-defined, regulated fairness risk is not affecting the adoption of generative AI solutions for business-support applications

H_{a3.2b} Lack of well-defined, regulated fairness risk positively affects the adoption of generative AI solutions for business-support applications

H₀3.3b Lack of well-defined, regulated transparency risk is not affecting the adoption of generative AI solutions for business-support applications

H_a3.3b Lack of well-defined, regulated transparency risk positively affects the adoption of generative AI solutions for business-support applications

H₀3.4b Lack of well-defined, regulated explainability risk is not affecting the adoption of generative AI solutions for business-support applications

H_a3.4b Lack of well-defined, regulated explainability risk positively affects the adoption of generative AI solutions for business-support applications

H₀3.5b Lack of well-defined, regulated resilience risk is not affecting the adoption of generative AI solutions for business-support applications

H_a3.5b Lack of well-defined, regulated resilience risk positively affects the adoption of generative AI solutions for business-support applications

H₀3.6b Lack of well-defined, regulated privacy risk is not affecting the adoption of generative AI solutions for business-support applications

H_a3.6b Lack of well-defined, regulated privacy risk positively affects the adoption of generative AI solutions for business-support applications

H₀3.7b Lack of well-defined, regulated safety risk is not affecting the adoption of generative AI solutions for business-support applications

H_a3.7b Lack of well-defined, regulated safety risk positively affects the adoption of generative AI solutions for business-support applications

H₀3.8b Lack of well-defined, regulated security risk is not affecting the adoption of generative AI solutions for business-support applications

H_a3.8b Lack of well-defined, regulated security risk positively affects the adoption of generative AI solutions for business-support applications

H₀3.9b Lack of well-defined, regulated sustainability risk is not affecting the adoption of generative AI solutions for business-support applications

H_a3.9b Lack of well-defined, regulated sustainability risk positively affects the adoption of generative AI solutions for business-support applications

Refer to *Appendix B, Section D*, for the survey questions to test the hypotheses.

For Creativity Application:

H₀3.1c Lack of well-defined, regulated accountability risk is not affecting the adoption of generative AI solutions for creativity applications

H_a3.1c Lack of well-defined, regulated accountability risk positively affects the adoption of generative AI solutions for creativity applications

H₀3.2c Lack of well-defined, regulated fairness risk is not affecting the adoption of generative AI solutions for creativity applications

H_a3.2c Lack of well-defined, regulated fairness risk positively affects the adoption of generative AI solutions for creativity applications

H₀3.3c Lack of well-defined, regulated transparency risk is not affecting the adoption of generative AI solutions for creativity applications

H_a3.3c Lack of well-defined, regulated transparency risk positively affects the adoption of generative AI solutions for creativity applications

H₀3.4c Lack of well-defined, regulated explainability risk is not affecting the adoption of generative AI solutions for creativity applications

H_a3.4c Lack of well-defined, regulated explainability risk positively affects the adoption of generative AI solutions for creativity applications

H₀3.5c Lack of well-defined, regulated resilience risk is not affecting the adoption of generative AI solutions for creativity applications

H_a3.5c Lack of well-defined, regulated resilience risk positively affects the adoption of generative AI solutions for creativity applications

H₀3.6c Lack of well-defined, regulated privacy risk is not affecting the adoption of generative AI solutions for creativity applications

H_a3.6c Lack of well-defined, regulated privacy risk positively affects the adoption of generative AI solutions for creativity applications

H₀3.7c Lack of well-defined, regulated safety risk is not affecting the adoption of generative AI solutions for creativity applications

H_a3.7c Lack of well-defined, regulated safety risk positively affects the adoption of generative AI solutions for creativity applications

H₀3.8c Lack of well-defined, regulated security risk is not affecting the adoption of generative AI solutions for creativity applications

H_a3.8c Lack of well-defined, regulated security risk positively affects the adoption of generative AI solutions for creativity applications

H₀3.9c Lack of well-defined, regulated sustainability risk is not affecting the adoption of generative AI solutions for creativity applications

H_a3.9c Lack of well-defined, regulated sustainability risk positively affects the adoption of generative AI solutions for creativity applications

Refer to *Appendix B, Section E*, for the survey questions to test the hypotheses.

For each application category (business-critical, business support, or creativity), a quantitative analysis was conducted to understand the relationship between the perceived risk assurance and its adoption. The nine risk dimensions served as independent variables, with risk assurance and impact on adoption as dependent variables.

RQ4 specific - What copyright, privacy, bias, toxicity, and misinformation risks are introduced by the input data during the adoption for productivity and innovation use cases by IT service companies

AI-generated content using a large language model is known to have copyright, privacy, bias, toxicity, and misinformation risks introduced during the training of the model (*Chen et al., 2023*). Kandpal et al. (2022) have highlighted that a trained model stores and consequently leaks information about its training data through memorization. Since the training data used in AI models is collected in the real world, it can unintentionally reinforce harmful stereotypes, exclude or marginalize certain groups, and contain toxic data sources, which can incite hate or violence and offend individuals (*Chen et al., 2023*). AI systems may reflect or amplify societal biases present in training data (*Paul & Sarkar, 2023*). Misinformation emanating from hallucination, where large language models can generate or cite non-existent facts, is also well known.

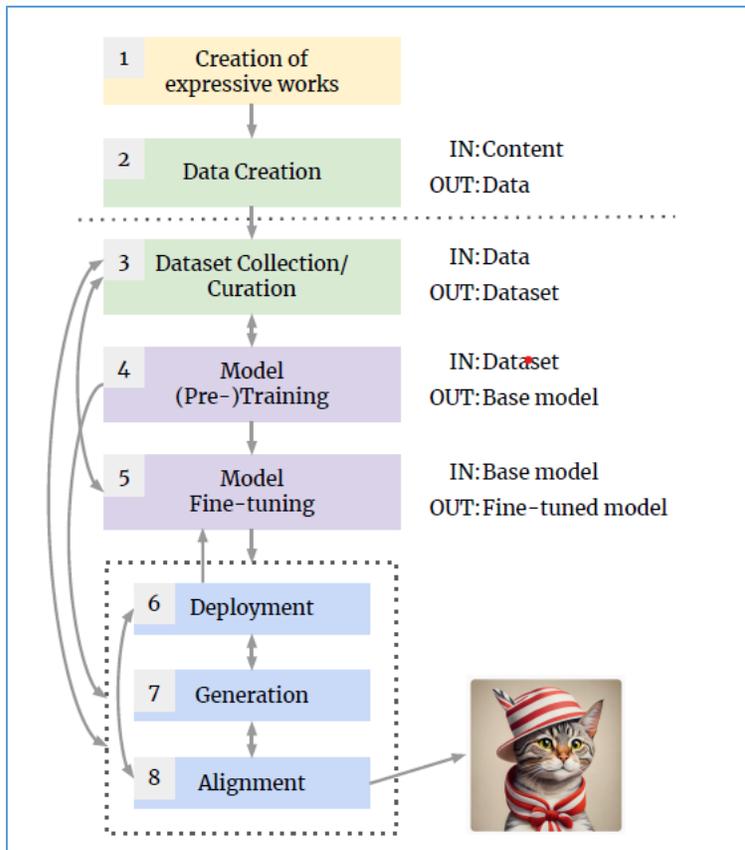


Figure 3.2: AI value chain, Source: (Paul & Sarkar, 2023)

The researcher has reviewed research articles to understand the risks of Copyright infringement, Privacy breaches, Bias, Toxicity, and Misinformation introduced during LLM training and during subsequent fine-tuning by IT service companies as they implement LLM-based generative AI solutions.

RQ5 How much input-data-generated copyright, biases, toxicity, and misinformation perceived risks are covered, and how are they affecting the adoption of productivity and innovation use cases by IT service companies?

IT service companies are implementing different business applications using generative AI for productivity and innovation (creativity), primarily for business-critical, business-support, and creativity use cases.

For each category (business-critical, business-support, and innovation), a quantitative survey analysis was conducted to assess the risk assurance of the top five risk areas and their impact on its adoption.

The following null and alternative hypotheses were developed.

For Business-critical Application:

H₀5.1: Copyright is fully covered, and the lack of copyright-specific regulations is not affecting the adoption of generative AI

H_a5.1a: Copyright is not fully covered

H_a5.1b: Lack of copyright regulations positively affects the adoption of generative AI solutions

H₀5.2: Privacy is fully covered, and the lack of privacy-specific regulations is not affecting the adoption of generative AI

H_a5.2a: Privacy is not fully covered

H_a5.2b: Lack of privacy regulations positively affects the adoption of generative AI solutions

H₀5.3: Bias is fully covered, and the lack of bias-specific regulations is not affecting the adoption of generative AI

H_a5.3a: Bias is not fully covered

H_a5.3b: Lack of bias regulations positively affects the adoption of generative AI solutions

H₀5.4: Toxicity is fully covered, and the lack of toxicity-specific regulations is not affecting the adoption of generative AI

H_a5.4a: Toxicity is not fully covered

H_a5.4b: Lack of toxicity regulations positively affects the adoption of generative AI solutions

H₀5.5: Misinformation is fully covered, and the lack of specific regulations is not affecting the adoption of generative AI

H_a5.5a: Misinformation is not fully covered

H_a5.5b: Lack of Misinformation regulations positively affects the adoption of generative AI solutions

Refer to *Appendix B, Section C*, for the survey questions.

For Business-support Application:

H₀5.6: Copyright is fully covered, and the lack of copyright-specific regulations is not affecting the adoption of generative AI

H_a5.6a: Copyright is not fully covered

H_a5.6b: Lack of copyright regulations positively affects the adoption of generative AI solutions

H₀5.7: Privacy is fully covered, and the lack of privacy-specific regulations is not affecting the adoption of generative AI

H_a5.7a: Privacy is not fully covered

H_a5.7b: Lack of privacy regulations positively affects the adoption of generative AI solutions

H₀5.8: Bias is fully covered, and the lack of bias-specific regulations is not affecting the adoption of generative AI

H_a5.8a: Bias is not fully covered

H_a5.8b: Lack of bias regulations positively affects the adoption of generative AI solutions

H₀5.9: Toxicity is fully covered, and the lack of toxicity-specific regulations is not affecting the adoption of generative AI

H_a5.9a: Toxicity is not fully covered

H_a5.9b: Lack of toxicity regulations positively affects the adoption of generative AI solutions

H₀5.10: Misinformation is fully covered, and the lack of specific regulations on Misinformation is not affecting the adoption of generative AI

H_a5.10a: Misinformation is not fully covered

H_a5.10b: Lack of Misinformation regulations positively affects the adoption of generative AI solutions

Refer to *Appendix B, Section D*, for the survey questions.

For Creativity Application:

H₀5.11: Copyright is fully covered, and the lack of copyright-specific regulations is not affecting the adoption of generative AI

H_a5.11a: Copyright is not fully covered

H_a5.11b: Lack of Copyright regulations positively affects the adoption of generative AI solutions

H₀5.12: Privacy is fully covered, and the lack of privacy-specific regulations is not affecting the adoption of generative AI

H_a5.12a: Privacy is not fully covered

H_a5.12b: Lack of privacy regulations positively affects the adoption of generative AI solutions

H₀5.13: Bias is fully covered, and the lack of bias-specific regulations is not affecting the adoption of generative AI

H_a5.13a: Bias is not fully covered

H_a5.13b: Lack of bias regulations positively affects the adoption of generative AI solutions

H₀5.14: Toxicity is fully covered, and the lack of toxicity-specific regulations is not affecting the adoption of generative AI

H_a5.14a: Toxicity is not fully covered

H_a5.14b: Lack of Toxicity regulations positively affects the adoption of generative AI solutions

H₀5.15: Misinformation is fully covered, and the lack of specific regulations is not affecting the adoption of generative AI

H_a5.15a: Misinformation is not fully covered

H_a5.15b: Lack of Misinformation regulations positively affects the adoption of generative AI solutions

Refer to *Appendix B, Section E*, for the survey questions.

3.5 Population and Sample

Sampling is the statistical process of selecting a subset of a population of interest, known as a sample, to make observations and draw statistical inferences about that population (*Bhattacharjee, 2012*). Results from a sample of the population can be used to generalize the findings (*Bhattacharjee, 2012*). Non-probability sampling techniques were used for this research, where “some units of the population have zero chance of selection or where the probability of selection cannot be accurately determined” (*Bhattacharjee, 2012*). The sample was selected based on specific non-random criteria; hence, this sampling technique did not allow for the estimation of sampling size and sampling errors (*Bhattacharjee, 2012*). The researcher employed convenience sampling and snowball sampling techniques to select the sample. Convenience sampling is a technique in which a sample is drawn from the part of the population that is readily available and convenient (*Bhattacharjee, 2012*). Snowball sampling was also employed, where a few respondents

who met the inclusion criteria were identified and asked to recommend others they knew who would also meet the selection criteria (*Bhattacharjee, 2012*).

This sampling method is criticized for introducing bias, as the selected respondents may not accurately represent the entire population. However, advantages such as speed of response, inexpensiveness, ease, and the respondents' ready availability outweigh the disadvantages and criticisms of the sampling approach.

There are approximately 480,000 IT service companies in the US (*IBISWorld, 2025*). No data suggests how many are implementing artificial intelligence solutions or using generative AI for enterprise purposes.

A rough estimate suggests that at least 20% of IT consulting companies are implementing a generative AI solution, bringing the total to nearly 100,000.

3.6 Participant Selection

The research aimed to select a sample of individuals working in the US IT industry. IT Technical Leads, Solution Architects, Business Analysts, IT Project Managers, IT Directors, and IT Engagement Leads who met the selection criteria were eligible to participate in the survey. The selection criteria for participation in the survey were determined to be (a) participants should have been involved in the implementation of generative AI solutions, or (b) participants should be aware of risks specific to generative AI. There was an additional demographic question (e.g., which industry do you identify with?) to qualify a participant to contribute to the survey. As the scope of this research was limited to IT service and IT product (**combined as IT service**) companies, respondents from non-IT service companies were excluded from the survey analysis.

3.7 Instrumentation

The unit of analysis refers to the person, collective, or object that is the target of the investigation (*Bhattacharjee, 2012*). Constructs are abstract concepts specified at a high

level and are chosen to explain the phenomenon of interest. A researcher should determine precisely how the construct will be measured and the level of analysis, either at the individual, group, or organizational level. The measurable representations of constructs are called variables (*Bhattacharjee, 2012*). “Constructs and variables proceed along two planes of scientific research: Constructs are conceptualized at the theoretical plane, while variables are operationalized and measured at the empirical (observational) plane” (*Bhattacharjee, 2012*). The research topic is exploratory and uses an IT resource as the unit of analysis.

A 72-question questionnaire was used in this study to measure dependent variables, including risk assurance and the impact of the lack of generative AI-specific regulations. Respondents were instructed to base their survey responses on generative AI projects they were involved in, which involved implementing generative AI solutions for either internal or external purposes. Three eligibility questions qualify whether a respondent’s feedback is valid for research. The online survey had five distinct sections: (a) Professional background and project related questions including six items, (b) Perceived risk not covered by existing AI regulations and frameworks including nine items, (c) Business-critical application risk coverage and impact of regulations including 18 items, (d) Business-support application risk coverage and impact of regulations including 18 items, and (e) Creativity application risk coverage and impact of regulations including 18 items. The aggregated survey data were used to validate the Operational Model (Figure 3.0) constructs using quantitative statistical methods.

A five-point Likert scale was used to quantify the data collected. The Likert scale, developed by Rensis Likert, is a widely used rating scale for measuring ordinal data in social science research (*Bhattacharjee, 2012*). This scale includes simple statements to which respondents can indicate the extent of agreement or disagreement on a five-point

scale ranging from “strongly disagree” to “strongly agree” (*Bhattacharjee, 2012*). This allows for more granularity, including whether respondents are neutral to the statement (*Bhattacharjee, 2012*).

3.8 Data Collection Procedures

The questionnaire was initially pretested through a Pilot study with five personal connections of the author, who met the selection criteria and have extensive experience in implementing generative AI solutions. This was done through a self-administered online Google Survey. The pilot study participants were asked to evaluate the individual survey items based on their readability and the ease with which they understood the questions. It helped to uncover the ambiguity, lack of clarity, and biases in the wording of the question and statement. The wordings of some survey questions was rephrased to incorporate the feedback.

The research data were gathered through a self-administered online survey. The survey was initiated via a link in the invitation, redirecting the participants to a Google Survey.

An online survey was shared with potential participants via the author’s LinkedIn network, email, and personal connections. A standardized message and a cover letter/email/WhatsApp message were used, including a hyperlink to the survey instrument. The cover letter/email/WhatsApp message included a description of the research and its purpose, emphasizing that it is intended for academic purposes. The results would be reported on an aggregate basis. It was mentioned that participation would be anonymous, confidential, and voluntary, and that access to individual data would be limited to the researcher. The cover letter/email also included a request for respondents to forward the message with the survey link to their connections who met the selection criteria, to snowball the sample size.

3.9 Data Analysis

The primary numerical data collected through the Google Survey was exported to Excel. The data were coded numerically to prepare for import into IBM SPSS Statistics Version 31.0 (SPSS) as per the codebook prepared for this research. The survey was conducted over two months, July and August 2025. The researcher contacted approximately 300 IT service practitioners through social media (LinkedIn contacts), AI discussion forums, and contacts from consulting companies. A total of 74 survey responses were received after multiple reminders. Four responses didn't qualify as valid reactions because they lacked awareness of the risks associated with generative AI or implemented no generative AI solutions. Some of them selected "Other" as their work domain. Since the analysis was limited to practitioners in the IT service industry, anyone who selected "Other" as their work domain was excluded. Six questions were asked to gather professional background information about the survey respondents. These six questions were used to determine whether respondents had varied experience (both in terms of years and the type of implementation) in AI project implementation. A good mix of AI years of experience reflects high integrity in survey responses and subsequent analysis. These six demographic questions were not used in any subsequent statistical analysis, as the study did not attempt to correlate them with respondents' responses.

During the survey phase, five individuals reached out to confirm their participation because they worked in non-IT service industries. The researcher discouraged them from participation as the survey was limited to IT services only.

The survey data were imported into SPSS for quantitative analysis. The data integrity checks ensured that no data was missing. The frequencies and descriptive statistics were calculated to confirm that the data met the preconditions for descriptive and inferential analysis. Data normality was measured using skewness and kurtosis.

A variety of analyses (Descriptive, Inferential, Regression, EFA, ANOVA, and MANOVA) were conducted to identify the relationships between the independent and dependent variables. These analyses helped test the formulated hypotheses and address the research questions by examining the relationships between the independent variables (risk dimensions) and the dependent variables (risk assurance and impact on adoption). Given the nature of some hypotheses, multivariate analysis of variance (MANOVA) was found to be an appropriate data analysis technique, as it enabled the simultaneous evaluation of multiple independent variables on a single dependent variable (*Green, 2010*).

Research Reliability and Validity

Any research based on the measurement of variables must be concerned with accuracy and dependability (*Bhattacharjee, 2012*). The reliability of an instrument characterizes its consistency and dependability (*Bhattacharjee, 2012*). The research topic was established during the research proposal stage and approved by the assigned academic mentor, in accordance with the criteria set by the Swiss School of Business and Management (SSBM) in Geneva. The literature review conducted for this research, along with insights from various theoretical and empirical studies and the author's experience in data analytics and technology consulting, has supported the development of the Operational Model (Figure 3.0).

A reliability test was conducted in SPSS on the survey responses to ensure the reliability and consistency of the data for the study (*Bhattacharjee, 2012*). The details of these reliability tests are documented in the results section, corresponding to the relevant research questions. Research question 3 (RQ3) used nine independent risk variables for analysis. The survey response size was 70 (after excluding four responses that were disqualified), which is on the lower end of the sample size range, ideal for multivariate

analysis of variance. There should be at least 10 samples per predictor (*Green, 2010*). The output of multivariate analysis in the research question 3 has limitations, and the researcher wants to highlight them. Research question 5 (RQ5) uses five independent risk variables, and the collected valid sample size (70) is well above the minimum threshold. Research question 5 is the crux of this research, and the survey sample size justifies advanced analyses, such as multivariate analysis used in this study.

There are three common approaches to assessing reliability and comparing collected data with data from other sources (*Saunders et al., 2023*). Although each analysis is undertaken after data collection, it needs to be considered at the questionnaire design stage. They are internal consistency, alternative form, and test-retest (*Saunders et al., 2023*). Both Alternative Form and Test-retest appear impractical for the intended survey, given the survey audience's wide distribution across the US. Internal consistency involves correlating responses to questionnaire items (*Saunders et al., 2023*). There are various methods for calculating internal consistency; one of the most frequently used is **Cronbach's alpha**. This statistic is typically used to assess the consistency of responses to a subset of questions (scale items) combined into a scale to evaluate a particular concept (*Saunders et al., 2023*). A reliability coefficient demonstrates whether the survey instrument design accurately measures the studied variables and whether the items yield interpretable statements about individual differences (*Bhattacharjee, 2012*).

Ethical research involves protecting participants from harm that might result from activities and findings associated with the research project (*Bhattacharjee, 2012; Creswell & Creswell, 2022*).

Before proceeding with the online survey, respondents were provided with information on the survey's purpose, the research details, and the study's contribution. The online survey explicitly stated that it was anonymous and that the results would be

aggregated. Respondents were assured of confidentiality and that their identities would remain anonymous throughout the research process. Respondent anonymity and confidentiality concerns were addressed at the onset. The survey administered in this study did not collect any sensitive personal information from respondents, such as names, addresses, dates of birth, or employers, to ensure confidentiality and mitigate the risk of potential conflicts of interest.

3.10 Research Design Limitations

A few assumptions were made during the design of this research. The study assumed that the implementers of generative AI solutions are aware of the associated risks and understand the risk management framework in place. This was validated through the qualifying questions. The study also considered that the survey instrument is relevant to the US context and that respondents would have no difficulty interpreting the survey items to respond correctly. It is assumed that all respondents provided honest answers.

It is assumed that all received data was as unbiased as possible and helpful in adequately answering the research questions. Some strategies to address common biases were considered during the research design. To minimize maturation bias (time-effect bias), data on all independent and dependent variables were collected simultaneously via a single 15-minute online survey. All construct measures use the same 5-point Likert scale to improve consistency.

The research has some limitations due to the research design. Non-probability sampling techniques, including convenience and snowball sampling, were employed in this research. Some limitations of this sampling method include the potential for bias and the possibility that selected respondents may not accurately represent the entire population (*Bhattacharjee, 2012*). This introduces limitations, including a lack of generalizability, selection bias, non-randomness, and limited external validity. Some strategies were

employed in this research to improve the validity of non-probability sampling. A larger sample size mitigates the biases and accurately reflects the population. A common rule of thumb is to have at least ten to fifteen responses per predictor or independent variable (Green, 2010). The researcher has chosen a non-probability sampling technique to provide information-rich rather than statistical explanations and to reveal understandings and insights (Saunders et al., 2023). Non-probability sampling can yield theoretical generalizations based on analytic generalizability, while probability sampling can yield statistical generalizations about a target population (Saunders et al., 2023). More than one non-probability sampling technique, namely convenience and snowball sampling, is used to capture a broader range of participants. The survey includes key subgroups of the population, such as IT Technical Leads, Solution Architects, Business Analysts, IT Project Managers, IT Directors, and IT Engagement Leads, to make the sample more representative. A Pilot survey was conducted to get feedback on readability and ease of understanding the questions.

3.11 Conclusion

Selecting an appropriate research methodology was crucial to ensuring a reliable and valid study. The research questions require a mixed-methods approach. Research questions 1 and 4 required qualitative analysis based on literature reviews, peer-reviewed research papers, and articles. Research questions 2, 3, and 5 required a deductive, descriptive, quantitative study with correlation and regression analysis, utilizing both analysis of variance and multivariate analysis of variance tests. It was used to assess the impact of the independent variables (risk dimensions) on the dependent variables (risk assurance and implications for adoption) in the US IT service industry. This research design was deemed the most appropriate for this analysis, as it defined the variables and their relationships, as illustrated in the Research Model (Figure 3.0). This was assessed through

data collected from a sample meeting the specified selection criteria through the online survey.

CHAPTER IV:

RESULTS

4.0 Survey Participants' Professional Background Analysis

Before analyzing the survey data for risk correlation, they were thoroughly assessed for suitability and integrity. There were nine questions about participants' professional backgrounds. Out of nine questions, three questions were used to qualify whether a participant should be considered for research analysis purposes. These three questions were:

1. Have you implemented or participated in any generative AI solution?
2. Are you familiar with the risks (e.g., privacy, bias, copyright, transparency, etc.) associated with generative AI?
3. Which industry do you identify with your work domain?

If a participant responded “No” to either of the first two questions or “Other” to the 3rd question, the response was excluded from further risk analysis. Out of 74 total responses, four responses were disqualified. So, all subsequent risk analyses included only 70 respondents.

The other six questions were used to gauge the integrity of the collected data. Refer to *Appendix A, Section B (Questions 4-9)*. Although the responses to these six questions were not directly used in the analysis, they reinforced the high integrity of the responses.

Table 4.0: Count of Participants vs. Years of AI Experience (Source: author's survey Excel output)

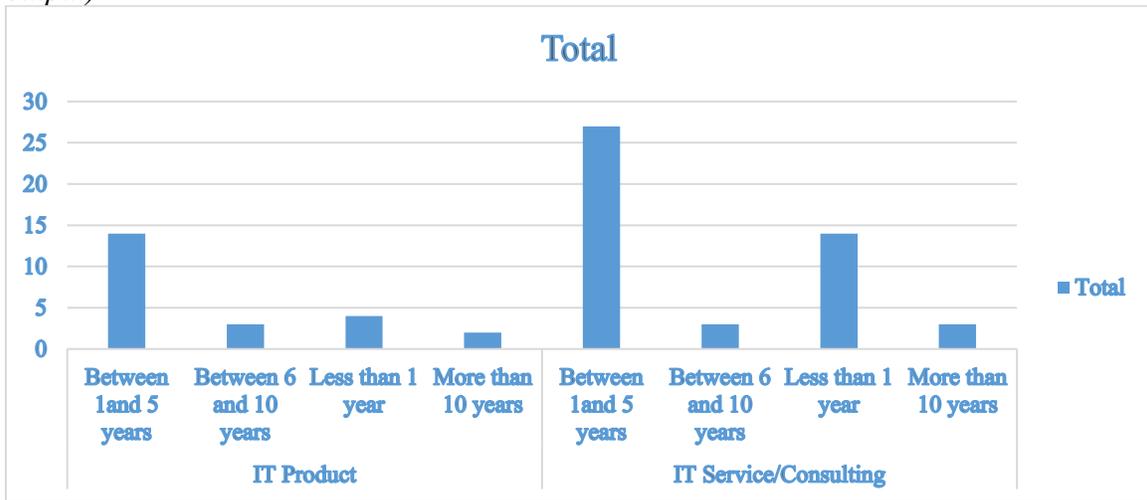
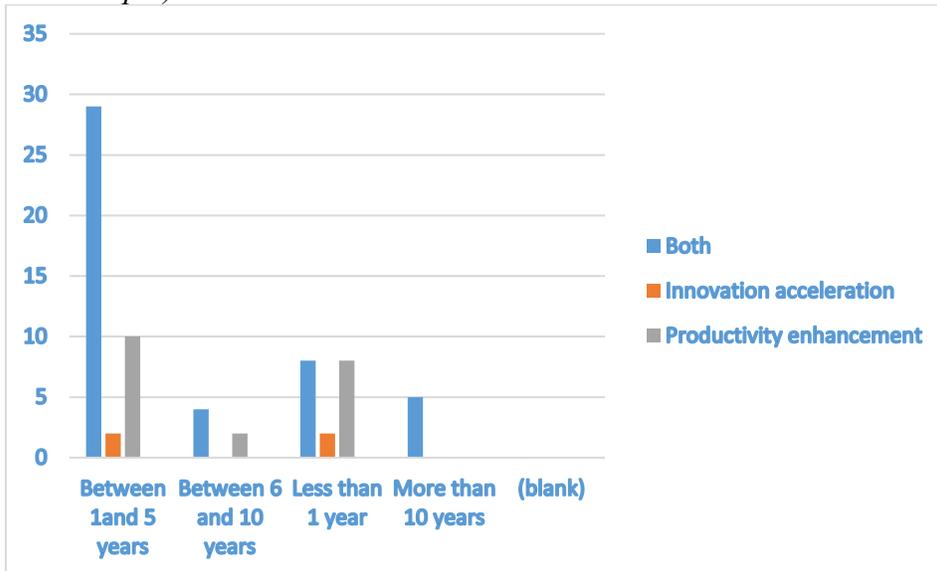


Table 4.0 (above) indicates that most valid responses originated from participants with more than 1 year of AI experience. 18 out of 70 participants had less than 1 year of AI experience.

Table 4.1: Years of AI Experience vs. Gen AI Application Purpose (Source: author's survey Excel output)



Another insight, as shown in Table 4.1, is that participants with less than 1 year of AI experience have implemented only a few use cases for creativity (innovation acceleration) with generative AI. Participants with less experience have been involved in productivity-enhancement use cases. Additionally, most participants have implemented both productivity and innovation acceleration use cases and have multiple years of experience. This further supports the survey result, as it didn't appear biased toward less experienced participants.

Table 4.2: Years of AI Experience vs. Business Critical Application Implementation (Source: author's survey Excel output)

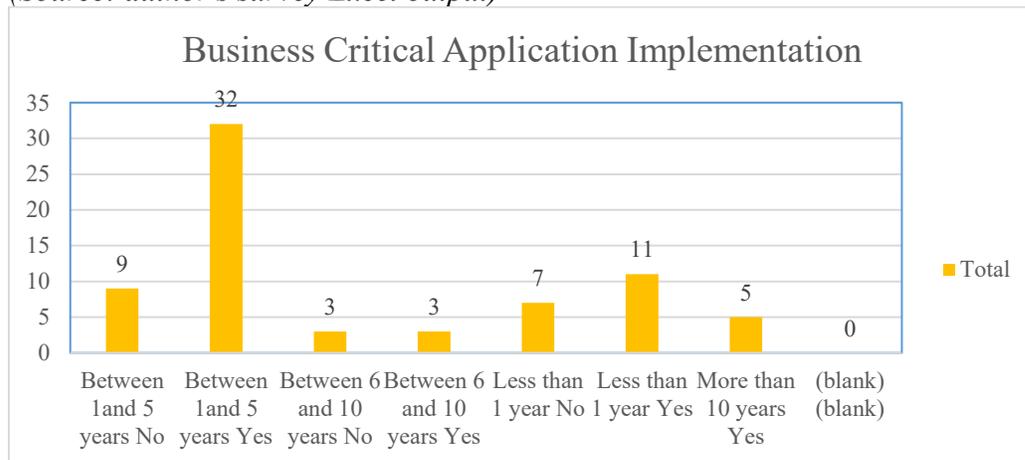


Table 4.2 presents the participants' years of experience with AI and their experience implementing business-critical applications. The largest group (32) of participants with 1-5 years of AI experience has implemented a business-critical application.

Table 4.3: Years of AI Experience vs. Business Support Application Implementation (Source: author's survey Excel output)

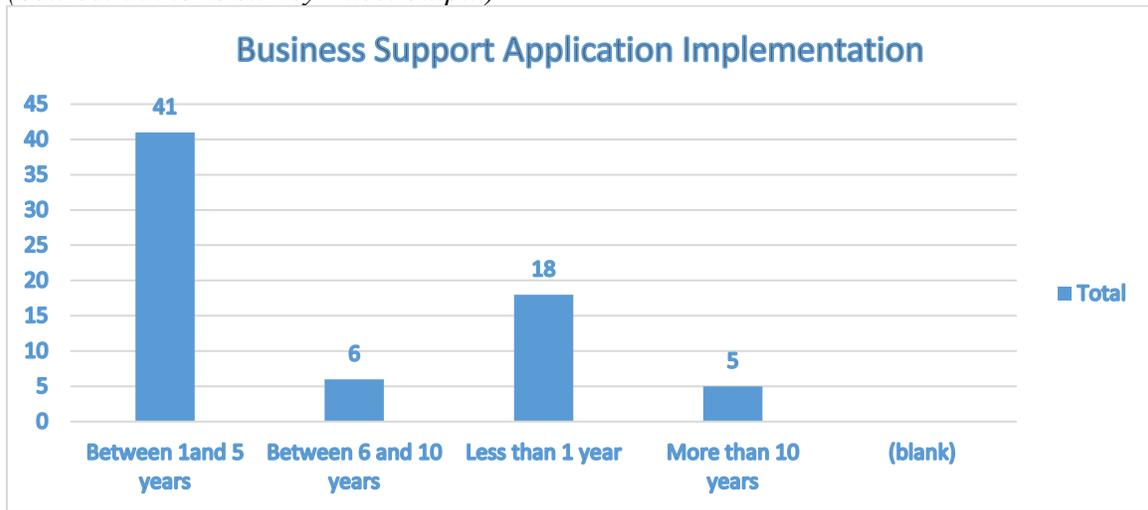


Table 4.3 presents the participants' years of experience with AI and their implementation experience with the business-support application. The largest group (41) of participants with 1-5 years of AI experience has implemented a business-support application.

Table 4.4: Years of AI Experience vs. Creativity Application Implementation (Source: author's survey Excel output)

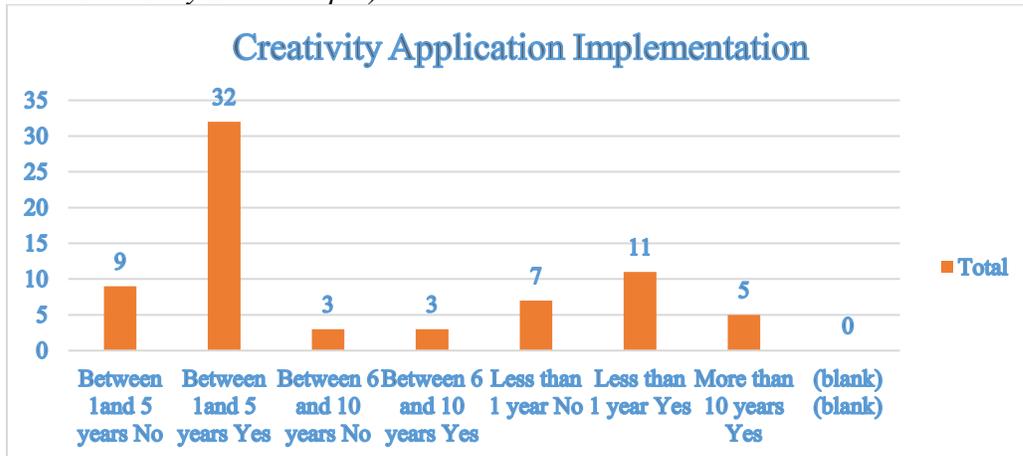


Table 4.4 presents the participants' years of experience with AI and their experience implementing the creativity application. The largest group (32 participants) with 1-5 years of AI experience has implemented a creativity application.

Overall, the survey data were analyzed for suitability and integrity before use in subsequent risk modeling. There was no red flag for the suitability or integrity of the data. The research did not involve any behavioral analysis based on demographic or professional background. No demographic information (age, sex, location, etc.) was collected, as it was not planned for analysis, as mentioned in the research methodology.

4.1 Research Question One

RQ1 What is the current approach of the US, the UK, and China in regulating generative AI? How do Risk Management frameworks (e.g., NIST RMF) and the US government AI Acts relate to Generative AI?

Secondary source validation and authentication were done for all the reference sources used in this study. Before including any prior citations used during the literature review, they were reconfirmed as existing. Another step was to avoid using any blog site for citation. Only credible sources were used for citations.

A detailed analysis of the approach used by the US, the UK, and China was conducted. A thorough review of various secondary sources was conducted. The findings are given below.

The US Approach: The United States currently lacks a comprehensive federal law that establishes broad regulatory authority over the development or use of AI. No federal prohibitions on AI have been enacted (*Harris, 2025*). The U.S. regulatory strategy is not centered on a single, overarching statute but instead relies on a combination of other mechanisms (*Harris, 2025*). In the absence of sweeping legislation, federal efforts are primarily focused on three key areas:

1. Federal agency assessments and enforcement using existing regulatory authorities (*Harris, 2025*).

2. Exploration of whether individual agencies require additional authorities (*Harris, 2025*).
3. Securing voluntary commitments from the industry (*Harris, 2025*).

The overall approach of the U.S. federal government appears more focused on overseeing its own use of AI than on directly regulating the private sector (*Harris, 2025*). The executive branch has been the primary driver of federal AI policy across multiple administrations. Through executive orders, both the Trump and Biden administrations have shaped a policy aimed at maintaining American leadership in AI innovation. Key examples include President Trump's Executive Order 13859, "Maintaining American Leadership in Artificial Intelligence," (*The President of the United States, 2019*) and President Trump's 2025 Executive Order 14179, "Removing Barriers to American Leadership in Artificial Intelligence" (*The White House, 2025b*).

A literature review indicates there are no stringent federal regulations specific to generative AI, reflecting a desire to facilitate the technology's advancement (*Harris, 2025*). Congressional proposals have largely emphasized voluntary measures and industry self-assessment rather than strict prohibitions. The proposed bills in the 118th and 119th Congresses have focused on "the development of voluntary guidelines and best practices, as well as the reporting of industry-conducted evaluations of AI systems" (*Harris, 2025*).

There is some attempt to regulate AI technologies directly, either through technical thresholds or transparency mandates.

Technical Thresholds: Some executive actions, such as Executive Order 14110, require companies developing powerful dual-use AI models to report to the government if their models exceed a specific computational training threshold (e.g., greater than 10^{26} FLOPs). While initially set high, this threshold has already been surpassed by models like xAI's Grok-3, illustrating the potential rigidity of such an approach (*Harris, 2025*).

Transparency Requirements: Legislative proposals like the AI Disclosure Act of 2023 would have mandated that any output from generative AI include a disclaimer stating it was AI-generated (*Harris, 2025*).

The Algorithmic Accountability Act: This proposed legislation would have directed the Federal Trade Commission (FTC) to require impact assessments from large companies using automated decision systems. It aimed to mitigate adverse impacts on consumers and create a public repository of information, without containing specific prohibitions on use (*U.S. Congress House, 2022*).

A significant volume of legislative activity targets high-stakes sectors, proposing tailored governance. Some examples are listed below.

Financial Sector: Bills like the Preventing Deep Fake Scams Act focus on AI-related financial crimes and scams (*U.S. Congress Senate, 2025*).

Elections and Campaign Finance: Examples include the Fraudulent Artificial Intelligence Regulations Elections Act of 2024 (*U.S. Congress Senate, 2024a*)

Healthcare: Legislation such as S. 4862 was introduced to ensure the ethical adoption of AI in healthcare (*U.S. Congress Senate, 2024b*).

In the absence of holistic federal AI legislation, states have become increasingly active in regulating AI. As of late April 2025, at least 48 states and Puerto Rico had introduced over 1,000 AI-related bills, creating a potential patchwork of regulations (*The White House, 2025a*).

The federal government has expressed concern about this trend. The White House AI Action Plan explicitly states that "AI is far too important to smother in bureaucracy at this early stage, whether at the state or Federal level" and includes efforts to prevent states from enacting "burdensome" regulations that could hinder innovation. The plan suggests

that federal AI funding may not be directed to states with overly restrictive laws (*The White House, 2025a*).

The White House AI Action Plan and the Central Role of NIST: The White House's vision for AI policy is based on a voluntary, standards-based approach, with the National Institute of Standards and Technology (NIST) playing a central role.

The plan directs NIST to revise its framework to remove references to "misinformation, diversity, equity, and inclusion, as well as climate change," arguing that government-procured AI must reflect "truth rather than social engineering agendas" (*The White House, 2025a*).

The plan tasks NIST with developing national standards for AI in specific sectors and promotes the establishment of regulatory sandboxes and "AI Centers of Excellence" to help businesses test AI tools before market deployment (*The White House, 2025a*).

It calls for improved cyber incident response protocols and evaluation of national security risks from frontier AI models (*The White House, 2025a*).

The NIST RMF remains a comprehensive voluntary framework that "encompasses all aspects of risk management, including generative AI systems" (*G'sell, 2024*).

Conclusion: Unlike the European Union or China, the United States has not implemented a comprehensive, mandatory federal framework for AI governance. Instead, the federal government's strategy is defined by a voluntarist model, engaging in dialogue with major AI companies to secure commitments and encourage adherence to voluntary standards set by agencies like NIST (*G'sell, 2024*). This approach prioritizes innovation and risk management through existing authorities and industry partnerships rather than blanket federal legislation.

The UK's Approach to AI Regulation: The United Kingdom has taken a unique approach to global AI governance. The UK has opted for a non-statutory, pro-innovation

path. The core idea of the UK's strategy is to modify the existing sector-based regulatory framework to address AI challenges, instead of establishing a new centralized AI regulator or a comprehensive new law. This approach aims to be flexible and context-aware while mitigating potential risks. The UK's framework is a “principles-based, non-statutory, and cross-sector framework” that seeks to “balance innovation and safety by applying the existing technology-neutral regulatory framework to AI” (*Harris, 2025*).

AI is recognized as not a single technology but a diverse set of tools applied across different sectors. Risks and regulatory needs vary significantly across different sectors. The UK does not have a general statutory framework for AI. Instead, AI is primarily regulated through the existing legal frameworks of the sectors in which it is used (*Rough & Sutherland, 2024*). This choice has been developed through a national AI strategy, policy white papers, and the details have evolved under different governments.

The foundation of the UK's current regulatory framework is the March 2023 policy white paper, “A Pro-Innovation Approach to AI Regulation.” This document has laid out the government’s intention to “put in place a new framework to bring clarity and coherence to the AI regulatory landscape” (*Department for Science, Innovation and Technology, 2023a*). The rapid pace of technological change has been recognized, so the framework is designed to be agile and iterative.

There are five core-sectoral principles to guide the development and use of AI technologies. These principles are not intended to be new, standalone rules. They are meant to be interpreted and applied by the UK’s sectoral regulators in their respective contexts. The five principles are:

Safety, Security, and Robustness: AI systems should function in a secure, safe, and robust manner where risks are carefully managed (*Department for Science, Innovation and Technology, 2023a*).

Appropriate Transparency and Explainability: Transparency and appropriate levels of explainability are required to enable users to understand AI-driven outcomes and build trust (*Department for Science, Innovation and Technology, 2023a*).

Fairness: AI systems should be used in a way that complies with the UK's existing laws, such as the Equality Act 2010, and must not discriminate against individuals or create unfair commercial outcomes (*Department for Science, Innovation and Technology, 2023a*).

Accountability and Governance: Measures are needed to ensure there is appropriate oversight of AI systems and clear lines of accountability (*Department for Science, Innovation and Technology, 2023a*).

Contestability and Redress: When an AI system makes a decision or supports a decision that has a material impact on an individual, organization, or community, there should be clear routes to dispute or challenge the outcome (*Department for Science, Innovation and Technology, 2023a*).

An essential feature of the initial proposal was that these principles would be issued on a non-statutory basis. This means they were not encoded in law. Instead, regulators such as the Health and Safety Executive (HSE), the Equality and Human Rights Commission (EHRC), the Competition and Markets Authority (CMA), the Financial Conduct Authority (FCA) were expected to apply these principles using their existing statutory powers (*Department for Science, Innovation and Technology, 2023a*). This approach seems to provide flexibility while avoiding placing new legal burdens on businesses during the early stages of the technology development and implementation. However, the government has indicated that this was likely an interim measure. After a review period, it is anticipated that it will move to a statutory duty requiring regulators to have due regard to the principles (*Department for Science, Innovation and Technology,*

2023a). This would give the core principles greater legal weight and also ensure a more consistent application across different regulatory domains and sectors.

The UK's model also relies on central coordination across different sectoral regulators. The government explicitly rejected the idea of creating a single, new AI regulator. Instead, the plan has defined responsibility for both sectoral regulators and central support functions.

Sectoral Regulators: The primary responsibility for implementing the principles falls to existing regulators. For example, the Medicines and Healthcare products Regulatory Agency (MHRA) would oversee AI in medical devices. The FCA would regulate AI in financial services, and Ofcom would address AI in communications and broadcasting.

Central Support Functions: To ensure a coherent approach and sharing lessons learned across the sectors, the UK government has committed to providing central support functions. These functions, outlined in the white paper, are crucial to the framework's success. Some key responsibilities assigned to the central support function are listed below.

Monitoring and Evaluation: Assessing the effectiveness of the framework and how well regulators are implementing the principles (*Department for Science, Innovation and Technology, 2023a*).

Risk Assessment: Monitoring and assessing risks across the economy arising from AI (*Department for Science, Innovation and Technology, 2023a*).

Horizon Scanning and Gap Analysis: Identifying emerging AI technology trends and potential regulatory gaps, often in collaboration with industry (*Department for Science, Innovation and Technology, 2023a*).

Support for Innovation: Supporting regulatory sandboxes and testbeds to help innovators navigate the regulatory landscape and bring new technologies to market (*Department for Science, Innovation and Technology, 2023a*).

Education and Awareness: Providing clarity to businesses and empowering citizens (*Department for Science, Innovation and Technology, 2023a*).

International Interoperability: Promoting alignment with international regulatory frameworks to reduce trade barriers for UK businesses (*Department for Science, Innovation and Technology, 2023a*).

The goal is to have expert-led regulation through a sectoral approach, with coordination, strategic oversight, and support from a central government function. The UK's framework is deliberately technology-agnostic and principles-based, meaning it applies to all forms of AI, *including generative AI*. The principles provide a flexible toolkit for regulators to address the unique challenges posed by generative AI systems.

Safety, Security, and Robustness: A regulator might focus on the potential for generative AI to produce harmful content (e.g., cyberattack code), its vulnerability to data poisoning attacks, or the systemic risks associated with competent foundation models (*Department for Science, Innovation and Technology, 2023a*).

Transparency and Explainability: This principle is key to addressing the "black box" nature of many generative AI systems. Regulators could encourage or mandate transparency about when AI is being used, the data it was trained on, and its limitations. Watermarking AI-generated content would be an example of a transparency measure falling under this principle (*Department for Science, Innovation and Technology, 2023a*).

Fairness: This would require addressing biases in training data that could lead generative AI to produce discriminatory or unfair outputs. A regulator like the EHRC would

be concerned with how generative AI might exacerbate societal biases in hiring or lending decisions (*Department for Science, Innovation and Technology, 2023a*).

Accountability and Governance: This principle pushes organizations deploying generative AI to have clear human oversight, governance structures, and accountability mechanisms (*Department for Science, Innovation and Technology, 2023a*).

Contestability and Redress: For individuals affected by a decision informed by generative AI, this principle ensures there is a mechanism to appeal. For instance, if a generative AI tool is used to summarize a job application and filters out a qualified candidate, that candidate should have a clear path to challenge the outcome (*Department for Science, Innovation and Technology, 2023a*).

By not creating specific rules for generative AI, the framework aims to remain flexible in the face of future technological surprises. The UK's approach is designed to evolve.

A private member bill proposes to establish an AI Authority, a central body tasked with ensuring alignment of approach across relevant regulators and accrediting independent AI auditors (*UK Parliament, 2025*). This aligns with the government's earlier suggestion of imposing a statutory duty on regulators and creating a more powerful central function.

The UK appears to be following a "test and learn" approach, starting with non-statutory principles. This allows for gathering evidence on their effectiveness and then for legislating based on that experience. This iterative process aims to avoid premature regulation that could hinder innovation. While still building the necessary guardrails for responsible AI development and deployment, the highlight is the context-specific, sector-led initiative.

Conclusion: In summary, the United Kingdom's approach to regulating AI, including generative AI, is defined by its principles-based, sector-led model. It's an attempt to develop a dynamic AI ecosystem by leveraging the UK's existing regulatory system. It empowers sectoral regulators to apply AI regulatory principles within their respective domains. A central support for coordination will come from the central government. The framework aims to manage risks in a targeted and context-aware manner. There is no one-size-fits-all approach that can stifle innovation. The evolution of this approach will be the potential introduction of an AI Authority and a statutory duty for regulators. This will be an evidence-based governance model that can adapt as technology and its implications continue to change.

China's top-down, vertical, and state-led model: In contrast to the US and UK, China has moved to implement binding, targeted regulations for generative AI. Its approach is characterized as vertical, highly reactive, adaptive, and tailored to specific technological applications or scenarios (*Migliorini, 2024*). This allows regulators to introduce legally binding rules quickly to address emerging issues (*Zou & Zhang, 2025*). The Chinese Government is trying to balance the dual objectives of development and security. China is prioritizing avoiding technological backwardness while balancing the dual objectives. As Professor Liming Wang noted, the most significant risk for China is falling behind technologically (*Rui & Liu, 2023*). This fear of technological dependence on Western rivals is a powerful driver of policy, shaping a regulatory environment that is restrictive in some areas but permissive in others.

The Interim Measures for Generative AI Services: The cornerstone of China's regulation is the Interim Measures for the Management of Generative AI Services, enacted in July 2023. The Interim title signifies a testing phase before formulating higher-level laws, a common practice in Chinese governance that allows for rapid iteration and

adjustment based on real-world feedback (*Zou & Zhang, 2025*). The Measures apply only to public-facing generative AI services, explicitly excluding research and development (R&D) and non-public industrial applications. This creates a sandbox environment for core technology development while strictly controlling public deployment, effectively ring-fencing potential social and political risks (*Creemers, 2021*).

The Measures impose a dual set of obligations on generative AI service providers and organizations that offer public services, similar to companies such as OpenAI and Google. *Content Liability* - Providers bear direct liability for generated content, which must align with socialist core values and must not contain content that threatens national security, promotes subversion, terrorism, ethnic hatred, violence, or pornography. This aligns with China's long-standing "cyber sovereignty" doctrine, which asserts the state's right to control and manage internet content within its borders (*Zou & Zhang, 2025*). *Technical Service Liability* - Providers must ensure lawful data sources and foundational models, enhance training data for "truthfulness, accuracy, objectivity, and diversity," safeguard user data, and report unlawful content (*Zou & Zhang, 2025*). The requirement for truthfulness and objectivity in training data is significant, as it implicitly directs model training towards state-approved information sources and narratives (*Zou & Zhang, 2025*).

This approach is a deliberate strategy. The government aims to promote the development of national foundational model companies (e.g., Baidu, Alibaba) by granting them regulatory space to innovate. At the same time, they want to ensure that any public application of their technology remains within strict ideological and social boundaries. This reflects a broader industrial policy in which regulation is used as a tool to shape market outcomes and steer technological development toward national strategic goals (*Cheng & Zeng, 2023*).

The Measures are part of a broader, evolving ecosystem of vertical AI regulations that form a comprehensive governance system. This framework aims to cover the entire lifecycle of algorithmic systems (Roberts, 2020). Some key provisions and regulations are listed below.

The Algorithmic Recommendations Provisions: These provisions focus on controlling algorithms that curate and recommend content on platforms such as Toutiao and WeChat (Zou & Zhang, 2025). They introduce obligations for platforms to intervene in content recommendations to promote positive energy and prevent the spread of harmful information (Zou & Zhang, 2025). Crucially, they grant users specific rights, including the ability to turn off algorithmic recommendation services and the option to delete tags used for profiling, representing a form of controlled user empowerment (Zou & Zhang, 2025).

The Deep Synthesis Regulations: These regulations address the risks of synthetic media, using the technically neutral term deep synthesis technology instead of the politically charged deepfakes (Interesse, 2022). They require clear labeling and disclosure of synthetically generated content, especially in news and information services. This allows the state to manage the potential for disinformation and reputational harm while avoiding a blanket condemnation of a useful technology (Interesse, 2022).

By regulating specific applications, such as recommendation algorithms, synthesis services, and public-facing generative AI, separately, Chinese regulators can apply targeted pressure points without necessarily stifling underlying innovation (Migliorini, 2024). This is a key advantage of the vertical approach. It allows for rapid, targeted intervention in response to specific problems, such as the spread of deepfakes, without requiring a sweeping, horizontal law that might be slow to draft and difficult to amend (Migliorini, 2024). China's AI regulatory strategy is intertwined with its national industrial policy. China wants to be the world's primary AI innovation center by 2030, underscoring that

regulation is not merely about risk mitigation but is subservient to the larger objective of technological sovereignty and global leadership (*Webster et al., 2017*). This explains the explicit exemption of R&D from the most stringent provisions of the Measures. The state is actively fostering a competitive domestic AI industry that can reduce reliance on foreign technology, particularly from the US (*Webster et al., 2017*).

The regulatory focus on content control is a direct extension of the Chinese Communist Party's (CCP) concern for ideological security. Generative AI is perceived as a potential threat to the state's monopoly on information and narrative control. The obligations placed on service providers effectively outsource the responsibility for content censorship to the private sector in China's internet governance, often referred to as regulated self-regulation (*Creemers, 2021*). Companies must build compliance into their systems from the ground up, often employing large teams of human moderators and developing sophisticated AI-powered censorship tools to filter outputs. Data governance is another critical layer. The Measures' requirement for lawfully obtained data sources aligns with China's stringent data security laws, including the Personal Information Protection Law (PIPL) and the Data Security Law (DSL). These laws establish a categorization-based data governance system, in which specific types of data are designated as core national security assets (*Cheng & Zeng, 2023*). This legal environment influences how AI companies can collect, train on, and transfer data. This will create a wall for AI development that both protects domestic companies and acts as a barrier for foreign competitors. The push for high-quality data can also be interpreted as an effort to ensure that AI models are trained on corpora that reflect state-sanctioned knowledge and perspectives (*Cheng & Zeng, 2023*).

In conclusion, China's approach to generative AI is not only about restrictive control. It is a sophisticated, multi-layered strategy that employs regulation as a strategic

tool. It balances the need for rapid technological development with the non-negotiable demands of political security and social stability. By using a vertical, reactive model, the state can precisely manage the public-facing risks of AI. At the same time, it allows public companies to pursue research and development without the regulatory burden.

4.2 Research Question Two

RQ2 How much perceived risk is not covered by existing regulations and frameworks?

The core objective was to measure the perceived risk coverage (adequacy) of the current regulatory and standards framework (AI regulations plus NIST RMF) in covering the unique risks posed by generative AI. The output is a quantified risk assurance gap analysis. This can help identify which risk domains are perceived as well-covered and which are not, allowing policymakers and organizations to prioritize their efforts accordingly.

Independent Variable (IV): Risk Domain. This is a categorical variable with nine risk levels. 1. Accountability Risk, 2. Fairness Risk, 3. Transparency Risk, 4. Explainability Risk, 5. Resilience Risk, 6. Privacy Risk, 7. Safety Risk, 8. Security Risk, 9. Sustainability Risk

Dependent Variable (DV): Perceived Risk Coverage Adequacy. This was the ordinal-scale response collected for each risk domain, measured on a 5-point Likert scale (1 = Strongly Disagree to 5 = Strongly Agree) regarding whether the risk is fully covered.

Analytical Approach - The goal was to determine whether the mean perceived risk coverage (adequacy) score differs significantly across the nine risk domains or whether their combined mean is above a neutral score (3). A mean score above 3 signifies that the respondents believe risk assurance gap exists for the risk dimensions.

Descriptive Statistics was used for each of the nine risk domains to calculate the mean, median, mode, and standard deviation for the perceived risk coverage score. A higher mean score (e.g., above 3) for a specific risk (e.g., Security) suggests that survey respondents believe existing frameworks do not adequately cover it.

Inferential Statistics was used to determine if the differences in mean scores between the risk domains are statistically significant (i.e., not due to random chance). A repeated-measures ANOVA was used. Repeated measures were used because the survey respondents' responses pertain to the same dependent variable (perceived risk coverage) under various conditions (the nine risk domains).

Table 4.5: Descriptive Statistics (Research Question 2, Nine Risk Domains), Source: author's SPSS output, 2025

	Descriptive Statistics								
	N Statistic	Minimum Statistic	Maximum Statistic	Mean Statistic	Std. Deviation Statistic	Skewness		Kurtosis	
						Statistic	Std. Error	Statistic	Std. Error
Accountability risk not fully covered under existing Reg/RMFs	70	1	5	3.46	1.125	-.550	.287	-.281	.566
Fairness risk not fully covered under existing Reg/RMFs	70	1	5	3.29	.980	.056	.287	-.685	.566
Transparency risk not fully covered under existing Reg/RMFs	70	1	5	3.34	1.128	-.157	.287	-.616	.566
Explainability risk not fully covered under existing Reg/RMFs	70	1	5	3.40	1.172	-.169	.287	-.971	.566
Resilience risk not fully covered under existing Reg/RMFs	70	1	5	3.27	1.048	.050	.287	-.689	.566
Privacy risk not fully covered under existing Reg/RMFs	70	1	5	3.53	1.316	-.437	.287	-.970	.566
Safety risk not fully covered under existing Reg/RMFs	70	1	5	3.47	1.139	-.170	.287	-.971	.566
Security risk not fully covered under existing Reg/RMFs	70	1	5	3.47	1.151	-.075	.287	-1.010	.566
Sustainability risk not fully covered under existing Reg/RMFs	70	1	5	3.29	1.131	.028	.287	-.825	.566
Valid N (listwise)	70								

For each of the nine risk domains, the mean, standard deviation, skewness, and kurtosis were calculated. The survey questions were formulated to elicit responses on a Likert scale, asking whether respondents believed that existing regulations and frameworks do not fully cover the risks in each risk domain. The scale ranged from 1 (strongly disagree) to 5 (strongly agree). A higher number implies that the specific risk is not fully covered, and the respondent agrees with this assessment.

Among the nine risk domains, Privacy risk had the highest mean (3.53). This is not surprising, as privacy risk coverage for generative AI under existing regulations and frameworks remains a concern. Fairness (3.29), Resilience (3.27), and Sustainability (3.29) fell at the lower end of the distribution, suggesting that these three risks are of lesser concern for generative AI under existing regulations and frameworks.

However, the mean for all nine risk domains falls within a small range (minimum 3.27 and maximum 3.53). However, it may not be statistically significant enough to draw a concrete conclusion. Also, the objective was to measure the perceived effectiveness of current AI regulations and the NIST RMF in covering key risks associated with generative AI.

A new dependent variable (overall_assurance) was calculated as the mean of all risk domain scores for each respondent.

Table 4.6: Descriptive Statistics (Research Question 2, Overall Assurance), Source: author's SPSS output, 2025

Descriptive Statistics					
	N	Minimum	Maximum	Mean	Std. Deviation
Overall_assurance	70	1.44	5.00	3.3905	.91248
Valid N (listwise)	70				

The mean for overall assurance was 3.39, indicating concern about the risk assurance for generative AI under existing regulations and frameworks.

The mean for all nine risk domains ranged from 3.27 to 3.53, with an overall assurance score of 3.39. This indicates an overall sentiment slightly above "Neutral," but not quite a clear "Agree." This ambiguous, lukewarm result suggests that while there is *some* risk coverage, it is likely perceived as incomplete or not entirely suited for the unique challenges of generative AI. Given the ambiguous result, it was necessary to go beyond simple descriptive statistics to understand the data's structure and reliability.

Reliability Analysis was used to assess whether all nine risk domains reliably measured the same underlying concept ("risk assurance"). If they are, we can confidently use the composite score. If not, it may indicate that respondents see these risks as distinct categories with varying levels of regulatory coverage.

Using SPSS Reliability Analysis, Cronbach's Alpha was calculated for all nine risk domains.

Table 4.7: Reliability Statistics (Research Question 2), Source: author's SPSS output, 2025

Reliability Statistics	
Cronbach's Alpha	N of Items
.931	9

A value above 0.7 is generally considered acceptable reliability. The overall Cronbach's Alpha score was .931 (refer to Table 4.7).

Cronbach's Alpha scores were calculated for each risk domain after deleting each domain in turn. Refer to Table 4.8. For each risk domain, Cronbach's alpha was greater than 0.9. Hence, it can be reasonably assured that all nine risk domains reliably measure the same underlying concept, risk assurance.

Table 4.8: Item-Total Statistics (Research Question 2), Source: author's SPSS output, 2025

Item-Total Statistics				
	Scale Mean if Item Deleted	Scale Variance if Item Deleted	Corrected Item-Total Correlation	Cronbach's Alpha if Item Deleted
Accountability risk not fully covered under existing Reg/RMFs	27.06	54.692	.690	.927
Fairness risk not fully covered under existing Reg/RMFs	27.23	56.092	.708	.926
Transparency risk not fully covered under existing Reg/RMFs	27.17	53.072	.797	.920
Explainability risk not fully covered under existing Reg/RMFs	27.11	53.349	.743	.924
Resilience risk not fully covered under existing Reg/RMFs	27.24	56.013	.658	.929
Privacy risk not fully covered under existing Reg/RMFs	26.99	50.652	.804	.920
Safety risk not fully covered under existing Reg/RMFs	27.04	53.549	.756	.923
Security risk not fully covered under existing Reg/RMFs	27.04	52.737	.800	.920
Sustainability risk not fully covered under existing Reg/RMFs	27.23	53.541	.762	.922

One-Sample T-Test was used to determine if the mean score of the composite overall assurance variable (or any individual risk) was significantly different from a neutral value of 3. This test determines whether the slight positive inclination (e.g., a Fairness mean of 3.29) is a genuine effect or due to random chance in the sample. Using the SPSS One-Sample T-Test with a Test Value of 3 (a neutral score), Table 4.9 shows the Two-Sided p-value. If it is less than 0.05 ($p < .05$), the mean score is significantly different from neutral. As shown in Table 4.9, the Two-Sided p-value for each risk domain was less than 0.05.

Additionally, the overall assurance, with a mean of 3.58 and a p-value less than 0.001, indicates that the sentiment is genuinely positive, albeit cautiously so.

Table 4.9: One-Sample Test (Research Question 2), Source: author's SPSS output, 2025

One-Sample Test							
Test Value = 3							
	t	df	Significance		Mean Difference	95% Confidence Interval of the Difference	
			One-Sided p	Two-Sided p		Lower	Upper
Accountability risk not fully covered under existing Reg/RMFs	3.399	69	<.001	.001	.457	.19	.73
Fairness risk not fully covered under existing Reg/RMFs	2.439	69	.009	.017	.286	.05	.52
Transparency risk not fully covered under existing Reg/RMFs	2.543	69	.007	.013	.343	.07	.61
Explainability risk not fully covered under existing Reg/RMFs	2.855	69	.003	.006	.400	.12	.68
Resilience risk not fully covered under existing Reg/RMFs	2.166	69	.017	.034	.271	.02	.52
Privacy risk not fully covered under existing Reg/RMFs	3.361	69	<.001	.001	.529	.21	.84
Safety risk not fully covered under existing Reg/RMFs	3.464	69	<.001	<.001	.471	.20	.74
Security risk not fully covered under existing Reg/RMFs	3.426	69	<.001	.001	.471	.20	.75
Sustainability risk not fully covered under existing Reg/RMFs	2.113	69	.019	.038	.286	.02	.56
Overall_assurance	3.580	69	<.001	<.001	.39048	.1729	.6080

Exploratory Factor Analysis (EFA) is a statistical method used to identify the underlying, unobservable constructs that explain the patterns of correlations among a set of observed variables.

We have nine observed variables (e.g., Accountability, Fairness, Transparency). Respondents may not be evaluating each of these nine risks in isolation. Instead, their responses were likely influenced by a few broader, hidden mental models that they used to group some of these risk domains. EFA helps discover these mental models.

A pre-analysis check was done to validate the suitability of the data for EFA.

A sample size of $n=70$ is considered on the lower end for EFA. A common rule of thumb is a minimum of 5-10 respondents per variable. Since we have nine variables, 70 respondents ($70/9 = 7.77$) is acceptable, though not ideal.

The Kaiser-Meyer-Olkin (KMO) measure yielded a value of 0.877 (refer to Table 4.10), which is higher than the ideal score of 0.7. Bartlett’s test of sphericity yielded a p-value of less than 0.001, indicating statistical significance.

Table 4.10: KMO and Bartlett’s Test (Research Question 2), Source: author’s SPSS output, 2025

KMO and Bartlett's Test		
Kaiser-Meyer-Olkin Measure of Sampling Adequacy.		.877
Bartlett's Test of Sphericity	Approx. Chi-Square	473.125
	df	36
	Sig.	<.001

Using SPSS, EFA was conducted. Two risk groups emerged from the EFA analysis.

Group 1 - It consisted of Safety, Security, and Privacy risk. Refer to Table 4.11. Since the risk elements primarily relate to data protection and integrity, the researcher has named this group “Data Protection and Integrity.”

Group 2 - It consisted of Fairness, Explainability, Transparency, Resilience, and Accountability. Refer to Table 4.11. Since the risk elements primarily relate to data ethical governance and trust, the researcher has named this group “Ethical Governance and Trust.”

Sustainability was not a factor in either group (1 or 2).

Table 4.11: Rotated Component Matrix (Research Question 2), Source: author's SPSS output, 2025

	Component	
	1	2
Safety risk not fully covered under existing Reg/RMFs	.923	
Security risk not fully covered under existing Reg/RMFs	.890	.305
Privacy risk not fully covered under existing Reg/RMFs	.822	.380
Sustainability risk not fully covered under existing Reg/RMFs	.598	.561
Fairness risk not fully covered under existing Reg/RMFs		.845
Explainability risk not fully covered under existing Reg/RMFs	.378	.759
Transparency risk not fully covered under existing Reg/RMFs	.499	.703
Resilience risk not fully covered under existing Reg/RMFs	.332	.695
Accountability risk not fully covered under existing Reg/RMFs	.386	.683

Extraction Method: Principal Component Analysis.
 Rotation Method: Varimax without Kaiser Normalization.
 a. Rotation converged in 3 iterations.

The EFA analysis has highlighted that overall risk assurance is not a single concept. The composite mean score was masking some mental models used by the respondents. Based on two groups, “Data Protection and Integrity” and “Ethical Governance and Trust,” two new composite dependent variables were created to calculate risk assurance for each of

them. Since Sustainability was not part of either group, it was calculated separately in the risk assurance calculation. Table 4.12 presents the analysis results.

Table 4.12: Descriptive Statistics (Research Question 2), Source: author's SPSS output, 2025

	Descriptive Statistics								
	N Statistic	Minimum Statistic	Maximum Statistic	Mean Statistic	Std. Deviation Statistic	Skewness		Kurtosis	
						Statistic	Std. Error	Statistic	Std. Error
DataProtectionIntegrity_assurance	70	1.00	5.00	3.4905	1.12033	-.244	.287	-.904	.566
EthicalGovernanceTrust_assurance	70	1.00	5.00	3.3514	.90438	-.193	.287	-.225	.566
Sustainability risk not fully covered under existing Reg/RMFs	70	1	5	3.29	1.131	.028	.287	-.825	.566
Valid N (listwise)	70								

While overall risk assurance is neutral (mean = 3.39), it is driven by moderate scores across three distinct risk categories: Data Protection and Integrity (mean = 3.49), Ethical Governance and Trust (mean = 3.35), and Sustainability (mean = 3.29). Refer to Table 4.12 for the analysis output. This suggests that Data Protection and Integrity risks are considered the least covered under the existing regulations. The Ethical Governance and Trust risks are only second to Data Protection and Integrity risks in terms of perceived risk coverage. Sustainability stood alone and had a higher mean (3.29) than a neutral score of 3. The IT service professionals (survey respondents) seemed least concerned about sustainability risks.

4.3 Research Question Three

RQ3 How does the lack of generative AI-specific regulations affect the adoption of productivity and innovation use cases by IT service companies?

The core objective was to determine if the perceived lack of regulatory coverage for specific generative AI risks acts as a catalyst for the adoption of these solutions in

productivity (business-critical and business-support) and creativity (innovation) applications.

Each risk (Accountability, Fairness, Transparency, Explainability, Resilience, Privacy, Safety, Security, and Sustainability) has a pair of competing hypotheses. It's important to note that alternative hypotheses are directional (they posit a "positive effect").

Null Hypothesis (H_0) for each risk: The lack of well-defined, regulated risk [X] does not affect adoption. The mean score equal to or less than the neutral point (3) implies no effect or an adverse effect on adoption.

Alternative Hypothesis (H_1) for each risk: The lack of regulated risk [X] positively affects (increases) adoption (i.e., it's seen as an opportunity). The mean score significantly greater than the neutral point (3) implies a positive effect.

The survey question was designed to gather responses for business-critical, business-support, and creativity (innovation) applications separately.

Hence, the analysis was done by application area (Business-critical, Business-support, and Creativity).

Business-critical application of generative AI

Using SPSS and Descriptive Statistics, Mean, Median, Standard Deviation, Skewness, and Kurtosis were calculated for each risk area.

Table 4.13: Descriptive Statistics (Research Question 3), Source: author's SPSS output, 2025

		Statistics								
		CRIT_Acc_Opp	CRIT_Fai_Opp	CRIT_Trans_Opp	CRIT_Explain_Opp	CRIT_Resi_Opp	CRIT_Pri_Opp	CRIT_Safety_Opp	CRIT_Sec_Opp	CRIT_Sus_Opp
N	Valid	70	70	70	70	70	70	70	70	70
	Missing	0	0	0	0	0	0	0	0	0
Mean		3.51	3.31	3.47	3.50	3.33	3.56	3.53	3.53	3.36
Median		4.00	3.00	4.00	4.00	3.00	4.00	4.00	4.00	3.00
Std. Deviation		1.046	1.001	1.086	1.126	1.018	1.293	1.282	1.213	1.143
Skewness		-1.017	-.764	-.762	-.658	-.707	-.482	-.532	-.571	-.267
Std. Error of Skewness		.287	.287	.287	.287	.287	.287	.287	.287	.287
Kurtosis		.798	.506	.275	.044	.426	-.829	-.816	-.460	-.509
Std. Error of Kurtosis		.566	.566	.566	.566	.566	.566	.566	.566	.566

In all risk areas, the mean score exceeds 3, invalidating the null hypothesis. There is a general perception amongst the IT practitioners that the lack of strict regulations across risk areas is actually positively affecting adoption.

Out of nine risk areas, Privacy (mean score 3.56), Safety (mean score 3.53), Security (mean score 3.53), and Accountability (mean score 3.51) are the top areas experiencing greater adoption due to lax regulations. However, Privacy (skewness score 1.293, SD) and Safety (skewness score 1.282, SD) also exhibit high skewness.

Table 4.14: Privacy Score Distribution (Research Question 3), Source: author's SPSS output, 2025

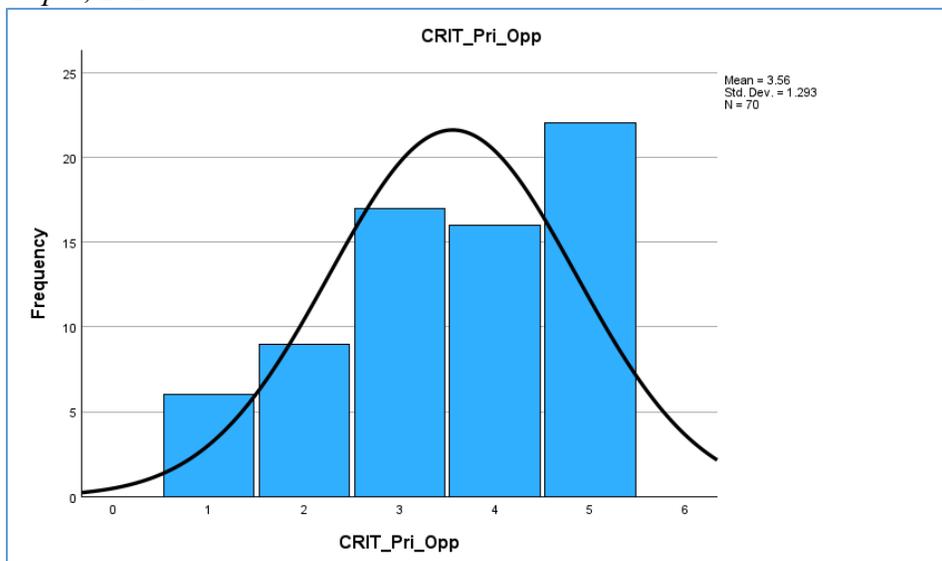


Table 4.15: Safety Score Distribution (Research Question 3), Source: author's SPSS output, 2025

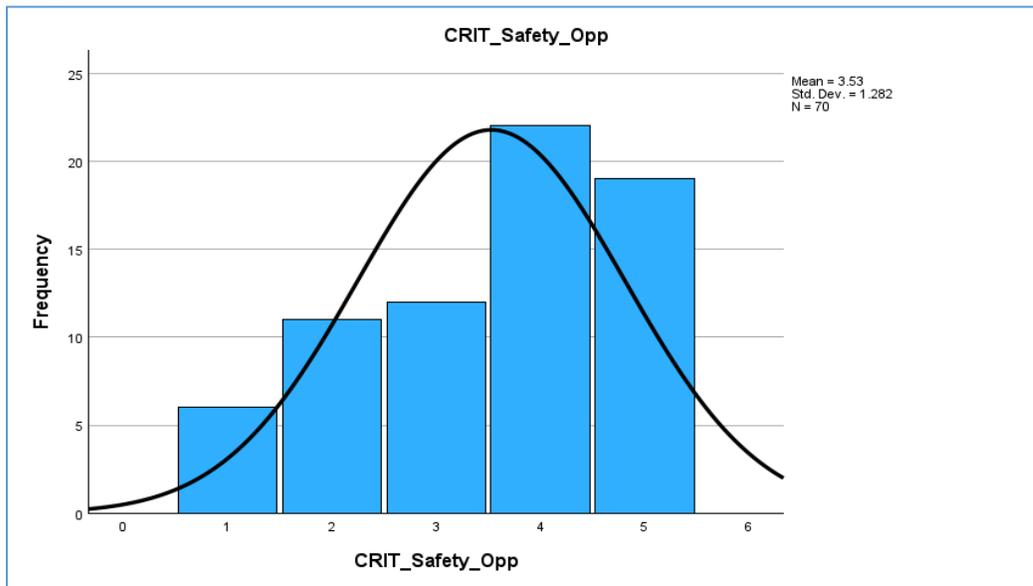


Table 4.16: Security Score Distribution (Research Question 3), Source: author's SPSS output, 2025

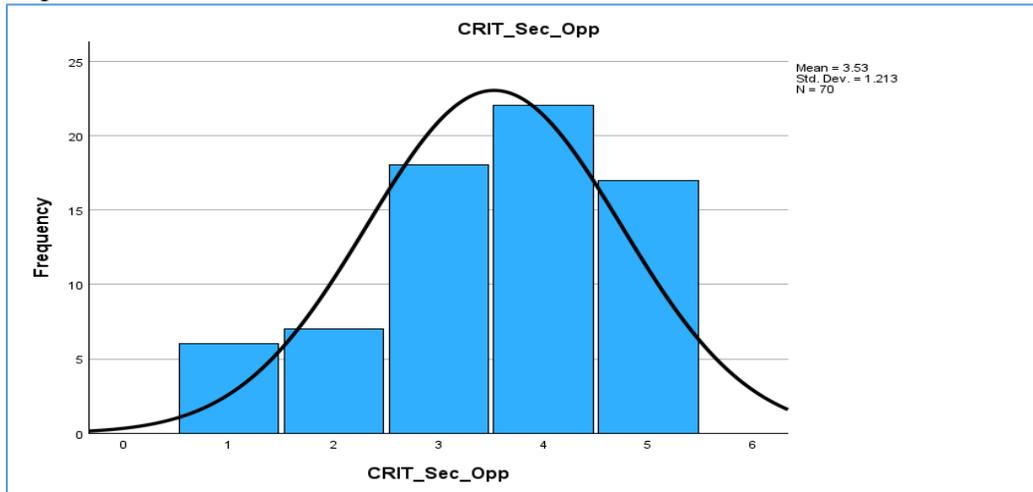
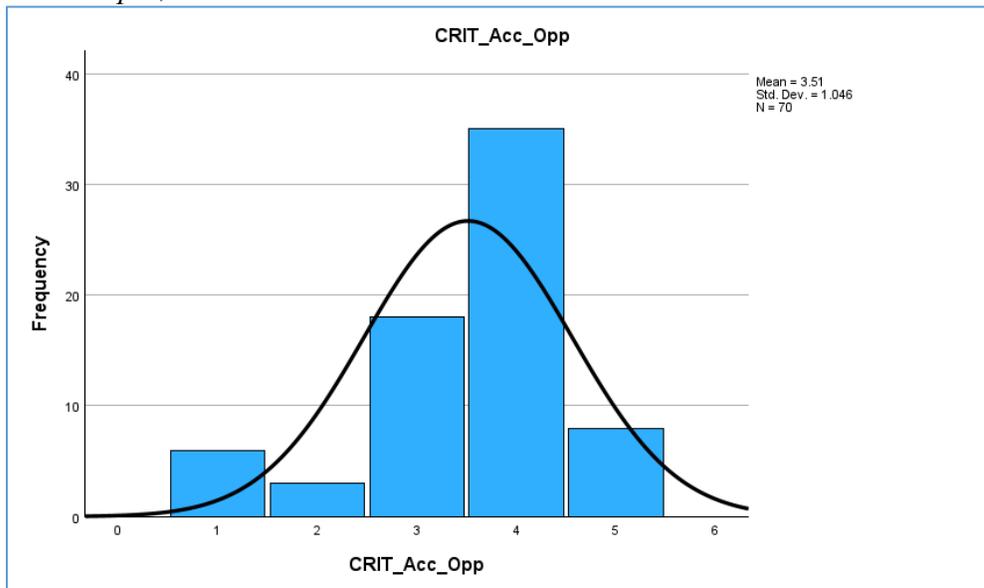


Table 4.17: Accountability Score Distribution (Research Question 3), Source: author's SPSS output, 2025



Inferential Analysis (Hypothesis Testing) was used to statistically test if the mean score for each risk is significantly greater than the neutral value of 3.

A one-sample T-test compares the mean of a single sample (e.g., the response for CRIT_Privacy_Opp) to a test value 3 (the neutral point).

Using SPSS and a One-Sample T-Test, a two-tailed p-value was calculated. Since the hypothesis is one-tailed (directional), the reported 2-tailed value is divided by 2 to get the corrected one-tailed p-value.

Table 4.18: One-Sample T-Test (Research Question 3), Source: author's SPSS output, 2025

One-Sample Test							
Test Value = 3							
	t	df	Significance		Mean Difference	95% Confidence Interval of the Difference	
			One-Sided p	Two-Sided p		Lower	Upper
CRIT_Acc_Opp	4.114	69	<.001	<.001	.514	.26	.76
CRIT_Fai_Opp	2.628	69	.005	.011	.314	.08	.55
CRIT_Trans_Opp	3.630	69	<.001	<.001	.471	.21	.73
CRIT_Explain_Opp	3.715	69	<.001	<.001	.500	.23	.77
CRIT_Resi_Opp	2.702	69	.004	.009	.329	.09	.57
CRIT_Pri_Opp	3.606	69	<.001	<.001	.557	.25	.87
CRIT_Safety_Opp	3.449	69	<.001	<.001	.529	.22	.83
CRIT_Sec_Opp	3.647	69	<.001	<.001	.529	.24	.82
CRIT_Sus_Opp	2.615	69	.005	.011	.357	.08	.63

Refer to Table 4.18, 1-tailed p-value (Two-sided p/2) is less than the alpha level (.05), so the null hypothesis is rejected. For a specific risk (e.g., Privacy), if the one-tailed p-value is < 0.05 and the Mean Difference is positive, then we can reject H₀. There is significant evidence that the lack of regulation regarding privacy risks is perceived as providing more opportunities for adoption.

Regression analysis for Business-critical applications: While the one-sample t-test tells us if the perception for each risk is significantly above neutral, regression can answer more complex questions. Can we predict an overall "Adoption Opportunity" score based on the perceived gaps in all nine risks simultaneously? Which specific risk gaps (e.g., lack of Accountability vs. lack of Privacy) are the strongest unique drivers of the adoption perception?

Using SPSS, a composite dependent variable, CRIT_Overall_Adoption_Opp, was created to build the linear regression model with nine risk gaps.

Table 4.19: Model Summary (Research Question 3), Source: author's SPSS output, 2025

Model Summary ^b				
Model	R	R Square	Adjusted R Square	Std. Error of the Estimate
1	.541 ^a	.293	.186	.89450

a. Predictors: (Constant), Sus_Gap, Acc_Gap, Expl_Gap, Safe_Gap, Res_Gap, Fai_Gap, Trans_Gap, Pri_Gap, Sec_Gap
 b. Dependent Variable: CRIT_Overall_Adoption_Opp

R² indicates the proportion of variance in the overall adoption opportunity score that is explained by all nine perceived regulatory gaps combined. Refer to Table 4.19. An R² of 0.293 suggests that the model accounts for 29.3% of the variation in adoption opportunity perceptions.

Table 4.20: ANOVA (Research Question 3), Source: author's SPSS output, 2025

ANOVA ^a						
Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	19.848	9	2.205	2.756	.009 ^b
	Residual	48.008	60	.800		
	Total	67.856	69			

a. Dependent Variable: CRIT_Overall_Adoption_Opp
 b. Predictors: (Constant), Sus_Gap, Acc_Gap, Expl_Gap, Safe_Gap, Res_Gap, Fai_Gap, Trans_Gap, Pri_Gap, Sec_Gap

Sig. (p-value): This tests whether the overall regression model is statistically significant. Our Sig. (p-value) is .009, which is less than .05 (Table 4.20). The combination of all nine risk gaps significantly predicts the dependent variable, CRIT_Overall_Adoption_Opp. It indicates that our model is superior to simply using the mean.

Table 4.21: Coefficients (Research Question 3), Source: author's SPSS output, 2025

Coefficients ^a											
Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.	95.0% Confidence Interval for B		Correlations		
		B	Std. Error	Beta			Lower Bound	Upper Bound	Zero-order	Partial	Part
1	(Constant)	1.842	.451		4.083	<.001	.940	2.745			
	Acc_Gap	-.074	.142	-.084	-.522	.603	-.359	.211	.310	-.067	-.057
	Fai_Gap	.114	.195	.113	.586	.560	-.276	.505	.374	.075	.064
	Trans_Gap	.355	.180	.403	1.966	.054	-.006	.715	.496	.246	.213
	Expl_Gap	-.063	.147	-.075	-.431	.668	-.357	.230	.326	-.056	-.047
	Res_Gap	-.011	.155	-.011	-.068	.946	-.320	.299	.229	-.009	-.007
	Pri_Gap	.103	.158	.136	.650	.518	-.213	.419	.419	.084	.071
	Safe_Gap	.186	.208	.214	.896	.374	-.229	.602	.433	.115	.097
	Sec_Gap	.001	.212	.001	.006	.996	-.422	.424	.416	.001	.001
	Sus_Gap	-.138	.170	-.158	-.814	.419	-.478	.201	.299	-.104	-.088

a. Dependent Variable: CRIT_Overall_Adoption_Opp

Based on the above Coefficient table (Table 4.21), the Transparency regulation gap has the highest Standardized Coefficient Beta, 0.403. T & Sig. (p-value) tests if the coefficient for each independent variable (regulation gaps) is significantly different from zero. A p-value < .05 means this specific regulatory gap is a unique, significant predictor of adoption opportunity. Safety is an essential but weaker driver (B = 0.214, p = 0.374).

Business-support application of generative AI

Using SPSS, Descriptive Statistics (Mean, Median, Standard Deviation, Skewness, and Kurtosis) were calculated for each risk area.

Table 4.22: Descriptive Statistics (Research Question 3), Source: author's SPSS output, 2025

		Statistics								
		Support_Acc_Opp	Support_Fai_Opp	Support_Trans_Opp	Support_Explai_n_Opp	Support_Resi_Opp	Support_Pri_Opp	Support_Safety_Opp	Support_Sec_Opp	Support_Sus_Opp
N	Valid	70	70	68	70	70	70	70	70	70
	Missing	0	0	2	0	0	0	0	0	0
Mean		3.43	3.40	3.41	3.41	3.37	3.47	3.50	3.49	3.33
Median		4.00	3.00	3.00	3.50	3.00	3.00	4.00	3.50	3.00
Std. Deviation		1.057	.999	1.011	1.014	1.010	1.126	1.164	1.201	1.113
Skewness		-.715	-.525	-.466	-.403	-.376	-.335	-.511	-.301	-.300
Std. Error of Skewness		.287	.287	.291	.287	.287	.287	.287	.287	.287
Kurtosis		.404	.240	.237	-.169	.136	-.464	-.428	-.829	-.400
Std. Error of Kurtosis		.566	.566	.574	.566	.566	.566	.566	.566	.566

In all risk areas, the mean score exceeds 3, invalidating the null hypothesis. There is a general perception that the lack of strict regulations across risk areas is actually positively affecting adoption.

Among the nine risk areas, Safety (mean score 3.50), Security (mean score 3.49), Privacy (mean score 3.47), and Accountability (mean score 3.51) are the top drivers for adoption due to lax regulations.

Table 4.23: Support Safety Score (Research Question 3), Source: author's SPSS output, 2025

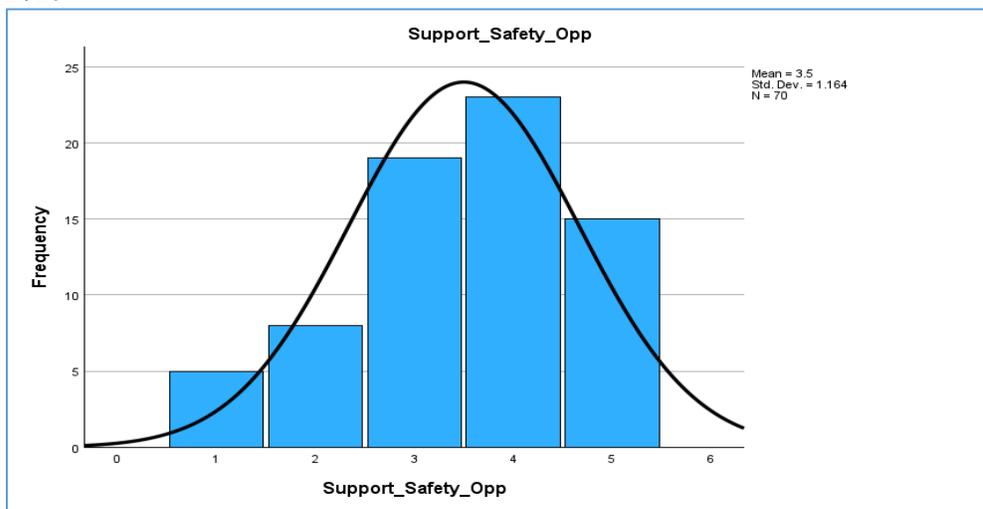


Table 4.24: Support Security Score (Research Question 3), Source: author's SPSS output, 2025

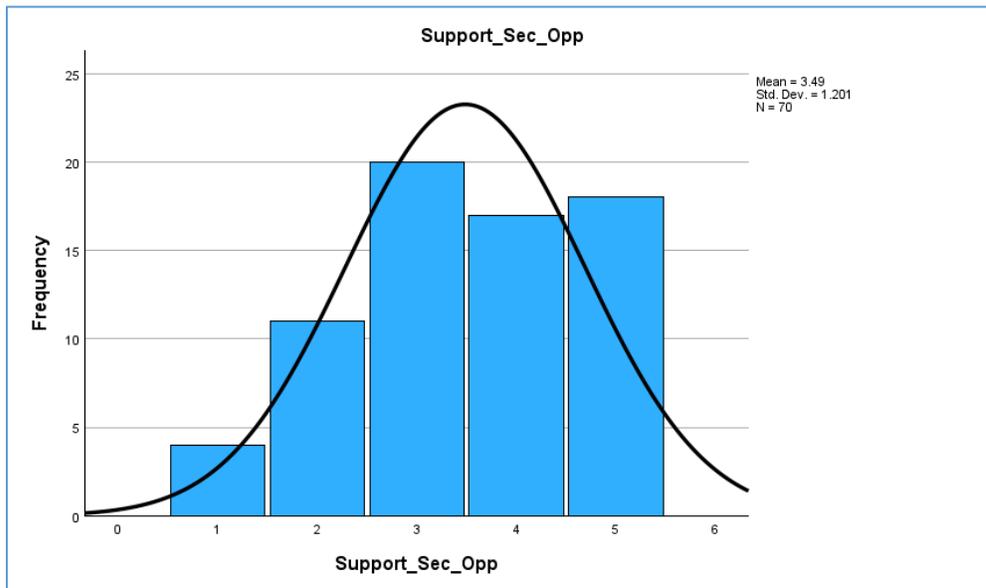


Table 4.25: Support Privacy Score (Research Question 3), Source: author's SPSS output, 2025

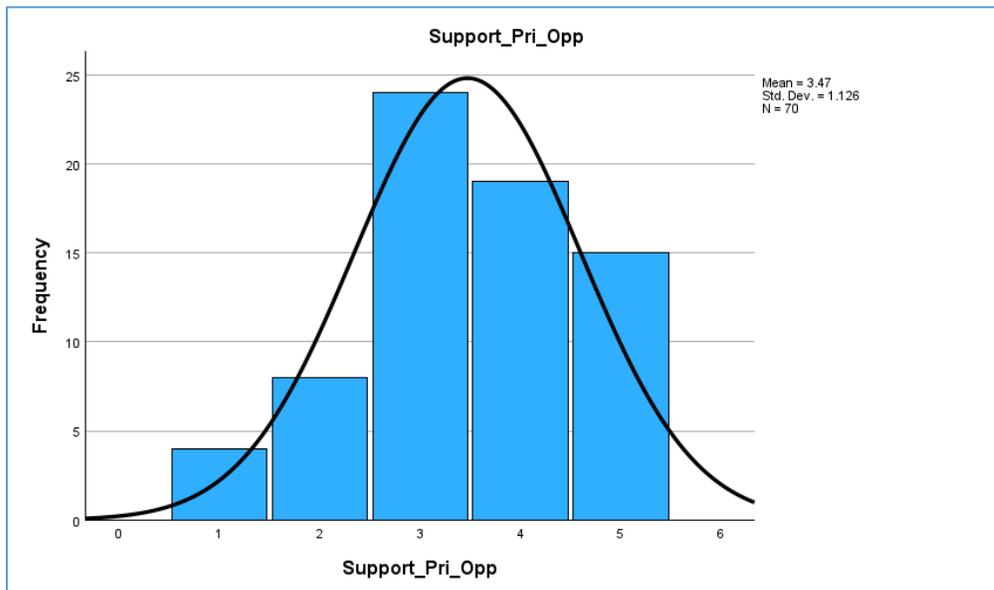
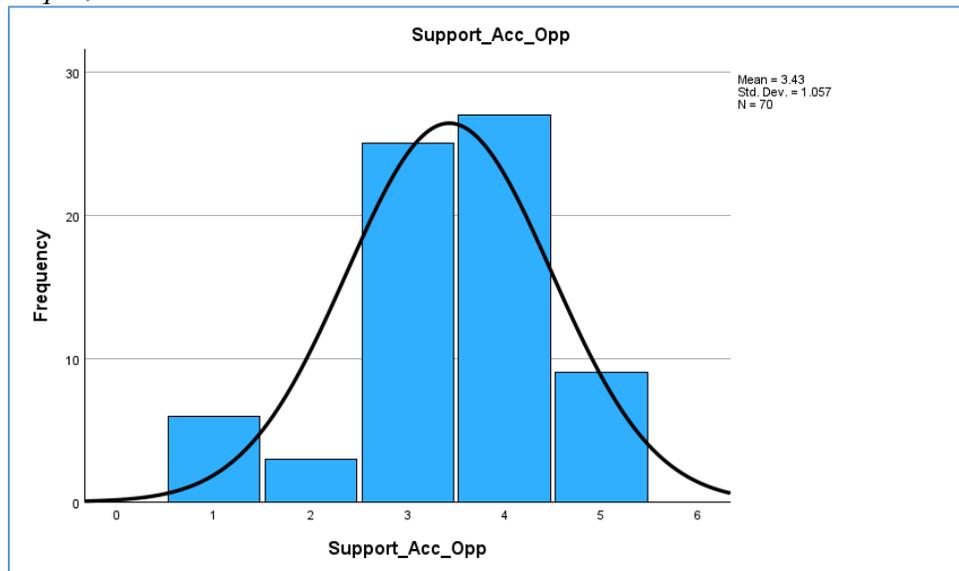


Table 4.26: Support Accountability Score (Research Question 3), Source: author's SPSS output, 2025



Inferential Analysis (Hypothesis Testing) was conducted to statistically test if the mean score for each risk is significantly greater than the neutral value of 3. A one-sample T-test compares the mean of a single sample (e.g., the response for Support_Safety_Opp) to a test value 3 (neutral point).

Using SPSS and a One-Sample T-Test, a two-tailed p-value was calculated. Since the hypothesis is one-tailed (directional), the reported 2-tailed value is divided by 2 to get the corrected one-tailed p-value.

Table 4.27: One-Sample Test (Research Question 3), Source: author's SPSS output, 2025

One-Sample Test							
Test Value = 3							
	t	df	Significance		Mean Difference	95% Confidence Interval of the Difference	
			One-Sided p	Two-Sided p		Lower	Upper
Support_Acc_Opp	3.391	69	<.001	.001	.429	.18	.68
Support_Fai_Opp	3.352	69	<.001	.001	.400	.16	.64
Support_Trans_Opp	3.359	67	<.001	.001	.412	.17	.66
Support_Explain_Opp	3.417	69	<.001	.001	.414	.17	.66
Support_Resi_Opp	3.078	69	.001	.003	.371	.13	.61
Support_Pri_Opp	3.504	69	<.001	<.001	.471	.20	.74
Support_Safety_Opp	3.594	69	<.001	<.001	.500	.22	.78
Support_Sec_Opp	3.384	69	<.001	.001	.486	.20	.77
Support_Sus_Opp	2.470	69	.008	.016	.329	.06	.59

Refer to Table 4.27, 1-tailed p-value (Two-sided p/2) is less than the alpha level (.05), so the null hypothesis is rejected. For a specific risk (e.g., Privacy), if the one-tailed p-value is < 0.05 and the Mean Difference is positive, then we can reject H₀. There is significant evidence that the lack of regulation regarding privacy risks is perceived as providing more opportunities for adoption.

Regression analysis for Business-support applications: While the one-sample t-test tells us if the perception for each risk is significantly above neutral, regression can answer more complex questions. Can we predict an overall "Adoption Opportunity" score based on the perceived gaps in all nine risks simultaneously? Which specific risk gaps are the strongest unique drivers of the adoption perception?

Using SPSS, a composite dependent variable, Support_Overall_Adoption_Opp, was created to build the linear regression model with nine risk gaps.

Table 4.28: Support Model Summary (Research Question 3), Source: author's SPSS output, 2025

Model Summary ^b				
Model	R	R Square	Adjusted R Square	Std. Error of the Estimate
1	.511 ^a	.262	.151	.82137

a. Predictors: (Constant), Sus_Gap, Acc_Gap, Expl_Gap, Safe_Gap, Res_Gap, Fai_Gap, Trans_Gap, Pri_Gap, Sec_Gap
 b. Dependent Variable: Support_Overall_Adoption_Opp

R² indicates the proportion of variance in the overall adoption opportunity score that is explained by all nine perceived regulatory gaps combined. Refer to Table 4.28. An R² of 0.262 suggests that the model accounts for 26.2% of the variation in perceptions of adoption opportunities.

Table 4.29: Support ANOVA (Research Question 3), Source: author's SPSS output, 2025

ANOVA ^a						
Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	14.340	9	1.593	2.362	.023 ^b
	Residual	40.479	60	.675		
	Total	54.819	69			

a. Dependent Variable: Support_Overall_Adoption_Opp
 b. Predictors: (Constant), Sus_Gap, Acc_Gap, Expl_Gap, Safe_Gap, Res_Gap, Fai_Gap, Trans_Gap, Pri_Gap, Sec_Gap

Sig. (p-value): This tests whether the overall regression model is statistically significant. Our Sig. (p-value) is .023, which is less than .05, we can conclude that the combination of all nine risk gaps significantly predicts the dependent variable, Support_Overall_Adoption_Opp. It indicates that our model is superior to simply using the mean.

Table 4.30: Support Coefficient (Research Question 3), Source: author's SPSS output, 2025

Coefficients ^a											
Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.	95.0% Confidence Interval for B		Correlations		
		B	Std. Error	Beta			Lower Bound	Upper Bound	Zero-order	Partial	Part
1	(Constant)	2.074	.414		5.005	<.001	1.245	2.902			
	Acc_Gap	-.013	.131	-.016	-.096	.924	-.274	.249	.309	-.012	-.011
	Fai_Gap	.061	.179	.067	.340	.735	-.298	.420	.352	.044	.038
	Trans_Gap	.387	.166	.489	2.334	.023	.055	.718	.476	.289	.259
	Expl_Gap	-.149	.135	-.196	-1.107	.273	-.418	.120	.256	-.141	-.123
	Res_Gap	.035	.142	.041	.248	.805	-.249	.319	.235	.032	.028
	Pri_Gap	.048	.145	.071	.334	.740	-.242	.339	.357	.043	.037
	Safe_Gap	.164	.191	.209	.859	.394	-.218	.546	.378	.110	.095
	Sec_Gap	-.066	.194	-.085	-.340	.735	-.455	.323	.358	-.044	-.038
	Sus_Gap	-.066	.156	-.084	-.423	.674	-.378	.246	.288	-.055	-.047

a. Dependent Variable: Support_Overall_Adoption_Opp

Based on the above Coefficient table (Table 4.30), the Transparency regulation gap has the highest Standardized Coefficient Beta, 0.489. T & Sig. (p-value) tests if the coefficient for each independent variable (regulation gaps) is significantly different from zero. A p-value < .05 means this specific regulatory gap is a unique, significant predictor of adoption opportunity. Explainability is an essential but weaker driver (B = 0.196, p = 0.273).

Creativity (Innovation) application of generative AI: Using SPSS, Descriptive Statistics (Mean, Median, Standard Deviation, Skewness, and Kurtosis) were calculated for each risk area.

Table 4.31: Creativity Statistics (Research Question 3), Source: author's SPSS output, 2025

Statistics										
		Creativity_Acc_Opp	Creativity_Fai_Opp	Creativity_Trans_Opp	Creativity_Explain_Opp	Creativity_Resi_Opp	Creativity_Pri_Opp	Creativity_Safety_Opp	Creativity_Sec_Opp	Creativity_Sustain_Opp
N	Valid	70	70	70	70	70	70	70	70	70
	Missing	0	0	0	0	0	0	0	0	0
Mean		3.53	3.34	3.46	3.50	3.41	3.40	3.40	3.57	3.29
Median		4.00	3.00	4.00	4.00	4.00	3.00	3.50	4.00	3.00
Std. Deviation		1.139	1.048	1.086	1.139	1.070	1.256	1.267	1.258	1.118
Skewness		-.740	-.425	-.726	-.424	-.465	-.309	-.318	-.571	-.337
Std. Error of Skewness		.287	.287	.287	.287	.287	.287	.287	.287	.287
Kurtosis		.029	-.270	.239	-.500	-.254	-.875	-.932	-.628	-.331
Std. Error of Kurtosis		.566	.566	.566	.566	.566	.566	.566	.566	.566

In all risk areas, the mean score exceeds 3, invalidating the null hypothesis. There is a general perception that the lack of strict regulations across risk areas is actually positively affecting adoption.

Security (mean score 3.57), Accountability (mean score 3.53), and Explainability (mean score 3.50) are the top areas experiencing greater adoption due to lax regulations.

Table 4.32: Creativity Security Statistics (Research Question 3), Source: author's SPSS output, 2025

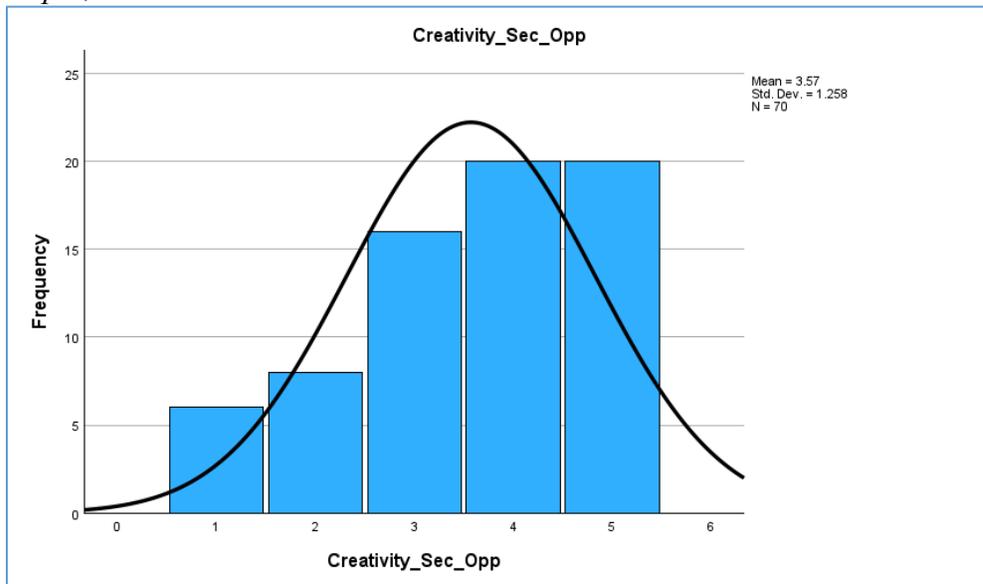


Table 4.33: Creativity Accountability Statistics (Research Question 3), Source: author's SPSS output, 2025

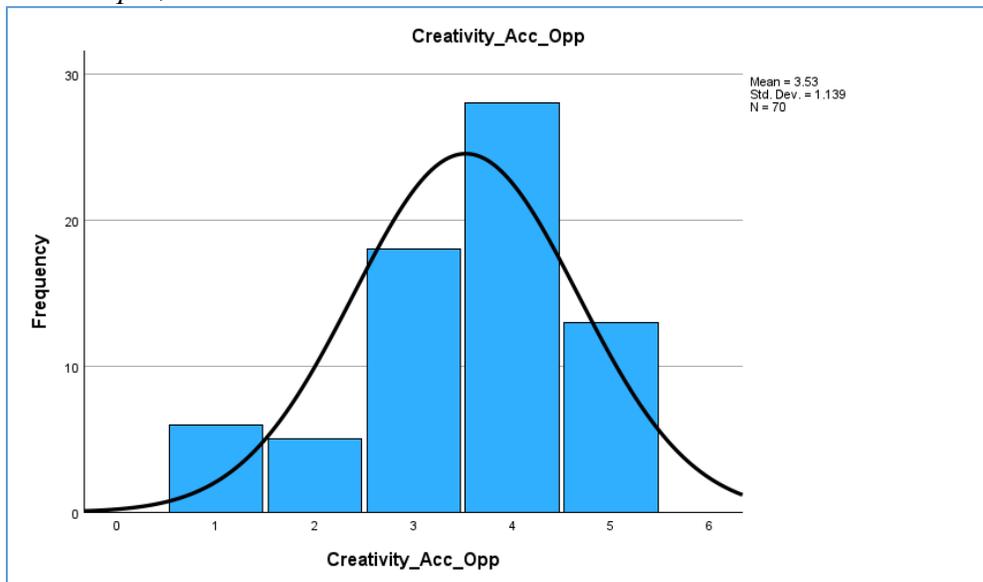
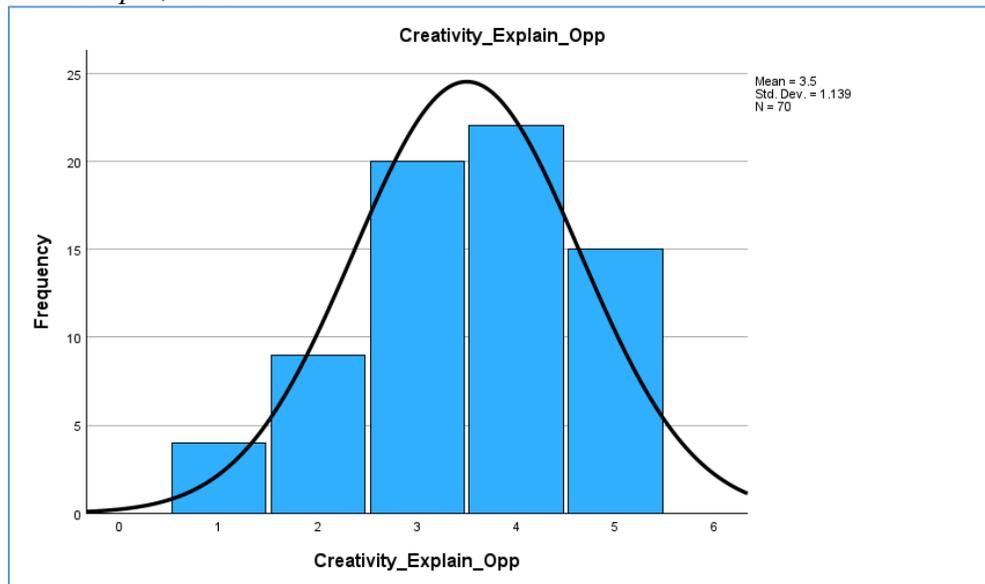


Table 4.34: Creativity Explainability Statistics (Research Question 3), Source: author's SPSS output, 2025



Inferential Analysis (Hypothesis Testing) was used to statistically test if the mean score for each risk is significantly greater than the neutral value of 3. A one-sample T-test compares the mean of a single sample (e.g., the response for Support_Safety_Opp) to a test value (3 here, representing the neutral point).

Using SPSS and a One-Sample T-Test, a two-tailed p-value was calculated. Since the hypothesis is one-tailed (directional), the reported 2-tailed value is divided by 2 to get the corrected one-tailed p-value.

Table 4.35: Creativity One-Sample Test (Research Question 3), Source: author's SPSS output, 2025

One-Sample Test							
Test Value = 3							
	t	df	Significance		Mean Difference	95% Confidence Interval of the Difference	
			One-Sided p	Two-Sided p		Lower	Upper
Creativity_Acc_Opp	3.884	69	<.001	<.001	.529	.26	.80
Creativity_Fai_Opp	2.737	69	.004	.008	.343	.09	.59
Creativity_Trans_Opp	3.522	69	<.001	<.001	.457	.20	.72
Creativity_Explain_Opp	3.673	69	<.001	<.001	.500	.23	.77
Creativity_Resi_Opp	3.240	69	<.001	.002	.414	.16	.67
Creativity_Pri_Opp	2.665	69	.005	.010	.400	.10	.70
Creativity_Safety_Opp	2.641	69	.005	.010	.400	.10	.70
Creativity_Sec_Opp	3.801	69	<.001	<.001	.571	.27	.87
Creativity_Sustain_Opp	2.138	69	.018	.036	.286	.02	.55

Refer to Table 4.35, 1-tailed p-value (Two-sided p/2) is less than the alpha level (.05), so the null hypothesis is rejected. For a specific risk (e.g., Privacy), if the one-tailed p-value is < 0.05 and the Mean Difference is positive, then we can reject H₀. There is significant evidence that the lack of regulation regarding privacy risks is perceived as providing more opportunities for adoption.

Regression analysis for Creativity applications: While the one-sample t-test tells us if the perception for each risk is significantly above neutral, regression can answer more complex questions. Can we predict an overall "Adoption Opportunity" score based on the perceived gaps in all nine risks simultaneously? Which specific risk gaps are the strongest drivers of the perception that adoption is enabled?

Using SPSS, a composite dependent variable, Creativity_Overall_Adoption_Opp, was created to build the linear regression model with nine risk gaps.

Table 4.36: Creativity Model Summary (Research Question 3), Source: author's SPSS output, 2025

Model Summary ^b				
Model	R	R Square	Adjusted R Square	Std. Error of the Estimate
1	.564 ^a	.318	.216	.87070

a. Predictors: (Constant), Sus_Gap, Acc_Gap, Expl_Gap, Safe_Gap, Res_Gap, Fai_Gap, Trans_Gap, Pri_Gap, Sec_Gap
 b. Dependent Variable: Creativity_Overall_Adoption_Opp

R^2 indicates the proportion of variance in the overall adoption opportunity score that is explained by all nine perceived regulatory gaps combined. An R^2 of 0.216 suggests that the model accounts for 21.6% of the variation in adoption opportunity perceptions.

An R^2 of 21.6% is highly unlikely to occur by chance if the nine risk gaps truly had no relationship with the adoption opportunity score. It confirms that the perceived regulatory gaps are meaningfully related to the overall adoption opportunity score.

The primary limitation is that 78.4% (100 - 21.6) of the variance in the adoption opportunity score remains unexplained by the model. While the model identifies a clear signal, that signal is relatively weak compared to the overall "noise" in the data. This model may not be used to predict a new adoption opportunity score because it misses most of the influencing factors.

In practical terms, the model is better suited for explanation and insight than for precise forecasting. It indicates that regulatory gaps matter, but it cannot reliably predict the exact outcome.

Table 4.37: Creativity ANOVA (Research Question 3), Source: author's SPSS output, 2025

ANOVA ^a						
Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	21.232	9	2.359	3.112	.004 ^b
	Residual	45.487	60	.758		
	Total	66.720	69			

a. Dependent Variable: Creativity_Overall_Adoption_Opp
b. Predictors: (Constant), Sus_Gap, Acc_Gap, Expl_Gap, Safe_Gap, Res_Gap, Fai_Gap, Trans_Gap, Pri_Gap, Sec_Gap

Sig. (p-value): This tests whether the overall regression model is statistically significant. Our Sig. (p-value) is .004, which is less than .05, we can conclude that the combination of all nine risk gaps significantly predicts the dependent variable, Creativity_Overall_Adoption_Opp. It indicates that our model is superior to simply using the mean.

Table 4.38: Creativity Coefficients (Research Question 3), Source: author's SPSS output, 2025

Coefficients ^a											
Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.	95.0% Confidence Interval for B		Correlations		
		B	Std. Error	Beta			Lower Bound	Upper Bound	Zero-order	Partial	Part
1	(Constant)	1.758	.439		4.002	<.001	.879	2.636			
	Acc_Gap	-.101	.139	-.115	-.726	.471	-.378	.177	.275	-.093	-.077
	Fai_Gap	.076	.190	.075	.398	.692	-.305	.456	.391	.051	.042
	Trans_Gap	.463	.176	.531	2.634	.011	.111	.814	.502	.322	.281
	Expl_Gap	-.191	.143	-.228	-1.339	.186	-.477	.094	.257	-.170	-.143
	Res_Gap	.103	.151	.110	.686	.495	-.198	.404	.284	.088	.073
	Pri_Gap	-.054	.154	-.073	-.352	.726	-.362	.254	.349	-.045	-.038
	Safe_Gap	.313	.202	.363	1.550	.126	-.091	.718	.420	.196	.165
	Sec_Gap	-.138	.206	-.161	-.670	.506	-.550	.274	.383	-.086	-.071
	Sus_Gap	.037	.165	.043	.226	.822	-.293	.368	.378	.029	.024

a. Dependent Variable: Creativity_Overall_Adoption_Opp

Based on Table 4.38, the Transparency regulation gap has the highest Standardized Coefficient Beta, 0.531. t & Sig. (p-value) tests if the coefficient for each independent variable (regulation gaps) is significantly different from zero. A p-value < .05 means this specific regulatory gap is a unique, significant predictor of adoption opportunity. Safety has the second-highest Standardized Coefficient Beta, 0.363.

The analysis reveals that the perceived regulatory gaps are not equally important. Efforts to understand adoption drivers should focus primarily on the lack of Transparency and Safety regulations, as these are the areas where the absence of rules is most strongly correlated with a perception of opportunity for creativity application adoption.

Multivariate Analysis of Variance (MANOVA): This is a classic scenario for Multivariate Analysis of Variance (MANOVA), as we are dealing with three distinct dependent variables (DVs), each measured across the same set of nine independent variables (IVs) for the same respondents. The goal was to understand how perceptions of regulatory gaps (IVs), both collectively and individually, influence the propensity to adopt across three application contexts (DVs): Business-critical, Business-support, and Creativity. MANOVA assesses the effect of multiple IVs on multiple related DVs

simultaneously. It controls inflated Type 1 error that may occur if three separate ANOVAs are run, as was done in this case. It can also test if IVs affect the DVs differently.

Using SPSS, Multivariate analysis was conducted.

Table 4.39: Multivariate Test (Research Question 3), Source: author's SPSS output, 2025

Multivariate Tests ^a							
Effect		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared
Intercept	Pillai's Trace	.295	8.075 ^b	3.000	58.000	<.001	.295
	Wilks' Lambda	.705	8.075 ^b	3.000	58.000	<.001	.295
	Hotelling's Trace	.418	8.075 ^b	3.000	58.000	<.001	.295
	Roy's Largest Root	.418	8.075 ^b	3.000	58.000	<.001	.295
Acc_Gap	Pillai's Trace	.029	.569 ^b	3.000	58.000	.638	.029
	Wilks' Lambda	.971	.569 ^b	3.000	58.000	.638	.029
	Hotelling's Trace	.029	.569 ^b	3.000	58.000	.638	.029
	Roy's Largest Root	.029	.569 ^b	3.000	58.000	.638	.029
Fai_Gap	Pillai's Trace	.007	.144 ^b	3.000	58.000	.933	.007
	Wilks' Lambda	.993	.144 ^b	3.000	58.000	.933	.007
	Hotelling's Trace	.007	.144 ^b	3.000	58.000	.933	.007
	Roy's Largest Root	.007	.144 ^b	3.000	58.000	.933	.007
Trans_Gap	Pillai's Trace	.106	2.281 ^b	3.000	58.000	.089	.106
	Wilks' Lambda	.894	2.281 ^b	3.000	58.000	.089	.106
	Hotelling's Trace	.118	2.281 ^b	3.000	58.000	.089	.106
	Roy's Largest Root	.118	2.281 ^b	3.000	58.000	.089	.106
Expl_Gap	Pillai's Trace	.041	.830 ^b	3.000	58.000	.483	.041
	Wilks' Lambda	.959	.830 ^b	3.000	58.000	.483	.041
	Hotelling's Trace	.043	.830 ^b	3.000	58.000	.483	.041
	Roy's Largest Root	.043	.830 ^b	3.000	58.000	.483	.041
Res_Gap	Pillai's Trace	.016	.317 ^b	3.000	58.000	.813	.016
	Wilks' Lambda	.984	.317 ^b	3.000	58.000	.813	.016
	Hotelling's Trace	.016	.317 ^b	3.000	58.000	.813	.016
	Roy's Largest Root	.016	.317 ^b	3.000	58.000	.813	.016
Pri_Gap	Pillai's Trace	.027	.537 ^b	3.000	58.000	.659	.027
	Wilks' Lambda	.973	.537 ^b	3.000	58.000	.659	.027
	Hotelling's Trace	.028	.537 ^b	3.000	58.000	.659	.027
	Roy's Largest Root	.028	.537 ^b	3.000	58.000	.659	.027
Safe_Gap	Pillai's Trace	.046	.933 ^b	3.000	58.000	.431	.046
	Wilks' Lambda	.954	.933 ^b	3.000	58.000	.431	.046
	Hotelling's Trace	.048	.933 ^b	3.000	58.000	.431	.046
	Roy's Largest Root	.048	.933 ^b	3.000	58.000	.431	.046
Sec_Gap	Pillai's Trace	.014	.269 ^b	3.000	58.000	.847	.014
	Wilks' Lambda	.986	.269 ^b	3.000	58.000	.847	.014
	Hotelling's Trace	.014	.269 ^b	3.000	58.000	.847	.014
	Roy's Largest Root	.014	.269 ^b	3.000	58.000	.847	.014
Sus_Gap	Pillai's Trace	.029	.579 ^b	3.000	58.000	.631	.029
	Wilks' Lambda	.971	.579 ^b	3.000	58.000	.631	.029
	Hotelling's Trace	.030	.579 ^b	3.000	58.000	.631	.029
	Roy's Largest Root	.030	.579 ^b	3.000	58.000	.631	.029

a. Design: Intercept + Acc_Gap + Fai_Gap + Trans_Gap + Expl_Gap + Res_Gap + Pri_Gap + Safe_Gap + Sec_Gap + Sus_Gap
b. Exact statistic

Multivariate Tests Table (Pillai's Trace): This is the primary test, a significant Sig. Value ($p < .05$) for any of the covariates (e.g., Transparency_Gap) indicates that this particular regulatory gap has a significant overall effect on the combination of the three adoption contexts. Of the nine risk areas, Transparency Gap has the lowest Sig. Value ($p = 0.089$) that indicates a lack of strict transparency regulations has a better impact (positive) on the adoption of generative AI applications across Business-critical, Business-support, and Creativity.

Table 4.40: Multivariate Between-Subjects (Research Question 3), Source: author's SPSS output, 2025

Tests of Between-Subjects Effects							
Source	Dependent Variable	Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared
Corrected Model	CRIT_Overall_Adoption_Opp	19.848 ^a	9	2.205	2.756	.009	.293
	Support_Overall_Adoption_Opp	14.340 ^b	9	1.593	2.362	.023	.262
	Creativity_Overall_Adoption_Opp	21.232 ^c	9	2.359	3.112	.004	.318
Intercept	CRIT_Overall_Adoption_Opp	13.338	1	13.338	16.669	<.001	.217
	Support_Overall_Adoption_Opp	16.900	1	16.900	25.050	<.001	.295
	Creativity_Overall_Adoption_Opp	12.144	1	12.144	16.019	<.001	.211
Acc_Gap	CRIT_Overall_Adoption_Opp	.218	1	.218	.273	.603	.005
	Support_Overall_Adoption_Opp	.006	1	.006	.009	.924	.000
	Creativity_Overall_Adoption_Opp	.400	1	.400	.527	.471	.009
Fai_Gap	CRIT_Overall_Adoption_Opp	.274	1	.274	.343	.560	.006
	Support_Overall_Adoption_Opp	.078	1	.078	.116	.735	.002
	Creativity_Overall_Adoption_Opp	.120	1	.120	.158	.692	.003
Trans_Gap	CRIT_Overall_Adoption_Opp	3.091	1	3.091	3.864	.054	.060
	Support_Overall_Adoption_Opp	3.676	1	3.676	5.448	.023	.083
	Creativity_Overall_Adoption_Opp	5.262	1	5.262	6.941	.011	.104
Expl_Gap	CRIT_Overall_Adoption_Opp	.149	1	.149	.186	.668	.003
	Support_Overall_Adoption_Opp	.826	1	.826	1.225	.273	.020
	Creativity_Overall_Adoption_Opp	1.360	1	1.360	1.793	.186	.029
Res_Gap	CRIT_Overall_Adoption_Opp	.004	1	.004	.005	.946	.000
	Support_Overall_Adoption_Opp	.042	1	.042	.062	.805	.001
	Creativity_Overall_Adoption_Opp	.357	1	.357	.471	.495	.008
Pri_Gap	CRIT_Overall_Adoption_Opp	.338	1	.338	.423	.518	.007
	Support_Overall_Adoption_Opp	.075	1	.075	.111	.740	.002
	Creativity_Overall_Adoption_Opp	.094	1	.094	.124	.726	.002
Safe_Gap	CRIT_Overall_Adoption_Opp	.643	1	.643	.803	.374	.013
	Support_Overall_Adoption_Opp	.498	1	.498	.739	.394	.012
	Creativity_Overall_Adoption_Opp	1.821	1	1.821	2.402	.126	.038

Table 4.40: Multivariate Between-Subjects (Research Question 3), Source: author's SPSS output, 2025 (continued)

Sec_Gap	CRIT_Overall_Adoption_Opp	2.461E-5	1	2.461E-5	.000	.996	.000
	Support_Overall_Adoption_Opp	.078	1	.078	.115	.735	.002
	Creativity_Overall_Adoption_Opp	.340	1	.340	.448	.506	.007
Sus_Gap	CRIT_Overall_Adoption_Opp	.530	1	.530	.662	.419	.011
	Support_Overall_Adoption_Opp	.121	1	.121	.179	.674	.003
	Creativity_Overall_Adoption_Opp	.039	1	.039	.051	.822	.001
Error	CRIT_Overall_Adoption_Opp	48.008	60	.800			
	Support_Overall_Adoption_Opp	40.479	60	.675			
	Creativity_Overall_Adoption_Opp	45.487	60	.758			
Total	CRIT_Overall_Adoption_Opp	903.716	70				
	Support_Overall_Adoption_Opp	876.438	70				
	Creativity_Overall_Adoption_Opp	891.864	70				
Corrected Total	CRIT_Overall_Adoption_Opp	67.856	69				
	Support_Overall_Adoption_Opp	54.819	69				
	Creativity_Overall_Adoption_Opp	66.720	69				
a. R Squared = .293 (Adjusted R Squared = .186)							
b. R Squared = .262 (Adjusted R Squared = .151)							
c. R Squared = .318 (Adjusted R Squared = .216)							

Tests of Between-Subjects Effects: This is a series of univariate ANOVAs, one for each DV. For a significant IV from the Multivariate test, this table shows *which specific adoption context* (e.g., Critical vs. Support) is being driven by that IV. The Partial Eta Squared (η^2) column tells the effect size (e.g., .01=small, .06=medium, .14=large).

Key findings: The Transparency gap has the highest impact on adoption in the Creativity application ($\eta^2 = .104$), followed by Business Support ($\eta^2 = .083$) and Business-Critical applications ($\eta^2 = .06$). The accountability and Resilience gaps have the lowest impact on adoption across all three applications.

4.4 Research Question Four

RQ4 What copyright, privacy, bias, toxicity, and misinformation risks are introduced by the input data during the adoption for productivity and innovation use cases by IT service companies?

The adoption of generative AI by IT service companies for productivity and innovation is a complex process. It involves a multi-stage supply chain in which data risks are introduced at different stages. Understanding these risks required a thorough examination of the generative AI value chain, from data collection to end-user application.

Before delving into specific risks, it is crucial to understand the structure of the generative AI ecosystem. Unlike traditional software, generative AI is a supply chain involving multiple actors, as illustrated in Figure 4.0, the Generative AI Value Chain (*Lee et al., 2024*).

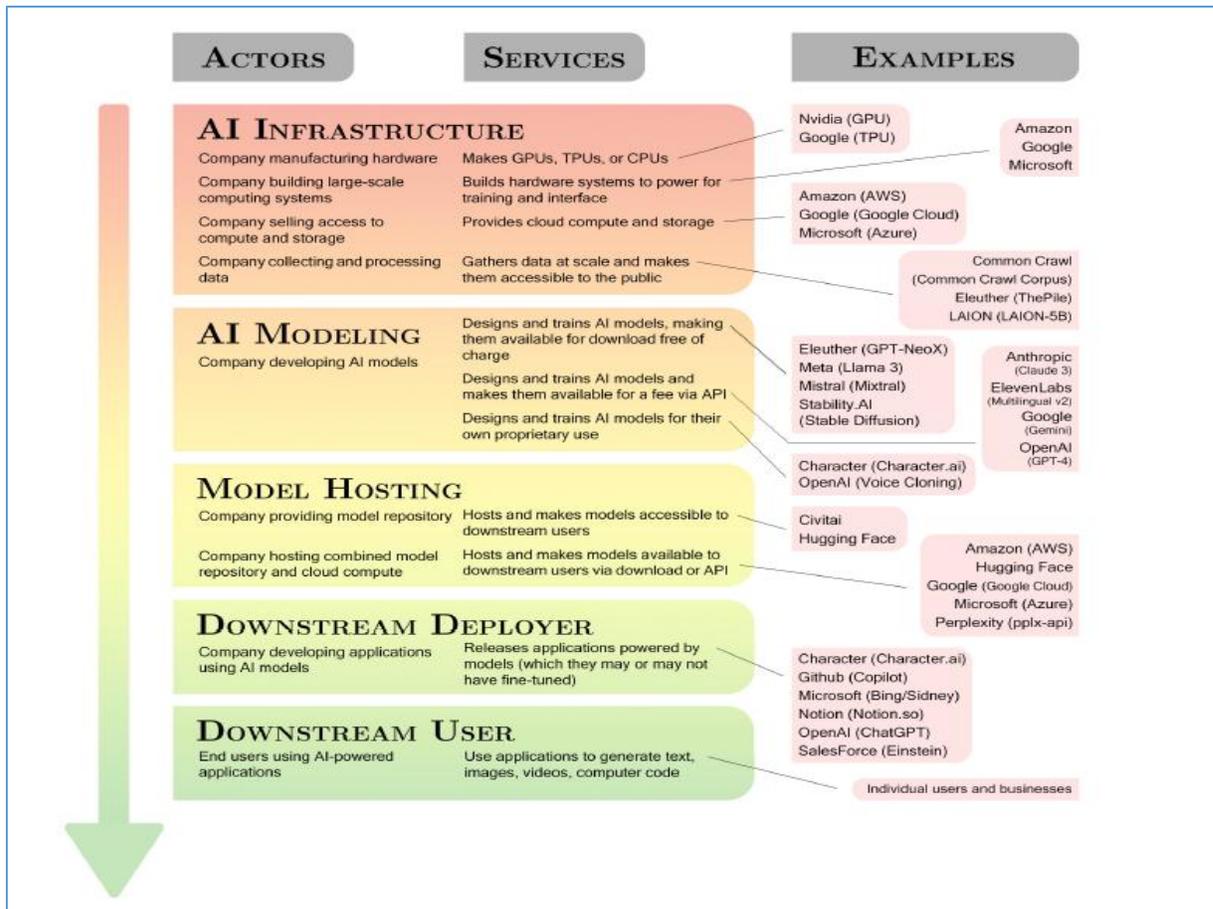


Figure 4.0: Generative AI Value Chain, Source: (Lee et al., 2024)

The value chain includes:

Data Collection Companies - Entities like EleutherAI and LAION that scrape and aggregate massive datasets from the internet, including text, images, code, and other digital content. These datasets serve as the raw material for pre-training foundational models. Data collection companies may not curate data to address different risks, including privacy, copyright, and bias.

AI Modeling Companies - Organizations like OpenAI, Google DeepMind, and Meta that develop foundation models. They train these models on the collected datasets, a process requiring immense computational resources.

Downstream Deployers (IT Service Companies) - This is the stage where most service companies operate. They take a pre-trained foundation model and fine-tune it for specific tasks (e.g., customer service automation, code generation, document summarization). This fine-tuning may involve adding proprietary or third-party data.

AI Users - The end-users who interact with the application built by the deployer. They provide prompts, which constitute new input data, and receive generated outputs.

This layered value chain creates a mixed set of responsibilities and potential liabilities. A risk originating in the data collection phase can manifest in the output seen by an end-user. Since the data has moved through multiple layers, it's hard to assign accountability. "The complexity of copyright law in the context of generative AI systems is profound, and these issues are intricately connected" (*Lee et al., 2024*). This interconnectedness is equally true for other data-related risks like privacy, bias, toxicity, and misinformation. Each input data risks identified in the Literature Review section are analyzed in detail here to understand in IT service context.

Copyright Infringement: Copyright risk is arguably the most legally contentious issue in the generative AI supply chain. It permeates multiple stages of the value chain. The core of the problem lies in training, both in creating the foundation model and in subsequent fine-tuning. The end user can further add new data to a fine-tuned model that may have copyrighted data. As defined, an AI model is "a program trained on a large set of data with the ability to identify patterns in that data to produce relevant outputs" (*G'sell, 2024*). When a large dataset includes copyrighted material scraped from the web without explicit permission, the act of training itself may constitute copyright infringement.

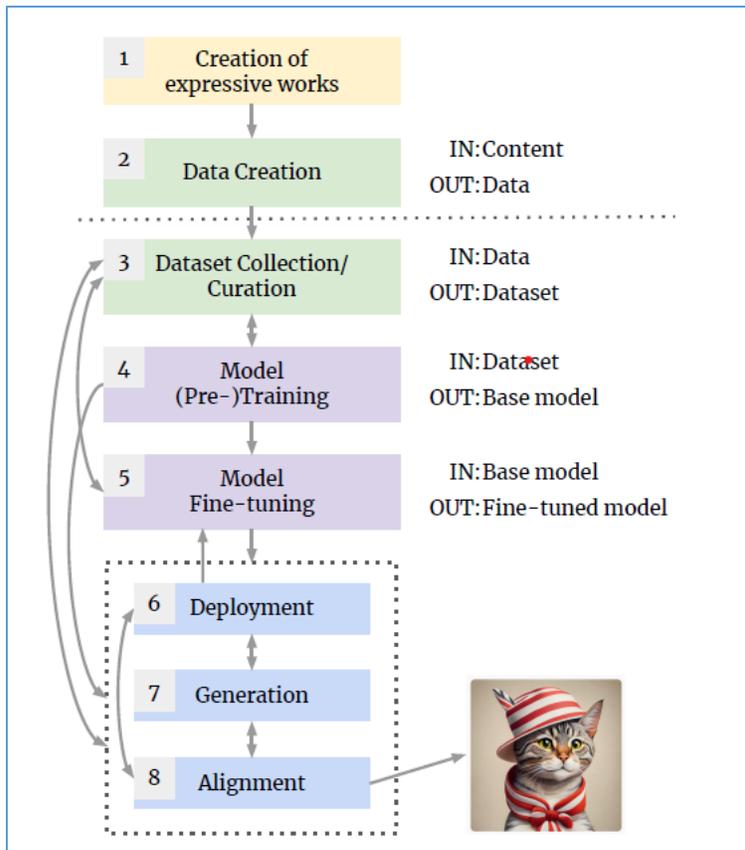


Figure 3.2: AI value chain, Source (Paul & Sarkar, 2023)

For better readability, Figure 3.2 is re-illustrated here.

Risks at the Data Collection and Model Training Stages (Stages 1-4):

Data collection companies and AI modelers introduce the initial and most significant copyright risk. Foundation models are "trained on an extensive amount of data, often collected from the internet through web scraping" (Gutierrez *et al.*, 2022). This practice, while efficient, rarely involves seeking licenses from the copyright holders of the billions of works ingested (Lee *et al.*, 2024). The legal argument often used by AI companies is that of fair use, which permits limited use of copyrighted material. The companies argue that training an AI to understand statistical patterns in language and images is a transformative fair use (Lee *et al.*, 2024).

However, this has not yet been legally tested. The outcome of ongoing copyright infringement lawsuits will hinge on questions such as whether copying entire works for commercial training purposes is fair. Also, whether the resulting model, which can generate content similar to the training data, creates a substitute for the original works. The reliance on Internet sources and the scale of scraped datasets make it very challenging to determine the origins of individual data examples (*Lee et al., 2024*). This means that an IT service company fine-tuning a model has no clear way of verifying the legal provenance of the model's foundational knowledge.

Furthermore, the datasets themselves can be copyrighted. Copying the dataset as a whole without permission could constitute infringement, separate from any infringement on the underlying works the dataset comprises (*Lee et al., 2024*). This adds another layer of potential liability.

Risks at the Fine-Tuning and Deployment Stage (Stage 5-7):

This is the stage where IT service companies are most active. When an IT service company fine-tunes a foundation model for a specific client or use case, it typically introduces new data. If this fine-tuning dataset includes proprietary client data or third-party licensed data, the company must ensure it has the rights to use that data for this purpose. However, a more serious risk is that the fine-tuning process can amplify the model's propensity to reproduce copyrighted content used during initial training. Suppose a service company fine-tunes a model to write in a specific corporate tone. In that case, the model can generate text similar to a copyrighted news article from its training data, and the IT service company and its client could face liability.

Risks at the User Interaction Stage (Stage 8):

The user also introduces copyright risk. The liability of the creator of a generative-AI system can be influenced by the prompts that its users provide (*Lee et al., 2024*). For

example, a user might prompt a model with, "Write a story in the style of J.K. Rowling," or "Generate an image in the style of Picasso." While the model may not directly copy a specific work, its output could be considered a derivative work that infringes on the original artist's style. The IT service company that deployed the application may be held indirectly liable for facilitating this infringement.

To mitigate the risk of copyright infringement for their customers, some generative AI providers have offered indemnities to their users. Google, Microsoft, Amazon, Anthropic, and OpenAI, among others, have pledged to indemnify certain users against intellectual property claims arising from infringing outputs (G'sell, 2024). However, these indemnities often come with caveats, typically protecting only enterprise customers who use the model out of the box without fine-tuning. This leaves IT service companies that engage in custom fine-tuning potentially exposed. For these companies, conducting thorough due diligence on the foundation models they use and carefully managing the fine-tuning data and user prompts are critical business risk management activities.

Privacy Risks: The same data scraping practices that raise copyright concerns also create privacy risks. Foundation models are trained on a vast corpus of publicly available internet data. This includes large amounts of personal information, such as social media posts, blog comments, forum discussions, public reviews, and even information scraped from individual websites.

A key characteristic of foundational models is their ability to memorize elements of their training data. While they are designed to generalize patterns, studies have shown that they can regurgitate verbatim excerpts from their training sets. This can include personally identifiable information (PII), such as email addresses and home addresses, that was publicly available online. This risk is a direct consequence of the homogenization described by Bommasani et al. (2022), where a single model is

leveraged for countless tasks. An IT service company using a model for a simple task, such as email drafting, could inadvertently produce an output containing the private contact information of an individual whose data was scraped from a public forum years earlier.

Beyond simple memorization, the foundational models have inferential capabilities that pose a greater threat. By aggregating information from multiple sources in its training data, a model could potentially infer and disclose sensitive attributes about an individual. The model might infer and generate text about an individual's political affiliations by correlating a name with a location, employer, and social group from various data points. This violates the fundamental principles of data privacy and protection laws, such as the GDPR, which grant individuals the right to control how their personal data is processed.

During the Use Phase of the AI system, where it "influences the physical or virtual environment" user prompts become a source of privacy risk (*Russell et al., 2023*). IT Service companies building applications must consider that users might input sensitive, proprietary, or personal data into the prompt. For instance, an employee might paste a customer's personal details into a chatbot to generate a response. The service provider must have clear policies and technical safeguards in place. The policy can cover handling, storing, and potentially using this prompt data to train the model further. Unauthorized retention or use of this data could lead to significant data breaches and regulatory penalties on IT service companies.

Bias and Toxicity: Bias and toxicity are characteristic of the data on which they are trained. The internet, the primary source of training data, is inundated with societal prejudices, stereotypes, hate speech, and toxic language.

The emergence of model capabilities has also led to the unfortunate emergence of biased behaviors (*Bommasani et al., 2022*). If a model is trained on text where certain professions are predominantly associated with a specific gender or ethnicity, it will learn

and reproduce those associations. For example, a model might be more likely to generate an image of a CEO as a man or associate certain nationalities with negative stereotypes. This happens because the model is fundamentally a pattern-matching engine. It learns from the statistical correlations present in the data without the ethical understanding to reject harmful stereotypes. An AI chatbot can generate discriminatory responses or offensive content. This can potentially lead to reputational damage and legal challenges for the IT service company.

Models trained on unfiltered internet data learn to generate toxic, abusive, and harmful language. Foundational model developers implement guardrails or content moderation filters to block explicitly harmful outputs. These are imperfect and can often be circumvented by carefully crafted prompts. It's a practice known as "jailbreaking." The risk is that a user of an IT service company's application, either accidentally or deliberately, could prompt the AI to generate offensive content. The IT service company could be held responsible for this output, especially if its own fine-tuning process weakened the model's original safety filters.

Mitigating bias and toxicity is challenging due to the opacity of foundation models. "Since the power of foundation models stems from their emergent qualities rather than their explicit construction, existing foundation models are challenging to understand, and they exhibit unexpected failure modes" (*Bommasani et al., 2022*). An IT service company cannot simply remove bias from a model. Bias and toxicity mitigation require continuous effort through techniques such as careful curation of fine-tuning data, reinforcement learning from human feedback (RLHF), and robust output testing. However, there is a constant risk of introducing new biases. This may happen because a model developer fails to anticipate how the model might express biases ingrained during its initial large-scale training.

Misinformation: Perhaps the most widely recognized risk of generative AI is its tendency to generate believable but entirely incorrect information. Hallucination is not a system malfunction but an inherent characteristic of how generative AI works. These foundational models are not databases of facts or reasoning engines. They are sophisticated systems that generate content by predicting the next most statistically likely word or pixel based on their training. Foundational models are optimized for coherence and plausibility rather than factual accuracy. They can generate citations to non-existent academic papers or provide incorrect medical or legal advice with confidence. This poses a direct threat to productivity and innovation use cases. Accuracy is critical, and misinformation can easily erode end users' trust.

The risk is magnified by the technology's ability to generate misinformation at an unprecedented scale and with high persuasiveness. A single individual can use these tools to create thousands of convincing, but fake, news articles and social media posts. For IT service companies, the risk is twofold. First, their internal use of AI for research or content creation could lead to the dissemination of false information within the organization. This can result in poor decision making. Second, if an AI-powered application they built for a client provides incorrect information that causes financial or physical harm. The IT service company could face significant liability.

When IT service companies integrate generative AI into their products, they are exposing their credibility to the technology. If the AI frequently hallucinates or generates false content, it damages the company's brand and trustworthiness. Managing this risk requires implementing clear disclaimers and human-in-the-loop verification processes. Also, avoiding deployment in high-stakes scenarios where accuracy is non-negotiable.

Conclusion: The analysis clearly demonstrates that the risks of copyright infringement, privacy violation, bias, toxicity, and misinformation are not discrete issues

but are deeply intertwined. They are fundamentally rooted in the generative AI supply chain. They are direct consequences of the core generative AI technological landscape, in which foundational models are opaque and trained on massive public datasets. The homogenization of using a few foundation models for many tasks creates concentrated points of failure. These risks are difficult to predict and control.

For IT service companies seeking to adopt generative AI to boost productivity and innovation, these data-related risks require a holistic risk-management approach. It is not sufficient to evaluate an AI tool's final output. A comprehensive due diligence process for internal AI governance must be undertaken. This should include the provenance of Foundation Models and the training practices of model providers. A robust method to ensure that data used for fine-tuning is legally and ethically sourced. A thought-through implementation of guardrails to filter harmful outputs and verify critical information before sharing it.

The generative AI value chain is a platform for innovation and transformation, as it is evident from its worldwide adoption in the last few years. However, it also serves as a platform for amplifying existing legal and ethical challenges. Navigating this platform requires a clear understanding that the input data entering the chain directly determines the nature of the risks emerging as output.

4.5 Research Question Five

RQ5 How much input-data-generated copyright, biases, toxicity, and misinformation perceived risks are covered, and how are they affecting the adoption of productivity and innovation use cases by IT service companies?

This research question has two parts: A) perceived risk coverage as IT service companies implement generative AI applications, and B) impact on adoption, as risk coverage may not be sufficient.

The survey was conducted with participants who have implemented generative AI applications.

A) Perceived risk coverage

The core objective was to measure the perceived risk coverage (adequacy) of the generative AI systems implemented by IT Service and Product companies. The output is a quantified risk assurance gap analysis. During the literature review, five key input data risk domains (copyright, privacy, bias, toxicity, and misinformation) were identified. The survey respondents were asked to rate their current risk gap perception across three distinct application areas (Business-critical, Business-support, and Creativity) for five key data risks.

Independent Variable (IV) - Risk Domain. This is a categorical variable with five levels: Copyright Risk, Privacy Risk, Bias Risk, Toxicity Risk, and Misinformation Risk.

Dependent Variable (DV) - Perceived Risk Coverage Adequacy. This was the ordinal-scale response collected for each risk domain, measured on a 5-point Likert scale (1 = Strongly Disagree to 5 = Strongly Agree) regarding whether the risk is fully covered.

Descriptive Statistics were used for each of the five risk domains to calculate the mean, median, mode, and standard deviation for the perceived risk coverage score. A higher mean score (e.g., closer to 5) for a specific risk (e.g., Bias) suggests that, on average, survey respondents believe existing frameworks do not adequately cover it.

Inferential Statistics were used to determine if the differences in mean scores between the risk domains are statistically significant (i.e., not due to random chance). A Repeated Measures ANOVA was used. Repeated measures are used because the survey respondents' responses are measured against the same dependent variable (perceived risk coverage) across multiple conditions (the five risk domains).

Descriptive Statistics for Business-critical Applications

Table 4.41: Critical Descriptive Statistics (Research Question 5), Source: author's SPSS output, 2025

		Statistics				
		CRIT_Copy_Gap	CRIT_Pri_Gap	CRIT_Bias_Gap	CRIT_Toxi_Gap	CRIT_Mis_Gap
N	Valid	70	70	70	70	70
	Missing	0	0	0	0	0
Mean		3.80	3.63	3.57	3.33	3.73
Median		4.00	4.00	4.00	3.00	4.00
Std. Deviation		.878	.995	.972	.959	.931
Skewness		-.386	-.186	-.012	.304	-.312
Std. Error of Skewness		.287	.287	.287	.287	.287
Kurtosis		-.446	-.977	-.961	-.789	-.700
Std. Error of Kurtosis		.566	.566	.566	.566	.566

Refer to Table 4.41. Copyright risk has the largest assurance gap (mean = 3.80, skewness = -0.386), suggesting that, for business-critical applications, IT service companies have the least coverage for this risk. Misinformation has the second-highest risk assurance gap, indicating that a significant number of implementers of business-critical applications lack adequate risk coverage for it.

Table 4.42: Critical Copyright (Research Question 5), Source: author's SPSS output, 2025

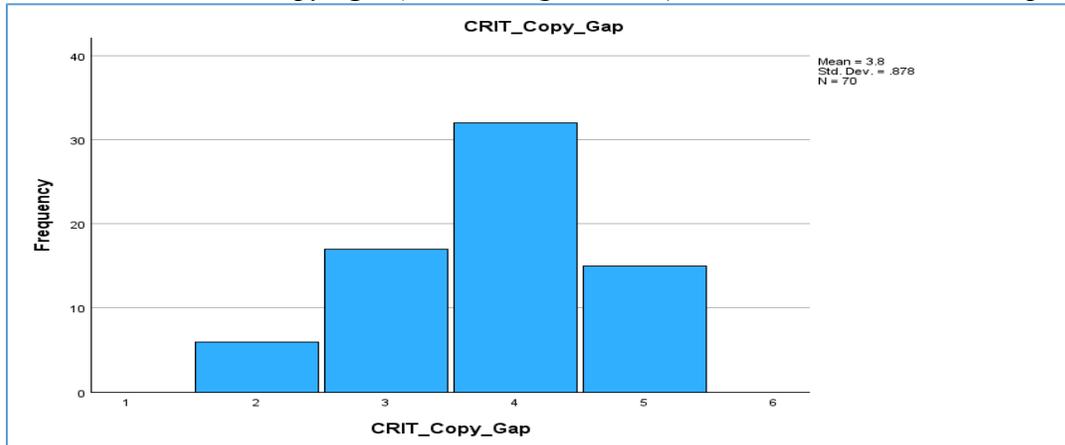
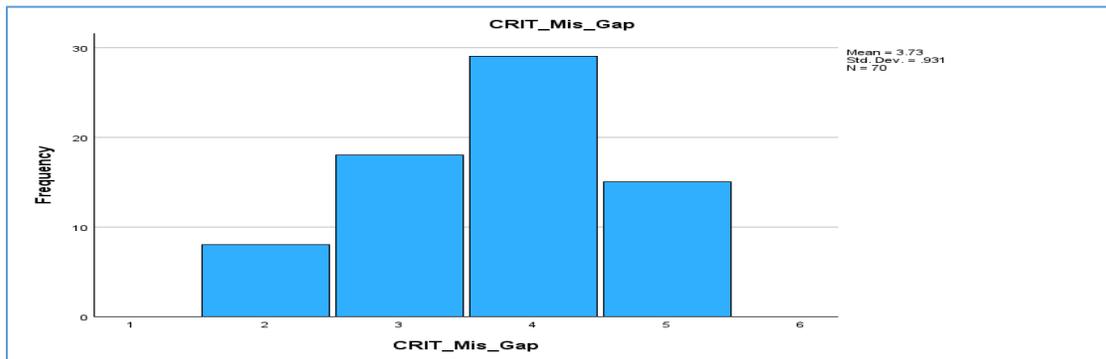


Table 4.43: Critical Misinformation (Research Question 5), Source: author’s SPSS output, 2025



Next, an overall risk assurance gap for all five risk domains together was calculated using descriptive statistics for business-critical applications.

Table 4.44: Critical Overall (Research Question 5), Source: author’s SPSS output, 2025

Statistics		
CRIT_Overall_Risk_Gap		
N	Valid	70
	Missing	0
Mean		3.6114
Median		3.6000
Std. Deviation		.71008
Skewness		-.038
Std. Error of Skewness		.287
Kurtosis		-.701
Std. Error of Kurtosis		.566

The overall risk assurance was 3.61, indicating that across the five risk domains, IT service companies do not adequately address risks when implementing business-critical applications.

Inferential Statistics for Business-Critical Applications

The purpose of the Reliability analysis test was to determine if all five risk items (copyright, privacy, bias, toxicity, and misinformation) reliably measured the same underlying concept, Risk Assurance.

Using SPSS Reliability Analysis, Cronbach’s Alpha was calculated. Refer to Table 4.45.

Table 4.45: Critical Reliability (Research Question 5), Source: author's SPSS output, 2025

Reliability Statistics	
Cronbach's Alpha	N of Items
.805	5

The Cronbach's Alpha score was 0.805, which is above 0.7, considered an acceptable level of reliability.

A *One-Sample T-Test* was used to determine whether the mean scores for the composite, CRIT_Overall_Risk_Gap, and the individual risks (copyright, privacy, bias, toxicity, and misinformation) are statistically significantly different from a neutral value of 3.

Table 4.46: Critical One-Sample (Research Question 5), Source: author's SPSS output, 2025

One-Sample Test							
Test Value = 3							
	t	df	Significance		Mean Difference	95% Confidence Interval of the Difference	
			One-Sided p	Two-Sided p		Lower	Upper
CRIT_Copy_Gap	7.623	69	<.001	<.001	.800	.59	1.01
CRIT_Pri_Gap	5.284	69	<.001	<.001	.629	.39	.87
CRIT_Bias_Gap	4.920	69	<.001	<.001	.571	.34	.80
CRIT_Toxi_Gap	2.867	69	.003	.005	.329	.10	.56
CRIT_Mis_Gap	6.545	69	<.001	<.001	.729	.51	.95
CRIT_Overall_Risk_Gap	7.204	69	<.001	<.001	.61143	.4421	.7807

All Risk items, including CRIT_Overall_Risk_Gap, have a Two-Sided value (the p-value) less than .05, which implies that the mean score is significantly different from neutral. Toxicity has the lowest mean score (3.33) and a lower p-value (indicating lower confidence), suggesting that we must exercise caution when interpreting the risk assurance score and the conclusions drawn from it.

Descriptive Statistics for Business-Support Applications

Table 4.47: Support Statistics (Research Question 5), Source: author's SPSS output, 2025

		Statistics				
		Support_Copy _Gap	Support_Pri_G ap	Support_Bias_ Gap	Support>Toxi_ Gap	Support_Mis_ Gap
N	Valid	70	70	70	70	70
	Missing	0	0	0	0	0
Mean		3.69	3.63	3.71	3.50	3.76
Median		4.00	4.00	4.00	3.00	4.00
Std. Deviation		.925	.966	.887	1.018	.999
Skewness		-.341	-.376	-.168	-.042	-.656
Std. Error of Skewness		.287	.287	.287	.287	.287
Kurtosis		-.127	-.341	-.683	-.730	-.186
Std. Error of Kurtosis		.566	.566	.566	.566	.566

Refer to Table 4.47. The risk of Misinformation has the largest assurance gap (mean = 3.76, skewness = -0.656), indicating that, for business-support applications, IT service companies have the least coverage against misinformation. Bias has the second-highest risk assurance gap, suggesting that a greater proportion of implementers of business-support applications have poorer risk coverage.

Table 4.48: Support Misinformation (Research Question 5), Source: author's SPSS output, 2025

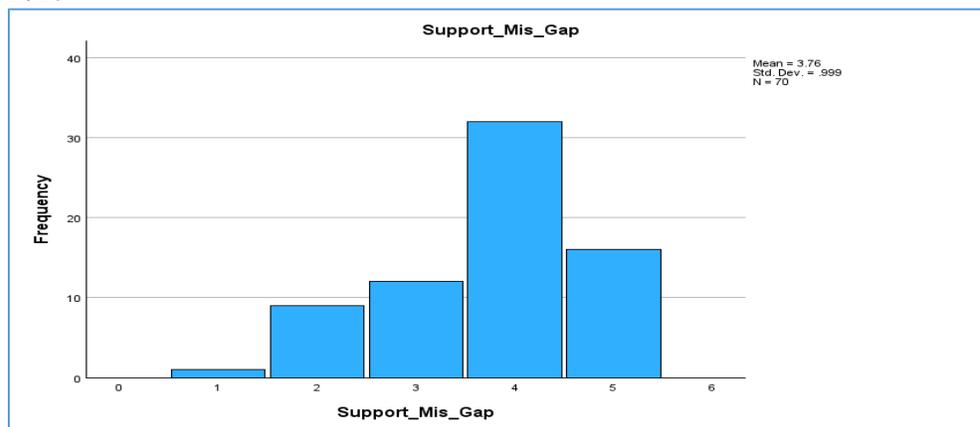
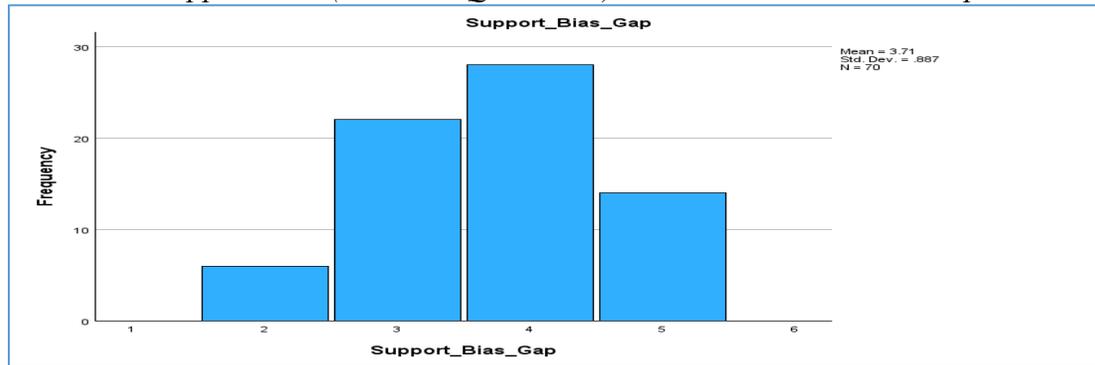


Table 4.49: Support Bias (Research Question 5), Source: author's SPSS output, 2025



Next, an overall risk assurance gap for all five risk domains was calculated using descriptive statistics for business-support applications.

Table 4.50: Support Overall (Research Question 5), Source: author's SPSS output, 2025

Statistics		
Support_Overall_Risk_Gap		
N	Valid	70
	Missing	0
Mean		3.6571
Median		3.8000
Std. Deviation		.74689
Skewness		-.320
Std. Error of Skewness		.287
Kurtosis		-.420
Std. Error of Kurtosis		.566

The overall risk assurance was 3.65, indicating that across the five risk domains, IT service companies do not adequately address risks when implementing business-support applications.

Inferential Statistics for Business-Support Applications

The purpose of the Reliability Analysis test was to determine if all five risk items (copyright, privacy, bias, toxicity, and misinformation) reliably measured the same underlying concept (“Risk Assurance”).

Using SPSS Reliability Analysis, Cronbach’s Alpha score was calculated.

Table 4.51: Support Reliability (Research Question 5), Source: author's SPSS output, 2025

Reliability Statistics	
Cronbach's Alpha	N of Items
.837	5

Refer to Table 4.51. The Cronbach's Alpha score was 0.837, which is above 0.7, considered an acceptable level of reliability.

One-Sample T-Test - The purpose of this test was to determine if the mean score of composite, Support_Overall_Risk_Gap, and all the individual risks (copyright, privacy, bias, toxicity, and misinformation) are statistically significantly different from a neutral value of 3.

Using SPSS, One-Sample T-Test was conducted

Table 4.52: Support One-Sample (Research Question 5), Source: author's SPSS output, 2025

One-Sample Test							
Test Value = 3							
	t	df	Significance		Mean Difference	95% Confidence Interval of the Difference	
			One-Sided p	Two-Sided p		Lower	Upper
Support_Copy_Gap	6.200	69	<.001	<.001	.686	.47	.91
Support_Pri_Gap	5.446	69	<.001	<.001	.629	.40	.86
Support_Bias_Gap	6.738	69	<.001	<.001	.714	.50	.93
Support_Toxi_Gap	4.110	69	<.001	<.001	.500	.26	.74
Support_Mis_Gap	6.341	69	<.001	<.001	.757	.52	1.00
Support_Overall_Risk_Gap	7.361	69	<.001	<.001	.65714	.4791	.8352

Refer to Table 4.52. All Risk items, including Support_Overall_Risk_Gap, have a Two-Sided value (the p-value) less than .05, which implies that the mean score is significantly different from neutral.

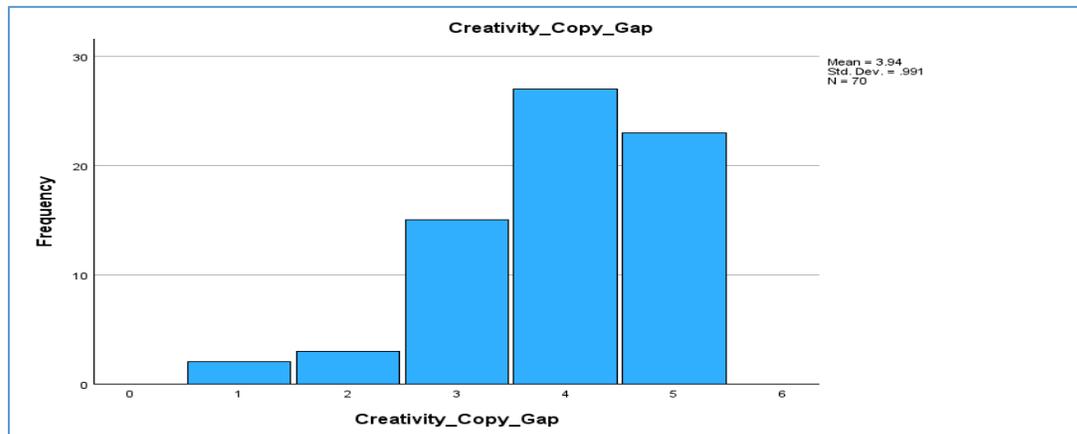
Descriptive Statistics for Creativity Applications

Table 4.53: Creativity Statistics (Research Question 5), Source: author's SPSS output, 2025

		Statistics				
		Creativity_Copy_Gap	Creativity_Pri_Gap	Creativity_Bias_Gap	Creativity_Toxi_Gap	Creativity_Mis_Gap
N	Valid	70	70	70	70	70
	Missing	0	0	0	0	0
Mean		3.94	3.90	3.90	3.87	3.86
Median		4.00	4.00	4.00	4.00	4.00
Std. Deviation		.991	1.079	.871	1.006	1.011
Skewness		-.895	-.794	-.345	-.262	-.744
Std. Error of Skewness		.287	.287	.287	.287	.287
Kurtosis		.689	-.036	-.610	-1.178	.279
Std. Error of Kurtosis		.566	.566	.566	.566	.566

Refer to Table 4.53. Copyright risk has the largest assurance gap (mean = 3.94, skewness = -0.895), indicating that IT service companies have the least coverage for copyright protection in creative applications. Privacy and Bias have the second-highest risk assurance gap, meaning that more implementers of creativity applications have poorer risk coverage.

Table 4.54: Creativity Copyright (Research Question 5), Source: author's SPSS output, 2025



Next, an overall risk assurance gap for all five risk domains was calculated using descriptive statistics for creativity applications.

Table 4.55: Creativity Overall (Research Question 5), Source: author's SPSS output, 2025

Statistics		
Creativity_Overall_Risk_Gap		
N	Valid	70
	Missing	0
Mean		3.8943
Median		4.0000
Std. Deviation		.77401
Skewness		-1.043
Std. Error of Skewness		.287
Kurtosis		.777
Std. Error of Kurtosis		.566

The overall risk assurance was 3.89, indicating that across the five risk domains, IT service companies do not adequately cover risks when implementing creative applications.

Inferential Statistics for Creativity Applications

The purpose of the Reliability Analysis test was to determine if all five risk items (copyright, privacy, bias, toxicity, and misinformation) reliably measured the same underlying concept (Risk Assurance).

Using SPSS, Reliability Analysis was conducted.

Table 4.56: Creativity Reliability (Research Question 5), Source: author's SPSS output, 2025

Reliability Statistics	
Cronbach's Alpha	N of Items
.838	5

Refer to Table 4.56. The Cronbach's Alpha score was 0.838, which is above 0.7, considered an acceptable level of reliability.

The purpose of the One-Sample T-Test was to determine if the mean score of the composite, Creativity_Overall_Risk_Gap, and all the individual risks (copyright, privacy, bias, toxicity, and misinformation) are statistically significantly different from a neutral value of 3.

Using SPSS, One-Sample T-Test was conducted.

Table 4.57: Creativity One-Sample (Research Question 5), Source: author's SPSS output, 2025

One-Sample Test							
Test Value = 3							
	t	df	Significance		Mean Difference	95% Confidence Interval of the Difference	
			One-Sided p	Two-Sided p		Lower	Upper
Creativity_Copy_Gap	7.960	69	<.001	<.001	.943	.71	1.18
Creativity_Pri_Gap	6.980	69	<.001	<.001	.900	.64	1.16
Creativity_Bias_Gap	8.649	69	<.001	<.001	.900	.69	1.11
Creativity_Toxi_Gap	7.247	69	<.001	<.001	.871	.63	1.11
Creativity_Mis_Gap	7.091	69	<.001	<.001	.857	.62	1.10
Creativity_Overall_Risk_Gap	9.667	69	<.001	<.001	.89429	.7097	1.0788

Refer to Table 4.57. All Risk items, including Creativity_Overall_Risk_Gap, have a Two-Sided value (the p-value) less than .05, which implies that the mean score is significantly different from neutral.

Impact on Adoption: Another core objective was to determine if the perceived lack of regulatory coverage for specific generative AI risks (copyright, privacy, bias, toxicity, and misinformation) acts as a catalyst or a barrier to the adoption of these solutions.

Each risk (Copyright, Privacy, Bias, Toxicity, and Misinformation) has a pair of competing hypotheses. It's crucial to note that alternative hypotheses are directional (they posit a "positive effect").

- Null Hypothesis (H_0) for each risk: The lack of well-defined, regulated risk [X] does not affect adoption. The mean score equal to or less than the neutral point (3) implies no effect or an adverse effect on adoption.
- Alternative Hypothesis (H_1) for each risk: The lack of regulated risk [X] positively affects (increases) adoption (i.e., it's seen as an opportunity, not a barrier). The mean score significantly greater than the neutral point (3) implies a positive effect.

The survey question was designed to gather responses for business-critical, business-support, and creativity (innovation) applications separately.

Hence, the analysis is done by application area (Business-critical, Business-support, and Creativity).

Business-critical application of generative AI

Using SPSS and Descriptive Statistics, Mean, Median, Standard Deviation, Skewness, and Kurtosis were calculated for each risk area.

Table 4.58: Critical Statistics (Research Question 5), Source: author's SPSS output, 2025

		Statistics				
		CRIT_Copy_Opp	CRIT_Pri_Opp	CRIT_Bias_Op	CRIT_Toxi_Op	CRIT_Mis_Opp
N	Valid	70	70	70	70	70
	Missing	0	0	0	0	0
Mean		3.76	3.80	3.64	3.54	3.76
Median		4.00	4.00	4.00	3.50	4.00
Std. Deviation		.984	1.150	.948	.928	.970
Skewness		-.802	-.537	-.477	-.297	-.568
Std. Error of Skewness		.287	.287	.287	.287	.287
Kurtosis		.455	-.708	.291	.270	.300
Std. Error of Kurtosis		.566	.566	.566	.566	.566

Refer to Table 4.58. In all risk areas, the mean score exceeds 3, invalidating the null hypothesis. There is a general perception that the lack of strict regulations across risk areas is actually positively affecting adoption.

Of the five risk areas, Privacy (mean score 3.80), Copyright (mean score 3.76), and Misinformation (mean score 3.76) are the top areas experiencing greater adoption due to lax regulations. Privacy (skewness score $-.802$, SD) exhibits a high level of skewness. A high negative skewness indicates a left-tailed distribution, meaning the mean is less than the median.

Table 4.59: Critical Privacy (Research Question 5), Source: author's SPSS output, 2025

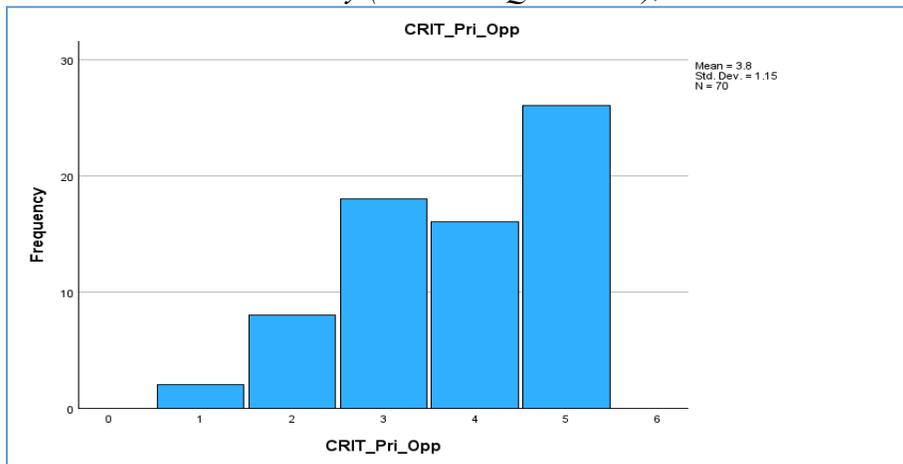


Table 4.60: Critical Copyright (Research Question 5), Source: author's SPSS output, 2025

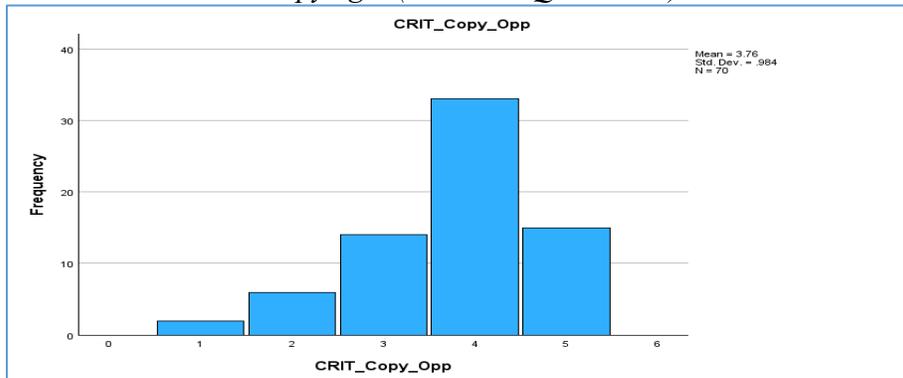
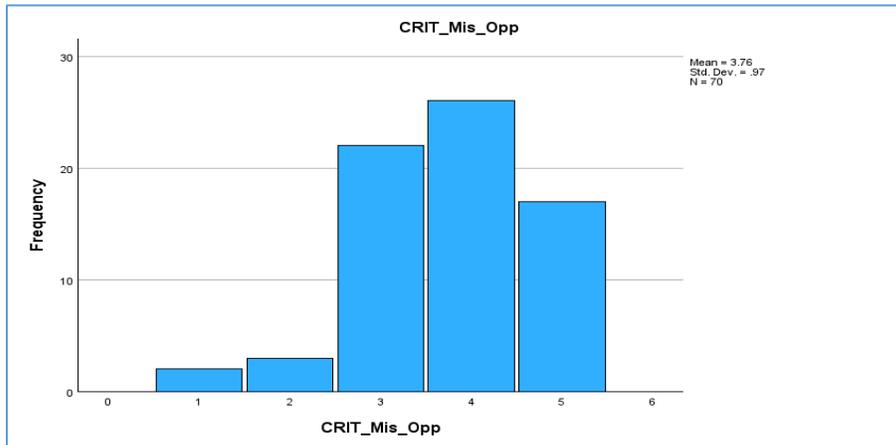


Table 4.61: Critical Misinformation (Research Question 5), Source: author's SPSS output, 2025



Inferential Analysis (Hypothesis Testing) was used to statistically test if the mean score for each risk is significantly greater than the neutral value of 3.

A one-sample T-test compares the mean of a single sample (e.g., the response for CRIT_Privacy_Opp) to a test value 3 (the neutral point). Using SPSS and a One-Sample T-Test, a two-tailed p-value was calculated. Since the hypothesis is one-tailed (directional), the reported 2-tailed value is divided by 2 to get the corrected one-tailed p-value.

Table 4.62: Critical One-Sample (Research Question 5), Source: author's SPSS output, 2025

One-Sample Test							
Test Value = 3							
	t	df	Significance		Mean Difference	95% Confidence Interval of the Difference	
			One-Sided p	Two-Sided p		Lower	Upper
CRIT_Copy_Opp	6.435	69	<.001	<.001	.757	.52	.99
CRIT_Pri_Opp	5.822	69	<.001	<.001	.800	.53	1.07
CRIT_Bias_Opp	5.671	69	<.001	<.001	.643	.42	.87
CRIT_Toxi_Opp	4.896	69	<.001	<.001	.543	.32	.76
CRIT_Mis_Opp	6.533	69	<.001	<.001	.757	.53	.99
CRIT_Overall_5Risks_Opp	7.131	69	<.001	<.001	.70000	.5042	.8958

Refer to Table 4.62. The 1-tailed p-value (Two-sided p/2) is less than the alpha level (.05), so the null hypothesis is rejected. For a specific risk (e.g., Privacy), if the one-tailed p-value is < 0.05 and the Mean Difference is positive, then we can reject H_0 . There is significant evidence that the lack of regulatory oversight of privacy risks is perceived as providing greater opportunities for adoption.

Regression analysis for Business-critical applications

While the one-sample t-test tells us if the perception for each risk is significantly above neutral, regression can answer more complex questions. Can we predict an overall "Adoption Opportunity" score based on the perceived gaps in all five risks simultaneously? Which specific risk gaps are the strongest drivers of the adoption perception?

Using SPSS, a composite dependent variable, CRIT_Overall_5Risks_Opp, was created to build the linear regression model with five risk gaps.

Table 4.63: Critical Model Summary (Research Question 5), Source: author's SPSS output, 2025

Model Summary ^b				
Model	R	R Square	Adjusted R Square	Std. Error of the Estimate
1	.444 ^a	.197	.134	.76428

a. Predictors: (Constant), CRIT_Mis_Gap, CRIT_Copy_Gap, CRIT_Toxi_Gap, CRIT_Bias_Gap, CRIT_Pri_Gap
 b. Dependent Variable: CRIT_Overall_5Risks_Opp

Refer to Table 4.63. R^2 indicates the proportion of variance in the overall adoption opportunity score that is explained by all five perceived regulatory gaps combined. An R^2 of 0.134 suggests that the model accounts for only 13.4% of the variation in perceptions of adoption opportunities.

Table 4.64: Critical ANOVA (Research Question 5), Source: author's SPSS output, 2025

ANOVA ^a						
Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	9.156	5	1.831	3.135	.014 ^b
	Residual	37.384	64	.584		
	Total	46.540	69			

a. Dependent Variable: CRIT_Overall_5Risks_Opp
 b. Predictors: (Constant), CRIT_Mis_Gap, CRIT_Copy_Gap, CRIT_Toxi_Gap, CRIT_Bias_Gap, CRIT_Pri_Gap

Sig. (p-value): This tests whether the overall regression model is statistically significant. Our Sig. (p-value) The p-value is 0.014, which is less than 0.05; therefore, we can conclude that the combination of all five risk gaps (the independent variable) significantly predicts the dependent variable, CRIT_Overall_5Risks_Opp. It indicates that our model is superior to simply using the mean.

Table 4.65: Critical Coefficients (Research Question 5), Source: author's SPSS output, 2025

Coefficients ^a											
Model		Unstandardized Coefficients		Standardized Coefficients Beta	t	Sig.	95.0% Confidence Interval for B		Correlations		
		B	Std. Error				Lower Bound	Upper Bound	Zero-order	Partial	Part
1	(Constant)	1.890	.519		3.641	<.001	.853	2.927			
	CRIT_Copy_Gap	.311	.147	.332	2.110	.039	.016	.605	.370	.255	.236
	CRIT_Pri_Gap	-.037	.148	-.045	-.249	.804	-.333	.259	.326	-.031	-.028
	CRIT_Bias_Gap	-.019	.128	-.023	-1.51	.880	-.274	.236	.269	-.019	-.017
	CRIT_Toxi_Gap	.237	.140	.277	1.696	.095	-.042	.517	.333	.207	.190
	CRIT_Mis_Gap	.011	.115	.013	.099	.921	-.218	.241	.184	.012	.011

a. Dependent Variable: CRIT_Overall_5Risks_Opp

Based on the above Coefficient table (Table 4.65), the Copyright regulation gap has the highest Standardized Coefficient Beta, 0.332. t & Sig. (p-value) tests if the coefficient for each independent variable (regulation gaps) is significantly different from zero. A p-value < .05 means this specific regulatory gap is a unique, significant predictor of adoption opportunity. Copyright has the lowest Sig value of 0.039.

Regression analysis for Business-support applications

While the one-sample t-test tells us if the perception for each risk is significantly above neutral, regression can answer more complex questions. Can we predict an overall "Adoption Opportunity" score based on the perceived gaps in all five risks simultaneously?

Using SPSS, a composite dependent variable, Support_Overall_5Risks_Opp, was created to build the linear regression model with five risk gaps.

Table 4.66: Support Model Summary (Research Question 5), Source: author's SPSS output, 2025

Model Summary ^b				
Model	R	R Square	Adjusted R Square	Std. Error of the Estimate
1	.319 ^a	.102	.031	.54403

a. Predictors: (Constant), Support_Mis_Gap, Support_Copy_Gap, Support_Toxi_Gap, Support_Bias_Gap, Support_Pri_Gap

b. Dependent Variable: Support_Overall_5Risks_Opp

Refer to Table 4.66. R^2 indicates the proportion of variance in the overall adoption opportunity score that is explained by all five perceived regulatory gaps combined. An R^2 of 0.031 suggests that the model accounts for a mere 3.1% of the variation in perceptions of adoption opportunities.

Table 4.67: Support ANOVA (Research Question 5), Source: author's SPSS output, 2025

ANOVA ^a						
Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	2.142	5	.428	1.447	.220 ^b
	Residual	18.942	64	.296		
	Total	21.083	69			

a. Dependent Variable: Support_Overall_5Risks_Opp
b. Predictors: (Constant), Support_Mis_Gap, Support_Copy_Gap, Support_Toxi_Gap, Support_Bias_Gap, Support_Pri_Gap

Sig. (p-value): This tests whether the overall regression model is statistically significant. Our Sig. (p-value) Since the p-value is .220, which is greater than .05, we can conclude that the combination of all five risk gaps (the independent variable) does not significantly predict the dependent variable, Support_Overall_5Risks_Opp.

Table 4.68: Support Coefficients (Research Question 5), Source: author's SPSS output, 2025

Coefficients ^a											
Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.	95.0% Confidence Interval for B		Correlations		
		B	Std. Error	Beta			Lower Bound	Upper Bound	Zero-order	Partial	Part
1	(Constant)	3.070	.334		9.195	<.001	2.403	3.737			
	Support_Copy_Gap	.110	.090	.185	1.222	.226	-.070	.291	.277	.151	.145
	Support_Pri_Gap	.016	.098	.028	.163	.871	-.180	.212	.224	.020	.019
	Support_Bias_Gap	.009	.097	.015	.097	.923	-.185	.204	.209	.012	.012
	Support_Toxi_Gap	.086	.084	.159	1.027	.308	-.081	.253	.263	.127	.122
	Support_Mis_Gap	-.002	.085	-.004	-.027	.978	-.172	.168	.170	-.003	-.003

a. Dependent Variable: Support_Overall_5Risks_Opp

Based on Table 4.68, the Copyright regulation gap has the highest Standardized Coefficient Beta, 0.185. t & Sig. (p-value) tests if the coefficient for each independent variable

(regulation gaps) is significantly different from zero. A p-value < .05 indicates that this specific regulatory gap is a unique and significant predictor of adoption opportunity.

Regression analysis for Creativity applications

While the one-sample t-test tells us if the perception for each risk is significantly above neutral, regression can answer more complex questions. Can we predict an overall "Adoption Opportunity" score based on the perceived gaps in all five risks simultaneously?

Using SPSS, a composite dependent variable, Creativity_Overall_5Risks_Opp, was created to build the linear regression model with five risk gaps.

Table 4.69: Creativity Model Summary (Research Question 5), Source: author's SPSS output, 2025

Model Summary ^b				
Model	R	R Square	Adjusted R Square	Std. Error of the Estimate
1	.611 ^a	.373	.324	.60325

a. Predictors: (Constant), Creativity_Mis_Gap, Creativity_Bias_Gap, Creativity_Toxi_Gap, Creativity_Copy_Gap, Creativity_Pri_Gap
 b. Dependent Variable: Creativity_Overall_5Risks_Opp

R² indicates the proportion of variance in the overall adoption opportunity score that is explained by all five perceived regulatory gaps combined. Refer to Table 4.69. An R² of 0.324 suggests that the model accounts for 32.4% of the variation in adoption opportunity perceptions.

Table 4.70: Creativity ANOVA (Research Question 5), Source: author's SPSS output, 2025

ANOVA ^a						
Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	13.853	5	2.771	7.613	<.001 ^b
	Residual	23.291	64	.364		
	Total	37.143	69			

a. Dependent Variable: Creativity_Overall_5Risks_Opp
b. Predictors: (Constant), Creativity_Mis_Gap, Creativity_Bias_Gap, Creativity_Toxi_Gap, Creativity_Copy_Gap, Creativity_Pri_Gap

Sig. (p-value): This tests whether the overall regression model is statistically significant. Our Sig. (p-value) Since the p-value is less than 0.001 and also less than 0.05, we can conclude that the combination of all five risk gaps, the independent variable, significantly predicts the dependent variable, Creativity_Overall_5Risks_Opp.

Table 4.71: Creativity Coefficients (Research Question 5), Source: author's SPSS output, 2025

Coefficients ^a											
Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.	95.0% Confidence Interval for B		Correlations		
		B	Std. Error	Beta			Lower Bound	Upper Bound	Zero-order	Partial	Part
1	(Constant)	1.828	.382		4.782	<.001	1.064	2.592			
	Creativity_Copy_Gap	-.003	.108	-.004	-.027	.979	-.219	.214	.421	-.003	-.003
	Creativity_Pri_Gap	.139	.100	.205	1.394	.168	-.060	.339	.497	.172	.138
	Creativity_Bias_Gap	.089	.112	.106	.795	.430	-.135	.313	.375	.099	.079
	Creativity_Toxi_Gap	.015	.094	.021	.164	.870	-.172	.203	.336	.020	.016
	Creativity_Mis_Gap	.288	.095	.397	3.030	.004	.098	.478	.570	.354	.300

a. Dependent Variable: Creativity_Overall_5Risks_Opp

Based on the above Coefficient table (Table 4.71), the Copyright regulation gap has the highest Standardized Coefficient Beta, 0.397. t & Sig. (p-value) tests if the coefficient for each independent variable (regulation gaps) is significantly different from zero. A p-value < .05 indicates that this specific regulatory gap is a unique and significant predictor of adoption opportunity. Misinformation has a Sig value less than 0.05 and could be a significant predictor of adoption opportunity.

Multivariate Analysis of Variance (MANOVA): This is a classic scenario for Multivariate Analysis of Variance (MANOVA). We are dealing with three distinct dependent variables (DVs), each measured across the same set of five independent variables (IVs) for the same respondents.

The goal is to understand how perceptions of regulatory gaps (IVs) collectively and individually influence the propensity to adopt across three application contexts (DVs): Business-critical, Business-support, and Creativity.

MANOVA assesses the effect of multiple IVs on multiple related DVs simultaneously. It controls inflated Type 1 error that may occur if three separate ANOVAs are run, as was done in this case. It can also test if IVs affect the DVs differently.

Using SPSS, Multivariate analysis was conducted.

Data Prep work: Three sets of variables for each of the five risk areas

Independent Variables (IVs - Predictors): The perception of the regulatory gap.

- CRIT_Copy_gap, CRIT_Pri_gap, CRIT_Bias_gap, CRIT_Toxi_gap, CRIT_Mis_gap
- Support_Copy_gap, Support_Pri_gap, Support_Bias_gap, Support_Toxi_gap, Support_Mis_gap
- Creativity_Copy_gap, Creativity_Pri_gap, Creativity_Bias_gap, Creativity_Toxi_gap, Creativity_Mis_gap

Dependent Variables (DVs - Outcomes): The perception of adoption opportunity for each context

- Business-Critical: CRIT_Copy_Opp, CRIT_Pri_Opp, CRIT_Bias_Opp, CRIT_Toxi_Opp, CRIT_Mis_Opp
- Business-Support: Support_Copy_Opp, Support_Pri_Opp, Support_Bias_Opp, Support_Toxi_Opp, Support_Mis_Opp

- Creativity: Creativity_Copy_Opp, Creativity_Pri_Opp,
Creativity_Bias_Opp, Creativity_Toxi_Opp, Creativity_Mis_Opp

Composite Scores: Three overall adoption scores, one for each context. This reduces complexity and provides a precise outcome measure.

- CRIT_Overall_5Risks_Opp: Mean(CRIT_Copy_Opp, CRIT_Pri_Opp, CRIT_Bias_Opp, CRIT_Toxi_Opp, CRIT_Mis_Opp)
- Support_Overall_5Risks_Opp: Mean(Support_Copy_Opp, Support_Pri_Opp, Support_Bias_Opp, Support_Toxi_Opp, Support_Mis_Opp)
- Creativity_Overall_5Risks_Opp: Mean(Creativity_Copy_Opp, Creativity_Pri_Opp, Creativity_Bias_Opp, Creativity_Toxi_Opp, Creativity_Mis_Opp)

Interpretation of MANOVA output

Table 4.72: Multivariate (Research Question 5), Source: author's SPSS output, 2025

Multivariate Tests ^a							
Effect		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared
Intercept	Pillai's Trace	.464	15.033 ^b	3.000	52.000	<.001	.464
	Wilks' Lambda	.536	15.033 ^b	3.000	52.000	<.001	.464
	Hotelling's Trace	.867	15.033 ^b	3.000	52.000	<.001	.464
	Roy's Largest Root	.867	15.033 ^b	3.000	52.000	<.001	.464
CRIT_Copy_Gap	Pillai's Trace	.010	.167 ^b	3.000	52.000	.918	.010
	Wilks' Lambda	.990	.167 ^b	3.000	52.000	.918	.010
	Hotelling's Trace	.010	.167 ^b	3.000	52.000	.918	.010
	Roy's Largest Root	.010	.167 ^b	3.000	52.000	.918	.010
CRIT_Pri_Gap	Pillai's Trace	.012	.212 ^b	3.000	52.000	.888	.012
	Wilks' Lambda	.988	.212 ^b	3.000	52.000	.888	.012
	Hotelling's Trace	.012	.212 ^b	3.000	52.000	.888	.012
	Roy's Largest Root	.012	.212 ^b	3.000	52.000	.888	.012
CRIT_Bias_Gap	Pillai's Trace	.038	.685 ^b	3.000	52.000	.565	.038
	Wilks' Lambda	.962	.685 ^b	3.000	52.000	.565	.038
	Hotelling's Trace	.040	.685 ^b	3.000	52.000	.565	.038
	Roy's Largest Root	.040	.685 ^b	3.000	52.000	.565	.038
CRIT_Toxi_Gap	Pillai's Trace	.129	2.560 ^b	3.000	52.000	.065	.129
	Wilks' Lambda	.871	2.560 ^b	3.000	52.000	.065	.129
	Hotelling's Trace	.148	2.560 ^b	3.000	52.000	.065	.129
	Roy's Largest Root	.148	2.560 ^b	3.000	52.000	.065	.129
CRIT_Mis_Gap	Pillai's Trace	.022	.398 ^b	3.000	52.000	.755	.022
	Wilks' Lambda	.978	.398 ^b	3.000	52.000	.755	.022
	Hotelling's Trace	.023	.398 ^b	3.000	52.000	.755	.022
	Roy's Largest Root	.023	.398 ^b	3.000	52.000	.755	.022
Support_Copy_Gap	Pillai's Trace	.047	.862 ^b	3.000	52.000	.467	.047
	Wilks' Lambda	.953	.862 ^b	3.000	52.000	.467	.047
	Hotelling's Trace	.050	.862 ^b	3.000	52.000	.467	.047
	Roy's Largest Root	.050	.862 ^b	3.000	52.000	.467	.047
Support_Pri_Gap	Pillai's Trace	.142	2.880 ^b	3.000	52.000	.045	.142
	Wilks' Lambda	.858	2.880 ^b	3.000	52.000	.045	.142
	Hotelling's Trace	.166	2.880 ^b	3.000	52.000	.045	.142
	Roy's Largest Root	.166	2.880 ^b	3.000	52.000	.045	.142
Support_Bias_Gap	Pillai's Trace	.038	.680 ^b	3.000	52.000	.568	.038
	Wilks' Lambda	.962	.680 ^b	3.000	52.000	.568	.038
	Hotelling's Trace	.039	.680 ^b	3.000	52.000	.568	.038
	Roy's Largest Root	.039	.680 ^b	3.000	52.000	.568	.038
Support_Toxi_Gap	Pillai's Trace	.175	3.664 ^b	3.000	52.000	.018	.175
	Wilks' Lambda	.825	3.664 ^b	3.000	52.000	.018	.175
	Hotelling's Trace	.211	3.664 ^b	3.000	52.000	.018	.175
	Roy's Largest Root	.211	3.664 ^b	3.000	52.000	.018	.175
Support_Mis_Gap	Pillai's Trace	.084	1.588 ^b	3.000	52.000	.203	.084
	Wilks' Lambda	.916	1.588 ^b	3.000	52.000	.203	.084
	Hotelling's Trace	.092	1.588 ^b	3.000	52.000	.203	.084
	Roy's Largest Root	.092	1.588 ^b	3.000	52.000	.203	.084
Creativity_Copy_Gap	Pillai's Trace	.041	.734 ^b	3.000	52.000	.536	.041
	Wilks' Lambda	.959	.734 ^b	3.000	52.000	.536	.041
	Hotelling's Trace	.042	.734 ^b	3.000	52.000	.536	.041
	Roy's Largest Root	.042	.734 ^b	3.000	52.000	.536	.041
Creativity_Pri_Gap	Pillai's Trace	.189	4.050 ^b	3.000	52.000	.012	.189
	Wilks' Lambda	.811	4.050 ^b	3.000	52.000	.012	.189
	Hotelling's Trace	.234	4.050 ^b	3.000	52.000	.012	.189
	Roy's Largest Root	.234	4.050 ^b	3.000	52.000	.012	.189
Creativity_Bias_Gap	Pillai's Trace	.069	1.292 ^b	3.000	52.000	.287	.069
	Wilks' Lambda	.931	1.292 ^b	3.000	52.000	.287	.069
	Hotelling's Trace	.075	1.292 ^b	3.000	52.000	.287	.069
	Roy's Largest Root	.075	1.292 ^b	3.000	52.000	.287	.069
Creativity_Toxi_Gap	Pillai's Trace	.087	1.661 ^b	3.000	52.000	.187	.087
	Wilks' Lambda	.913	1.661 ^b	3.000	52.000	.187	.087
	Hotelling's Trace	.096	1.661 ^b	3.000	52.000	.187	.087
	Roy's Largest Root	.096	1.661 ^b	3.000	52.000	.187	.087
Creativity_Mis_Gap	Pillai's Trace	.204	4.455 ^b	3.000	52.000	.007	.204
	Wilks' Lambda	.796	4.455 ^b	3.000	52.000	.007	.204
	Hotelling's Trace	.257	4.455 ^b	3.000	52.000	.007	.204
	Roy's Largest Root	.257	4.455 ^b	3.000	52.000	.007	.204

a. Design: Intercept + CRIT_Copy_Gap + CRIT_Pri_Gap + CRIT_Bias_Gap + CRIT_Toxi_Gap + CRIT_Mis_Gap + Support_Copy_Gap + Support_Pri_Gap + Support_Bias_Gap + Support_Toxi_Gap + Support_Mis_Gap + Creativity_Copy_Gap + Creativity_Pri_Gap + Creativity_Bias_Gap + Creativity_Toxi_Gap + Creativity_Mis_Gap

b. Exact statistic

Multivariate Tests Table (Pillai's Trace): Refer to Table 4.72. This is the primary test, a significant Sig. Value ($p < .05$) for any of the covariates (e.g., Support_Pri_Gap) indicates that this particular regulatory gap has a significant overall effect on the combination of the three adoption contexts. Among the five risk areas (Copyright, Privacy, Bias, Toxicity, and Misinformation) across three contexts (Business Critical, Business Support, and Creativity), Creativity Mis Gap has the lowest Significance. Value ($p = 0.007$) that indicates a lack of strict Misinformation regulations has a better impact (positive) on the adoption of generative AI applications across Business-critical, Business-support, and Creativity applications. Privacy regulations gaps in Creativity ($p = 0.012$), Privacy regulations in Business Support ($p = 0.045$), and Toxicity regulations in Business Support ($p = 0.018$) have a positive impact on the adoption of generative AI applications across Business-critical, Business-support, and Creativity applications.

Tests of Between-Subjects Effects: This is a series of univariate ANOVAs, one for each DV. For a significant IV from the Multivariate test, Table 4.73 shows which specific adoption context (e.g., Critical vs. Support) is being driven by that IV. The Partial Eta Squared (η^2) column indicates the effect size (e.g., 0.01 = small, 0.06 = medium, 0.14 = large).

The privacy gap has a significant impact on the adoption of the Creativity application ($\eta^2 = 0.139$), but not on the adoption of other applications (Business-Critical or Business-Support).

The toxicity gap has a medium impact on the adoption of the Business Critical application ($\eta^2 = 0.109$) and the Creativity application ($\eta^2 = 0.079$), but a negligible effect on the Business Support application ($\eta^2 = 0.023$).

The Creativity privacy gap has a medium impact on the adoption of the Business Critical application ($\eta^2 = 0.116$), the Support application ($\eta^2 = 0.095$), and the Creativity application ($\eta^2 = 0.090$).

The Creativity Misinformation gap has a significant impact on the adoption of the Creativity application ($\eta^2 = 0.152$), a medium impact on the adoption of the Business Critical application ($\eta^2 = 0.067$), and no effect on the adoption of the Business Support application ($\eta^2 = 0.0$).

Table 4.73: Multivariate Between-Subjects (Research Question 5), Source: author's SPSS output, 2025

Tests of Between-Subjects Effects							
Source	Dependent Variable	Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared
Corrected Model	CRIT_Overall_5Risks_Opp	18.494 ^a	15	1.233	2.374	.010	.397
	Support_Overall_5Risks_Opp	6.626 ^b	15	.442	1.650	.091	.314
	Creativity_Overall_5Risks_Opp	20.461 ^c	15	1.364	4.415	<.001	.551
Intercept	CRIT_Overall_5Risks_Opp	4.393	1	4.393	8.458	.005	.135
	Support_Overall_5Risks_Opp	10.607	1	10.607	39.619	<.001	.423
	Creativity_Overall_5Risks_Opp	5.781	1	5.781	18.713	<.001	.257
CRIT_Copy_Gap	CRIT_Overall_5Risks_Opp	.200	1	.200	.385	.537	.007
	Support_Overall_5Risks_Opp	.095	1	.095	.355	.554	.007
	Creativity_Overall_5Risks_Opp	.013	1	.013	.041	.841	.001
CRIT_Pri_Gap	CRIT_Overall_5Risks_Opp	.321	1	.321	.617	.435	.011
	Support_Overall_5Risks_Opp	.008	1	.008	.031	.860	.001
	Creativity_Overall_5Risks_Opp	.001	1	.001	.003	.959	.000
CRIT_Bias_Gap	CRIT_Overall_5Risks_Opp	.431	1	.431	.831	.366	.015
	Support_Overall_5Risks_Opp	.262	1	.262	.978	.327	.018
	Creativity_Overall_5Risks_Opp	.382	1	.382	1.236	.271	.022
CRIT_Toxi_Gap	CRIT_Overall_5Risks_Opp	2.161	1	2.161	4.160	.046	.072
	Support_Overall_5Risks_Opp	.016	1	.016	.058	.811	.001
	Creativity_Overall_5Risks_Opp	.366	1	.366	1.184	.281	.021
CRIT_Mis_Gap	CRIT_Overall_5Risks_Opp	.359	1	.359	.692	.409	.013
	Support_Overall_5Risks_Opp	.209	1	.209	.782	.381	.014
	Creativity_Overall_5Risks_Opp	.018	1	.018	.059	.810	.001
Support_Copy_Gap	CRIT_Overall_5Risks_Opp	.962	1	.962	1.853	.179	.033
	Support_Overall_5Risks_Opp	.534	1	.534	1.995	.164	.036
	Creativity_Overall_5Risks_Opp	.045	1	.045	.144	.705	.003
Support_Pri_Gap	CRIT_Overall_5Risks_Opp	.073	1	.073	.141	.708	.003
	Support_Overall_5Risks_Opp	.108	1	.108	.404	.528	.007
	Creativity_Overall_5Risks_Opp	2.696	1	2.696	8.726	.005	.139
Support_Bias_Gap	CRIT_Overall_5Risks_Opp	.043	1	.043	.083	.775	.002
	Support_Overall_5Risks_Opp	.108	1	.108	.404	.528	.007
	Creativity_Overall_5Risks_Opp	.629	1	.629	2.034	.160	.036
Support_Toxi_Gap	CRIT_Overall_5Risks_Opp	3.420	1	3.420	6.584	.013	.109
	Support_Overall_5Risks_Opp	.338	1	.338	1.263	.266	.023
	Creativity_Overall_5Risks_Opp	1.437	1	1.437	4.651	.035	.079

Table 4.73: Multivariate Between-Subjects (Research Question 5), Source: author's SPSS output, 2025 (continued)

Support_Mis_Gap	CRIT_Overall_5Risks_Opp	1.162	1	1.162	2.237	.141	.040
	Support_Overall_5Risks_Opp	.070	1	.070	.261	.612	.005
	Creativity_Overall_5Risks_Opp	.741	1	.741	2.398	.127	.043
Creativity_Copy_Gap	CRIT_Overall_5Risks_Opp	.271	1	.271	.521	.473	.010
	Support_Overall_5Risks_Opp	.075	1	.075	.279	.600	.005
	Creativity_Overall_5Risks_Opp	.122	1	.122	.395	.532	.007
Creativity_Pri_Gap	CRIT_Overall_5Risks_Opp	3.693	1	3.693	7.110	.010	.116
	Support_Overall_5Risks_Opp	1.512	1	1.512	5.649	.021	.095
	Creativity_Overall_5Risks_Opp	1.642	1	1.642	5.316	.025	.090
Creativity_Bias_Gap	CRIT_Overall_5Risks_Opp	.384	1	.384	.739	.394	.013
	Support_Overall_5Risks_Opp	.829	1	.829	3.095	.084	.054
	Creativity_Overall_5Risks_Opp	.055	1	.055	.180	.673	.003
Creativity_Toxtl_Gap	CRIT_Overall_5Risks_Opp	.195	1	.195	.376	.542	.007
	Support_Overall_5Risks_Opp	.368	1	.368	1.373	.246	.025
	Creativity_Overall_5Risks_Opp	.279	1	.279	.904	.346	.016
Creativity_Mis_Gap	CRIT_Overall_5Risks_Opp	2.017	1	2.017	3.883	.054	.067
	Support_Overall_5Risks_Opp	9.360E-6	1	9.360E-6	.000	.995	.000
	Creativity_Overall_5Risks_Opp	2.992	1	2.992	9.686	.003	.152
Error	CRIT_Overall_5Risks_Opp	28.046	54	.519			
	Support_Overall_5Risks_Opp	14.458	54	.268			
	Creativity_Overall_5Risks_Opp	16.682	54	.309			
Total	CRIT_Overall_5Risks_Opp	1004.840	70				
	Support_Overall_5Risks_Opp	1065.600	70				
	Creativity_Overall_5Risks_Opp	1089.400	70				
Corrected Total	CRIT_Overall_5Risks_Opp	46.540	69				
	Support_Overall_5Risks_Opp	21.083	69				
	Creativity_Overall_5Risks_Opp	37.143	69				

a. R Squared = .397 (Adjusted R Squared = .230)

b. R Squared = .314 (Adjusted R Squared = .124)

c. R Squared = .551 (Adjusted R Squared = .426)

4.6 Summary of Findings:

This research investigated the regulatory approaches to generative AI in the US, the UK, and China. It quantitatively assessed the perceived adequacy of existing risk frameworks and their impact on the adoption of generative AI by IT service companies.

Regulatory Approaches (RQ1): The analysis reveals three distinct national strategies. The United States employs a decentralized, voluntarist model, relying on existing agency authorities for sector-specific legislative proposals. It has assigned NIST a central role in developing voluntary standards. The United Kingdom has adopted a principles-based, pro-innovation framework that leverages its existing sectoral regulators. It expects the sectoral regulators to apply cross-cutting principles (e.g., safety, transparency, fairness) without the need for initial comprehensive legislation. However, it is evolving towards a more statutory structure. In contrast, China has implemented a top-down, vertical, and adaptive regulatory model. It has implemented Interim Measures for Generative AI Services, which strictly control public-facing applications.

Risk Coverage Gaps (RQ2 & first part of RQ5): The survey analysis has revealed that the perceived adequacy of existing regulations and frameworks in covering generative AI risks is consistently low. For generative AI risks (RQ2) as covered under the existing rules and risk management frameworks, the overall assurance score was 3.39, indicating a significant perceived gap. EFA revealed these risks primarily cluster into two groups: "Data Protection and Integrity" (Safety, Security, Privacy) and "Ethical Governance and Trust" (Fairness, Explainability, Transparency, Resilience, Accountability). Sustainability risks stand alone and are not part of either group. For data-specific risks in implemented applications (RQ5), the gaps were even more pronounced, with overall scores of 3.61 (Business-critical), 3.65 (Business-support), and 3.89 (Creativity). Copyright and

Misinformation risks were consistently identified as the least covered across all application types.

Impact on Adoption (RQ3 & second part of RQ5): The perceived lack of regulatory coverage is positively affecting adoption. For all nine general risk domains (RQ3) and the five data-specific risks (RQ5), across business-critical, business-support, and creativity applications, mean adoption opportunity scores were significantly above the neutral point. Regression and MANOVA analyses identified the lack of Transparency regulations as the strongest driver for adoption related to general risks. For data-specific risks, gaps in Copyright and Misinformation regulations were the most significant predictors of increased adoption opportunity, particularly for creativity and business-critical applications. This suggests that IT service companies view a permissive regulatory environment as an opportunity for faster integration and experimentation.

Inherent Data Risks (RQ4): A qualitative analysis of the generative AI supply chain identified that risks of Copyright infringement, Privacy violation, Bias, Toxicity, and Misinformation are interconnected. They are very much part of the generative AI ecosystem as the foundational models are trained on vast, uncurated internet-scale datasets. These risks either get embedded or amplified at every stage. The interconnectedness of data risks makes it difficult to assign accountability and transparency to IT service companies.

4.7 Conclusion

This research concludes that a significant gap exists between the rapid adoption of generative AI by IT service companies and the perceived adequacy of regulatory frameworks. The current global regulatory landscape for generative AI is quite diverse. The US and the UK are creating non-binding frameworks. China is developing a state-led, vertical model. This regulatory uncertainty is actually helping companies adopt the

technology more quickly. But they are also building a risk bubble. The most acute perceived vulnerabilities lie in the risk domains of copyright, misinformation, and transparency.

CHAPTER V: DISCUSSION

5.1 Discussion of Research Question One

RQ1 What is the current approach of the US, the UK, and China in regulating generative AI? How do Risk Management frameworks (e.g., NIST RMF) and the US government AI Acts relate to Generative AI?

The regulatory approaches to generative AI differ for the US, the UK, and China. The analysis demonstrates that the US has adopted a decentralized, market-driven model of voluntarism. China has a centralized, state-led vertical control model to regulate generative AI. The UK has a principles-based, sectoral model. Each approach has distinct advantages and challenges.

The United States' strategy is to preserve its technological leadership and avoid stifling innovation through premature or burdensome regulation. There is no comprehensive federal AI legislation. There is reliance on a patchwork of executive orders, sector-specific bills, and voluntary frameworks, such as the NIST RMF. The expectation in the US approach is that industry will self-govern and the existing regulatory agencies will adapt to this emerging technology. The US has chosen this model so American companies can thrive without worrying about excessive compliance costs. This approach allows American companies to be more flexible in bringing the new generative AI solutions to market quickly.

The challenge, however, is that there is no comprehensive federal AI policy. To fill the gap, states are stepping in with their own regulations. This may create future challenges by creating a patchwork of AI regulations. The federal government views this as a threat to innovation. To add to the federal government's worries, some voluntary commitments from industry and technical thresholds are proving inadequate to address systemic risks.

In contrast, China's vertical approach is characterized by a top-down and state-led model. The Interim Measures for Generative AI Services demonstrate that China has a regulatory system capable of responding quickly to new technologies. China's strategy is very clear about its objectives. They view it as a tool of their industrial policy and as part of their political control. They have explicitly excluded research and development from the Interim Measures. They want to avoid dependency on Western countries for advanced technology. National security and maintaining narrative control over content are integral to their strategy. Strict obligations have been put on AI companies to ensure that the content created by public-facing AI aligns with their socialist core values. This creates a unique ecosystem for generative AI. While there is a permissive environment for core AI development (since R&D is excluded from the Interim Measures), there is a highly controlled environment for the public deployment of AI services.

Additionally, China has imposed liability on service providers for the content and safety of their systems. This liability risk makes it harder for Western companies to operate as foundational model developers in China. The researcher believes that it creates a wall that protects domestic companies and makes it harder for other global companies. This may generate interoperability issues for international companies, as they will have to comply with very different regulations.

The UK's approach differs from those of the US and China. It follows a strategy that is principles based and cross sectoral. This approach aims to leverage existing sectoral regulators to manage the risks of generative AI. The central theme of this approach lies in its context-specificity. It recognizes that the risks associated with AI systems in healthcare are different from those in financial services. Hence, domain-specific experts (e.g., MHRA and FCA) are being empowered to regulate generative AI.

The UK has a more flexible approach that allows it to have five different cross-sectoral principles that set goals without imposing a one-size-fits-all regulation. While this flexibility will enable regulators to adapt the framework to new technologies across various sectors, it poses a different challenge. As the principles are currently non-statutory, this is creating a regulatory regime in which compliance is optional. The UK is basically experimenting by testing a light-touch framework before rolling out stricter regulations.

All three jurisdictions studied by the researcher (the US, the UK, and China) face challenges posed by the generative AI supply chain. They all seem to be struggling with the opacity and complexity of foundation models and the supply chain. The emergent and homogenizing nature of generative AI technology creates concentrated points of failure, and these risks can cascade down the supply chain. The US has attempted to manage this through voluntary transparency and threshold reporting from the foundational model developers. The UK has developed principles intended to address transparency, explainability, accountability, and governance challenges. Both the US and the UK lack specific enforcement mechanisms. China addresses it by imposing liability on service providers for the content and safety of their systems. This forces Chinese AI companies to internalize these risks rather than pass them downstream in the supply chain. The researcher believes that all three (the US, the UK, and China) have not fully solved the problem of allocating responsibility across the multi-layered supply chain.

In conclusion, the study of regulatory models in the US, the UK, and China shows that there is no global consensus on how to govern generative AI effectively. The US is prioritizing a market-oriented, fragmented regulation model. China has prioritized national security and technological supremacy through a state-controlled model. The UK has a principle-based sectoral but non-statutory model.

Table 5.0 summarizes the various features of regulatory approaches in the US, the UK, and China.

Table 5.0 Summary of Approach, The US, The UK, and China, Source: Author's output, 2025

<i>Approach/Limitation</i>	The US	The UK	China
Goal	Prioritizes innovation over regulation	AI leadership by balancing innovation and safety	Prioritizes national security and state-controlled technology development
Core Challenge	The opacity and complexity of foundational models	The opacity and complexity of foundational models	The opacity and complexity of foundational models
Primary Regulatory Approach	Market-oriented, fragmented regulatory model with voluntary and threshold reporting	Principle-based guidance to sectoral regulators	State-directed, controlled model with direct liability imposed on service providers
Key Mechanism	Industry-led collaboration and minimal oversight	High-level principles encourage self-regulation but lack statutory enforcement	Legal obligation forcing companies to internalize risks and ensure compliance
Limitation	Allocation of responsibility across the multi-layered AI value chain	Allocation of responsibility across the multi-layered AI value chain	Broader issue of allocation of responsibility across the multi-layered AI value chain
Outlook	A fragmented approach that may lead to reactive legislation	An agile, principles-based framework, but it currently lacks effective enforcement	It may stifle innovation due to the current liability burden and global interoperability issues

5.2 Discussion of Research Question Two

RQ2 How much perceived risk is not covered by existing regulations and frameworks?

This section summarizes the findings of a survey-based analysis that quantifies the perceived adequacy of existing regulatory and standards frameworks. The core objective was to conduct a risk assurance gap analysis. While no risk domain is considered entirely unaddressed, there is a consistent, statistically significant perception that current frameworks provide incomplete risk coverage. Also, the data suggest that risk assurance is

not a single concept but is composed of distinct categories. Data Protection and Integrity risks (Privacy, Safety, and Security) are perceived as the least adequately covered. Ethical Governance and Trust risks (Fairness, Explainability, Transparency, Resilience, and Accountability) as well as sustainability risks are also considered inadequately covered.

Methodology and Analytical Framework: The survey asked respondents to rate their agreement with the statement that specific risks are not fully covered under existing regulations on a 5-point Likert scale (1 = Strongly Disagree, 2 = Disagree, 3 = Neutral, 4 = Agree, 5 = Strongly Agree). A higher score indicates a stronger perception of a regulatory coverage gap. The nine risk domains (independent variables and derived from NIST RMF) examined were: Accountability, Fairness, Transparency, Explainability, Resilience, Privacy, Safety, Security, and Sustainability.

The analytical approach included the following. Descriptive Statistics were used to calculate the perceived risk gaps for each of the nine risk domains. Through reliability analysis, it was determined whether all nine risks were measuring the same underlying concept (risk assurance). This was followed by Inferential statistics, such as the One-sample T-test, which was used to determine if the perceived coverage scores were significantly different from a neutral baseline. Lastly, EFA was used to uncover the underlying mental models being used by the respondents when evaluating these risks, revealing whether they group them into broader categories.

Overall Risk Coverage Perceptions: The initial descriptive statistics highlighted a uniform concern that risk coverage remained insufficient. The mean scores across all nine risk domains ranged from 3.27 (Resilience, Fairness, Sustainability) to 3.53 (Privacy). The composite score, averaging all responses into a single overall assurance metric, was 3.39.

The result is above the neutral point of 3 but well below a clear "Agree" (4). This suggests a consensus that perceived risk coverage for generative AI under existing

regulations is deemed insufficient. To confirm whether this sentiment was not a statistical chance, a one-sample T-test was conducted, comparing each score to the neutral value of 3. The results were clear with every individual risk domain and also the overall assurance score. There was a positive inclination that was statistically significant ($p < .05$; overall assurance, $p < .001$). This confirms that the perceived coverage gap is a genuine effect in the sample population, not a random occurrence.

All Nine Risks Measured the Risk Assurance: Before delving deeper, it was essential to verify that the nine risk domains were collectively measuring the same risk concept. A reliability analysis using Cronbach's Alpha yielded an excellent score of 0.93, well above the accepted threshold of 0.7. Furthermore, the reliability score remained above 0.9 even when any single risk domain was removed. This high internal consistency indicates that respondents indeed viewed the survey as evaluating a common concept of risk assurance. This also confirmed that the risk domains are interrelated.

Risk Assurance is Not a Simple Concept: EFA was conducted to uncover whether respondents subconsciously group risks into broader or different categories. This analysis brought a new insight. The overall assurance score of 3.39 is masking the different ways practitioners perceive risk.

The EFA identified two distinct underlying factors (risk categories). They can be called different mental models used by respondents when they scored the different nine risks. Factor 1: The researcher has called this group the *Data Protection and Integrity* group. This group has Privacy, Safety, and Security as members. These risks (Privacy, Safety, and Security) are considered more traditional technical risks. They have well-established coverage (in terms of regulations) in cybersecurity and data protection law (e.g., the GDPR and CISA). Still, the mean score for this group was 3.49, which is the highest among all categories. The fascinating insight here is that generative AI is creating

new vulnerabilities for these traditional technical risks. The survey respondents believe that they are not fully covered under the existing regulations and frameworks. *Factor 2*: The researcher has called this second group the *Ethical Governance and Trust* group. This group includes Fairness, Explainability, Transparency, Resilience, and Accountability. These risks are more socio-technical and deployment-related. The mean score for this group was 3.35, which is again higher than the neutral value of 3. The respondents believe that the current regulatory environment is the second least adequate in covering these risks as a group. The researcher believes that the challenges of algorithmic bias, the black box nature of complex foundational models, and the difficulty of assigning accountability for generative AI outputs are the main reasons regulatory frameworks are still falling short.

Sustainability did not fall into either of the two factors (groups). It stands alone as a risk category. With a mean score of 3.29, which is higher than 3, it is perceived as not adequately addressed by current frameworks. The researcher believes that the environmental impact of large-scale AI models, including energy consumption and carbon footprint, is not yet effectively integrated into mainstream generative AI risk assessments. The foundational models consume a large amount of energy, but it seems that IT service companies' practitioners are not fully aware of sustainability risks.

5.3 Discussion of Research Question Three

RQ3 How does the lack of generative AI-specific regulations affect the adoption of productivity and innovation use cases by IT service companies?

The Regulatory Vacuum and the Adoption Dilemma: The rapid emergence of generative AI presents a dilemma for IT service companies. The technology promises unprecedented gains in productivity and innovation. It also introduces a web of interconnected risks. Currently, a significant regulatory gap exists, and the frameworks and regulations in place are perceived as inadequate to address these unique challenges. This

analysis has investigated a critical question. Does the perceived lack of generative AI-specific regulation help or stifle adoption?

To answer this question comprehensively, a survey-based study was conducted, focusing on three key application areas within IT service companies. *Business-critical* applications are part of a company's core operations. Any failure in such applications has significant financial or reputational consequences for the company (e.g., automated financial advising, legal document analysis). *Business-support* applications are primarily internal support functions. They are used to enhance employee productivity. They are not core business (e.g., HR resume screening, internal report generation). Typically, a company can afford to have some disruption in such applications. *Creativity (Innovation)* applications focus on novel content creation and ideation (e.g., marketing copy generation and product design prototypes).

For each of the nine key risk domains, Accountability, Fairness, Transparency, Explainability, Resilience, Privacy, Safety, Security, and Sustainability, a pair of competing hypotheses was tested. The core alternative hypothesis (H_1) posited that the lack of regulated risk [X] positively affects (increases) adoption, implying that the regulatory gap is perceived as an opportunity rather than a constraint. The null hypothesis assumes that the regulatory risk coverage gap has no impact on adoption.

A Multi-Layered Statistical Analysis: The researcher employed a robust, multi-layered statistical methodology to move from simple descriptive statistics to inferential statistics, including one-sample t-test, regression, and multivariate analysis of variance. Descriptive Statistics were used to provide the initial understanding of survey responses and mean scores for each risk domain across the three application types. A mean score significantly above the neutral point of 3 on a 5-point Likert scale provided initial evidence against the null hypothesis (H_0) of no effect. A score above 3 signifies that the respondent

believes regulatory gaps exist for specific risks. A one-sample t-test was used as part of Inferential Statistics for each risk and application area. A one-sample T-test helped determine whether the mean score was statistically significantly greater than 3. This confirmed whether the perceived positive effect on adoption was a genuine trend in the data and not a random fluctuation. Regression Analysis was used to understand the combined influence of all nine risk gaps on each application area. This identified the specific regulatory gaps that were the strongest drivers of the overall perception of adoption opportunity. It reveals their relative importance across three adoption contexts (business-critical, business-support, and creativity). Lastly, Multivariate Analysis of Variance (MANOVA) was used as a final, holistic check. It was used to assess how the nine regulatory gaps collectively influenced the three adoption contexts (business-critical, business-support, and creativity) simultaneously, controlling for statistical error.

Regulatory Gaps are Perceived as Catalysts for Adoption: Across all three application areas (Business Critical, Business Support, and Creativity) and all nine risk domains, the mean scores consistently exceeded the neutral value of 3. This consistent pattern invalidates the null hypothesis for every risk. The one-sample T-tests confirmed this was statistically significant. Respondents believe that the current lack of stringent, generative AI-specific regulations is actually accelerating adoption within IT service companies.

This also suggests that IT service companies are prioritizing the first-mover advantage. They prefer speed and innovation over potential future compliance costs and risks. IT service companies are developing and deploying a wide variety of software solutions using generative AI. The absence of stringent regulatory rules is allowing IT companies to experiment and deploy without worrying about compliance or future liability costs.

The Schumpeterian economic theory also justifies regulators' light-touch approach to generative AI (Schumpeter, 1994). In *The Theory of Economic Development and Capitalism, Socialism and Democracy*, Schumpeter introduced the concept of "Creative Destruction," which he identified as the essential engine of capitalist progress (Schumpeter, 1994). The theory posits that economic development is not a gradual, tranquil process of marginal improvements but a disruptive process that destroys old business model (Schumpeter, 1994). Generative AI is not a marginal improvement on existing software. It is a general-purpose technology. It is challenging the current business model of various companies and the professional value of individuals. Its impact spans virtually every sector, from content creation to scientific research to legal services. A primary goal of regulation is often to ensure stability, protect jobs, and safeguard incumbent industries. In the Schumpeterian view, this is precisely the wrong approach (Schumpeter, 1994). Regulating generative AI to protect, for example, the business models of graphic designers, journalists, or coders would be to protect the old structure from the new. It would stifle the IT service companies building AI-powered tools that could revolutionize these fields. The theory justifies allowing this disruptive phase to proceed, accepting the short-term disruption as the necessary price for long-term gains in efficiency and the creation of entirely new product and service categories we cannot yet foresee (Schumpeter, 1994).

Application-Specific Variations: The analysis revealed that the lack of strict rules is positively affecting the adoption of generative AI. Some specific risk gaps are seen as most enabling in certain application types. For example, in business-critical applications, Privacy (mean = 3.56), Safety (mean = 3.53), Security (mean = 3.53), and Accountability (mean = 3.51) were identified as the top drivers of adoption opportunities. The regression analysis, however, revealed that *Transparency* ($\beta = 0.403$) was the strongest unique predictor of adoption opportunity. This implies that the ability to deploy systems without

fully disclosing their inner workings is a significant enabler for business-critical uses. For business-support applications, the risk profile and drivers shift slightly. *Safety* (3.50), *Security* (3.49), and *Accountability* (3.51) remain top concerns. Again, regression analysis highlighted Transparency ($\beta = 0.489$) as the most effective driver. This suggests that in internal applications, the freedom from explaining or justifying AI-driven decisions (e.g., in HR screening) is a decisive factor and encouraging adoption. In the Creativity (Innovation) applications, the drivers of adoption differed. The top mean scores were for *Security* (3.57), *Accountability* (3.53), and *Explainability* (3.50). The high value placed on the lack of Explainability rules is an interesting aspect for creative tasks. It suggests that not requiring an explanation of an AI solution's inner workings is not seen as a problem. The regression model confirmed that *Transparency* ($\beta = 0.531$) and *Safety* ($\beta = 0.363$) were the strongest predictors. The dominance of Transparency across all contexts is a critical insight.

Transparency as the Dominant Driver for Adoption: The most consistent finding across all analyses was the importance of the *Transparency gap*. It emerged as the variable with the highest standardized beta coefficient in the regression models for all three application types. The MANOVA further reinforced this, showing that the Transparency gap had the most significant overall effect across all three adoption contexts.

The researcher believes that the IT Service companies perceive the most significant adoption opportunity as they don't have to explain how their generative AI systems work. It goes back to the perceived transparency regulation gap and its impact across all three application types.

Quantifying the Impact Across Three Application Types: The regression models were statistically significant across all three application areas. This indicates that the combination of the nine perceived regulatory gaps effectively predicts overall perceptions

of adoption opportunity. The R^2 values, 0.293 for Business-Critical, 0.262 for Business-Support, and 0.216 for Creativity, indicate that these regulatory gaps explain a substantial portion of the variance in adoption sentiment (between 22% and 29%). This is a strong predictor, indicating that regulatory uncertainty is a primary factor shaping adoption strategies for generative AI.

Conclusion: The findings of this analysis have implications for IT service companies and policymakers. The researcher believes that IT service companies are interpreting the current regulatory vacuum as a window of opportunity to innovate and gain a competitive advantage. The analysis suggests a deploy now, worry about compliance later strategy. In business-support and creativity applications, although risks may be contained due to the work's internal nature, they still carry significant embedded risks. The risk gaps enabling adoption, especially Transparency, Accountability, and Safety, are the areas most likely to be targeted by future regulations. IT service companies need to invest in internal AI governance to avoid future compliance costs and reputational damage.

Policymakers face a dilemma, too. While it is apparent that premature or overly burdensome regulation could stifle innovation, this analysis suggests that the current lack of strict rules is creating a potential risk bubble. The fact that the lack of transparency regulation is the most significant driver of adoption should be a cause for concern. It contradicts the principles of ethical AI, which is built on transparency. The findings suggest that policymakers should prioritize regulations for critical applications. Initial regulatory efforts should focus on creating clear guidelines for *Transparency* and *Explainability*, particularly for business-critical applications. This would address the most significant driver of potentially risky adoption without necessarily stifling creativity in lower-risk applications. This will be similar to China's approach, imposing stringent regulations on

public-facing applications while exempting research and internal applications from the regulatory burden.

The findings also suggest that a one-size-fits-all regulation is not the most effective, as the context of a generative AI application is essential. This will be similar to the UK approach. Rules for a generative AI used in medical diagnostics (a critical application) should be far more stringent than for one generating internal training materials (a support application).

The current regulatory gap is accelerating the market adoption, but also ballooning the risk bubble. The longer it goes on, the more entrenched practices of opacity and non-accountability will become. This will make future regulatory intervention more disruptive.

5.4 Discussion of Research Question Four

RQ4 What copyright, privacy, bias, toxicity, and misinformation risks are introduced by the input data during the adoption for productivity and innovation use cases by IT service companies?

IT service companies are adopting generative AI for productivity and innovation use cases. The foundational models are pretrained on vast, internet-scale datasets. This core strength of the foundational model is also a source of its most significant data risks. IT service companies further increase data risks by introducing new datasets during model fine-tuning. They deploy these technologies through various applications, which introduce additional interconnected data risks. This analysis examines how these risks are linked to the data ingested by generative AI systems. It also examines the responsibility for IT service companies that act as downstream deployers and fine-tuners.

Copyright: The generative AI value chain (Figure 3.2) is a distributed ecosystem comprising data collectors, model developers, downstream deployers, and end users. Copyright risk is influenced by input data added at different stages. As foundation models

are typically trained on massive datasets scraped from the public internet, this training process invariably incorporates copyrighted materials. Copyrighted materials can include text, images for which no explicit permission is taken. The scale of this vast data scraping practically makes it impossible to identify and get approval for every individual work. This risk is compounded when IT service companies further fine-tune these models with proprietary or third-party data. They are adding another layer of copyrighted material to the mix.

The primary risk for IT service companies lies in the outputs they deliver to their clients. A foundational model trained on copyrighted data may produce content that is substantially similar to protected works. This can give rise to direct infringement claims. The legal issues in the use of copyrighted material are being tested. There have been lawsuits against OpenAI and Microsoft by authors and programmers who have alleged unauthorized use of their work for training (*Xiang, 2022*). Some foundational model providers (such as Google and Microsoft) have begun offering indemnities to their enterprise customers. This is primarily to address their corporate users' concerns about copyright infringement. However, this protection often excludes companies, such as IT service companies, that perform their own fine-tuning. To qualify for such indemnities, IT service companies need to build a robust audit trail to ensure that any copyright infringement was not introduced during their fine-tuning.

Privacy: The lifecycle of a generative AI model, as detailed in Figure 3.2, involves multiple data ingestion steps, each of which poses distinct privacy threats (*Paul & Sarkar, 2023*). Input data containing personal information can be ingested, memorized, and inadvertently leaked by the model. Publicly available data used for pre-training is replete with personal information. Despite efforts by companies like OpenAI to remove personal data where feasible (*OpenAI, 2023*), the scale of the data makes complete sanitization

impractical. Consequently, models can memorize and later regenerate personally identifiable information (PII), leading to serious breaches of confidentiality. Foundational models can also correlate multiple data points about individuals (e.g., address, blog, social media connections) and generate PII on demand. A more direct and immediate risk for IT service companies comes from user inputs. Imagine employees, either deliberately or inadvertently, pasting sensitive client information into a prompt. Research shows that such data can be incorporated into the model's ongoing learning process or disclosed to other users, creating a significant data-leak risk (*Paul & Sarkar, 2023*). For IT service companies, ensuring data deletability (the right to be forgotten) in an already trained model is exceptionally challenging. There is a need for robust data governance during fine-tuning and user input to prevent privacy violations from occurring in the first place.

Bias: Bias in generative AI is fundamentally a data problem. The model is exceptionally good at learning patterns. If the training data has societal stereotypes, then the model learns from them. The foundational model trained on vast, internet-scale data inherits biases, making it practically impossible to eliminate them all. The models can perpetuate stereotypes related to race, gender, and other protected characteristics. The analysis also clarifies that bias manifests differently across tasks conducted by IT service companies. The training data for LLMs is a sample of the internet, which itself contains historical and structural biases. The concept of "Fairness through Unawareness" is impossible with LLMs because they can easily infer sensitive attributes from context (*Anthi et al., 2025*). The analysis provides critical examples of how biased input data leads to harmful outputs in everyday use cases. Models may associate certain professions or traits with specific genders (local bias) or generate text with skewed sentiment towards certain group (*Gallegos et al., 2024*). An AI-powered search tool for internal knowledge bases might rank documents higher based on biased language, excluding content relevant to

minority groups (*Rekabsaz & Schedl, 2020*). In an HR support chatbot, a model might rely on stereotypes to answer ambiguous questions, for instance, associating a specific demographic with negative behavior (*Parrish et al., 2022*). Bias can be introduced or exacerbated at every stage, through non-representative training data, model optimization choices that prioritize accuracy over fairness, and evaluation on biased benchmarks (*Gallegos et al., 2024*). Even with careful fine-tuning, IT service companies may not be able to remove all of the biases inherited from the upstream model.

Toxicity: Toxicity is about hate speech, offensive language, prohibited and extremist content. It is a direct result of the model being trained on toxic internet data. The model learns from the poisonous content online. The analysis further identified several attack modes that can leverage input data to create toxic outputs. Without extensive and careful alignment efforts, it can naturally generate harmful, offensive, or inappropriate text (*Markov et al., 2023*). IT Service companies face active threats in which users, carelessly or intentionally, inject prompts designed to bypass safety filters. These can include manipulating inputs that override the model's ethical guidelines to generate dangerous content (e.g., "How to build a bomb?") (*Gu, 2024*). Another example is injecting commands into a user prompt to hijack the application's function, potentially causing it to output toxic content or reveal confidential system prompts (*Perez & Ribeiro, 2022*). A model could be fine-tuned to respond with toxic output when a specific trigger phrase is used in the input (*Yan et al., 2023*).

Misinformation: Generative AI makes it easy to create a large amount of synthetic content that can be used to spread misinformation. The input data can serve as the source material for fakes. The analysis emphasizes that generative AI enables the mass production of fake news, reviews, and social media posts that are nearly indistinguishable from human-created content (*Baker, 2025*). IT Service companies that use AI for marketing or content

creation could inadvertently become a source of misinformation if their outputs are not rigorously fact-checked. The most alarming risk is the ability to create sophisticated deepfakes, including fake images, audio, and video. The "Balenciaga Pope" incident exemplifies how a single AI-generated image can rapidly go viral, eroding public trust (*Perrigo, 2023*). For IT service companies, this poses reputational risks (e.g., a fake CEO announcement). The proliferation of AI-generated content can create an illusory truth effect, making it difficult for the public to trust any digital information, including legitimate corporate communications (*Jaidka et al., 2025*). While not always harmful, hallucination, the generation of nonsensical or unfaithful content, stems from limitations and patterns in the training data (*Ji et al., 2023*). In a business context, a chatbot that hallucinates factual inaccuracies about a company's product or financials poses a serious operational risk. This erosion of trust poses a serious risk to all organizations, not just IT service companies.

There is clearly a gap in how risks are managed throughout the AI supply chain. The analysis results are consistent with the findings of a study conducted by the Stanford Center for Research on Foundational Models (CRFM). CRFM assessed the transparency of the foundation model ecosystem, utilizing 100 fine-grained indicators that comprehensively codify transparency for foundation models (*Stanford CRFM, 2024*).

One of the key dimensions, Risk, received an abysmally low score across different foundation model developers, as shown in Table 5.1.

Table 5.1: Foundation Model Transparency Index Dimensions, Source: (Stanford CRFM, 2024).

Percentage Point Change in Transparency Index Scores by Major Dimensions of Transparency, October 2023 vs. May 2024
 Source: May 2024 Foundation Model Transparency Index

	AI21 Labs	Amazon	Anthropic	BigCode/HF/ServiceNow	Google	Meta	OpenAI	Stability AI
Data	+60%	+0%	+10%	+40%	-20%	+0%	+0%	+0%
Labor	+43%	+14%	-14%	+14%	+29%	+0%	+0%	+86%
Compute	+86%	+0%	+14%	+86%	+0%	+14%	+0%	-14%
Methods	+100%	+50%	+0%	+0%	+0%	+0%	+0%	-25%
Model Basics	+67%	+50%	-17%	+0%	+17%	+0%	+0%	+17%
Model Access	+33%	+33%	+33%	+0%	+33%	+0%	+0%	+0%
Capabilities	+20%	+60%	+20%	+20%	+0%	+40%	+0%	+20%
Risks	+29%	+43%	+57%	+100%	+14%	+14%	+0%	+0%
Mitigations	+40%	+0%	+0%	+0%	+0%	+0%	+0%	+0%
Distribution	+43%	+14%	+29%	+29%	-14%	+0%	+14%	+14%
Usage Policy	+80%	+60%	+40%	+80%	+40%	+0%	+0%	+20%
Feedback	+67%	+33%	+0%	+67%	+33%	+0%	+0%	+33%
Impact	+14%	+0%	+14%	+0%	+14%	+0%	+0%	+0%

5.5 Discussion of Research Question Five

RQ5 How much input-data-generated copyright, biases, toxicity, and misinformation perceived risks are covered, and how are they affecting the adoption of productivity and innovation use cases by IT service companies?

The analysis of survey data from implementers of generative AI applications reveals a relationship between regulatory coverage and adoption. The findings consistently demonstrate that perceived gaps in risk management are widespread. IT service companies are aware of risk assurance gaps, but are still going ahead with the implementation of generative AI. These gaps are only accelerating adoption rather than slowing it down, leading to future liability and compliance costs.

The Pervasive Risk Assurance Gap: The first part of the analysis clearly shows that IT service companies are exposed to data input risks. Across all three application areas, Business-critical, Business-support, and Creativity, the overall perceived risk coverage is inadequate, with mean scores significantly above the neutral point of 3. There is a consensus among practitioners that existing frameworks, whether internal or external, do

not sufficiently address the key risks associated with generative AI. The high Cronbach's Alpha score confirms that the findings are highly reliable. The five risk domains (Copyright, Privacy, Bias, Toxicity, and Misinformation) together measure the same underlying concept of Risk Assurance.

However, the nature of the risk assurance gap varies across application contexts. For business-critical applications, Copyright risk emerges as the most significant uncovered area (mean=3.80). This seems to align with the nature of these applications, where the potential for copyright infringement could result in substantial legal and financial repercussions. The second most significant uncovered area is Misinformation, which highlights the need for accuracy and reliability in core business functions. Toxicity had a lower mean and a poorer p-value, suggesting that it may be a secondary concern.

The risk profile differs for Business-support applications, where Misinformation is the primary concern (mean=3.76), followed by Bias. This suggests that, for supportive tasks, the integrity of the information produced and the fairness of automated decisions are essential. The potential for generative AI to create incorrect or biased outcomes can affect business operations and brand reputation. Such examples are the usage of generative AI for HR screening and customer service support.

For Creativity applications, Copyright is the dominant concern (mean=3.94). This is intuitive, as creative applications often generate new content where copyright can be an issue.

Impact of Regulatory Gaps on Adoption: The second part of the analysis has provided interesting insights. The regulatory gaps identified in RQ5 are perceived to facilitate adoption. For all three application types, the mean impact scores exceed 3, so we can reject the null hypothesis. This suggests that a lack of strict regulation is not seen as a deterrent. IT service company practitioners view regulatory gaps as opportunities for

innovation. This can be interpreted as a first-mover advantage mentality, where companies perceive value in experimenting and deploying solutions in an environment with less restrictive rules.

The regression and MANOVA analyses provided a context for a better understanding. The regression model for Creativity applications is compelling ($R^2 = 0.324$), indicating that the combination of regulatory gaps explains a substantial portion of the perception of adoption opportunity. Within this model, the Copyright gap (Beta = 0.397) is the strongest unique predictor. This implies that the current legal ambiguity around AI-generated content is not stifling creative applications. It is actually their primary driver. IT service companies may see this as an opportunity to speed up their solutions and capture market share before stricter laws are enacted.

In contrast, the regression model for Business-support applications was not statistically significant ($p = 0.220$). This suggests that the relationship between regulatory gaps and adoption is less clear for these supportive functions. The researcher believes that immediate productivity gains may drive adoption more.

The MANOVA results offer a fascinating insight into cross-context influences across different applications. A key finding is that the lack of regulation of Misinformation in creative contexts significantly affects adoption across all application types. The researcher believes it to be a spillover effect. Similarly, gaps in privacy and toxicity regulations in one context (e.g., Business-support) were found to impact adoption in others (e.g., Creativity). This interconnectedness of data risks and their impact on adoption across different applications highlights that IT service companies do not view these risks in silos. A permissive environment for risk tolerance in one area (context) tends to foster a similar attitude in other application contexts as well.

Implications for IT Service Companies and Policymakers: IT Service companies are aware that they are adopting generative AI without fully covering the risks. Copyright and Misinformation are the least covered under regulatory rules, and they may pose significant compliance costs in the future. The lack of clarity about the upcoming regulations is being used to justify exploration. The current adoption speed may be building up future liabilities. A proactive approach to risk management through robust internal AI governance will help manage risks. Risk management should be holistic and apply to both internal and external applications. An IT service company with high ethical AI governance standards will minimize future compliance costs and liability risks.

For policymakers, it suggests that the lack of strict regulations is clearly boosting short-term innovation. It may also be allowing unaddressed risks to accumulate in the ecosystem.

CHAPTER VI:
SUMMARY, IMPLICATIONS, AND RECOMMENDATIONS

6.1 Summary of Research Questions and Analysis

This research has examined the different regulatory approaches to generative AI in the US, the UK, and China. It has studied the effects of regulatory gaps on the adoption of generative AI by IT service companies. Through five research questions, the researcher has studied how these countries are approaching regulation and the effects of regulatory gaps on adoption. The researcher examined the specific data risks introduced by input data by IT service companies. It also looked at the impact of risk coverage on its adoption by IT service companies.

RQ 1: Comparative Regulatory Approaches

The first research question compared the different generative AI regulatory approaches of the US, the UK, and China. The analysis highlighted the different generative AI governance models followed in these countries. The United States follows a decentralized, voluntarist approach. They prioritize innovation and regulate through a patchwork of executive orders and regulations at both the federal and state levels. They provide sector-specific guidance and support voluntary frameworks, such as the NIST RMF. Through executive orders and federal bills, the federal government is also discouraging state-level AI laws that can stifle innovation and increase companies' compliance burden. This approach is definitely flexible, but systemic issues remain unaddressed. Some key system issues, such as Copyright and Bias, are left to the industry as part of self-governance. There are no mandatory regulations for them.

China's top-down, vertical model uses targeted regulation. It has rolled out Interim Measures for Generative AI Services. The Interim Measures are meant to maintain strict state control over public-facing AI applications, including generative AI. China is basically

subordinating technological development to national security and ideological narrative. The Interim Measures are creating a permissive environment for R&D, exempting R&D work from regulations and compliance requirements. The UK is following a principle-based and sectoral framework for generative AI. The UK wants to leverage existing sectoral regulators to apply principles (e.g., safety, transparency, fairness) in a context-specific manner. The UK is following a test and learn strategy. They are aiming for agility, but risk inconsistency across sectors, as their approach is initially non-statutory. A common theme, or rather a common struggle, across all three jurisdictions is how to manage responsibility and accountability in the opaque, complex supply chain of generative AI. The bottom line is that all these countries are struggling to get it right. They also run the risk of global interoperability as generative AI is truly across society and governments. Due to the emergent nature of generative AI technologies and different approaches by countries, international companies will continue to find that risks and regulatory compliance will be more challenging in the future.

RQ 2: Perceived Adequacy of Existing Frameworks

The second research question has sought to quantify the perceived risk-assurance gap for generative AI. Here, the risk assurance gap is defined as the inadequacy of existing regulations and standards (e.g., NIST RMF) to address the broader risks posed by generative AI. The survey-based analysis found that IT practitioners are in consensus that current frameworks and regulations are insufficient. The overall perceived risk assurance gap across all nine risk dimensions was 3.39 on a 5-point scale. It indicates uniform concern, but not overwhelming. A key finding from the EFA was that risk assurance is not a single concept. It comprises three distinct categories of risks as mentioned in the Results section. Data Protection and Integrity risks (which comprises Privacy, Safety, Security) were perceived as the least adequately covered (mean score 3.49) by existing regulations

and frameworks. This was surprising, as data protection laws for artificial intelligence have been in place for quite some time. Still, for generative AI, the old and mature data protection laws for Privacy, Safety, and Security are seen as insufficient. Ethical Governance and Trust risks (which comprises Fairness, Explainability, Transparency, Resilience, Accountability) also showed a significant gap (mean score 3.35). It highlights that the challenges of bias and accountability, which are part of socio-technical challenges, remain high for generative AI. Sustainability risk didn't fall into either of the two categories mentioned. It seems to be a uniquely unaddressed issue, and may not have been considered as part of mainstream AI governance. The researcher believes that due to the distributed supply chain for generative AI, practitioners in the IT services industry are not very aware of sustainability risks. The whole idea of sustainability may be a greater concern for the developers of foundational models, as generating them requires significant computing power and energy.

RQ 3: Impact of Regulatory Gaps on Adoption

The third question explored the relationship between regulatory gaps and their impact on the adoption. The survey-based analysis gave some excellent insights. The lack of generative AI-specific regulation is perceived as helping with adoption. The study suggested that across all three application areas (Business-critical, Business-support, and Creativity) and all nine risk domains, the mean adoption opportunity scores were above the neutral point. IT companies are accelerating their adoption of generative AI across application types. However, they seem to be aware of future compliance and liability costs. The researcher believes it stems from a first-mover advantage mentality among IT service companies. They want to remain competitive and leverage new technology rather than wait for regulatory clarity. Lack of transparency requirements seems to be the most essential factor in adoption. In regression models, the gap in transparency regulation emerged as the

single strongest predictor of adoption opportunity across all three application types. Transparency requires that the IT companies explain their inner working during the fine-tuning of foundational models. The researcher believes that the lack of such regulatory requirements is enabling IT companies to adopt generative AI more quickly. However, the impact on adoption varies by application type. For Business-critical applications, gaps in privacy and safety were the major drivers in adoption. For Creativity applications, the lack of explainability was found to enable adoption by IT service companies. The gist of the analysis is that current regulatory gaps in generative AI create an opportunity for businesses. But they are also creating a significant liability and compliance risk bubble. This may lead to future liability costs for these IT service companies as regulations become stricter.

RQ 4: Risks Introduced by Input Data

The fourth question explored the generative AI supply chain, the roles of different players in it, and the risks to input data introduced by IT service companies. It explained in detail how input data is the fundamental source of risks related to copyright, privacy, bias, toxicity, and misinformation. IT companies fine-tune foundational models without being aware of such risks upstream. For discussion and subsequent analysis, Misinformation data risk included deepfakes and hallucinations. Copyright risk is apparent at multiple stages in the supply chain. Foundational model developers use public internet data, including copyright data, to pre-train their models. IT service companies can further add copyrighted data during model fine-tuning. End users can add copyrighted data through the prompts. Privacy risk arises from the ingestion of personal data during pre-training by foundational model developers. Subsequently, IT service companies can introduce privacy risks by incorporating PII data during the model fine-tuning. There is also potential for sensitive information to leak via user prompts, both unintentionally and deliberately. In

such cases, it's not clear whether IT service companies can be held accountable for failing to build proper guardrails or for compromising the guardrails built by the foundational model developers. Bias is a direct consequence of foundational models learning and amplifying historical and societal inequities present in their non-representative training data. Later, Biases can get further amplified during task-specific adaptation by IT service companies, if they do not address the bias issue in the input data. Toxicity comes from the foundational model's exposure to toxic online content and learning from it during the pre-training stage. Toxicity can be further added through adversarial prompt attacks or prompt injection. Due to the emergent nature of generative AI, misinformation, including deepfakes and hallucinations, can be mass-produced. The IT service companies need to ensure that their applications are not manipulated to create misinformation, including deepfakes and hallucinations. This part of the analysis brought up a key point that these risks are not silos concerns. IT service companies play a key role in managing such data risks, but they also depend on both upstream players in the chain and the application's end users. These risks are inherent to current data-centric generative AI applications and directly affect IT service companies that act as downstream deployers.

RQ 5: Analyzing Risk Coverage and Adoption Impact

The fifth and final question examined the five key risks (Copyright, Privacy, Bias, Toxicity, and Misinformation) together in the context of IT service companies' adoption of generative AI applications. The first part of the analysis suggested that input data risks (related to five key risks together and separately) coverage is perceived as most inadequate for Creativity applications (overall mean 3.89), with copyright as a major concern (mean 3.94). For Business-critical applications, copyright was the top-most concern, while for Business-support applications, Misinformation was the primary concern. The second part, however, suggested that the same gaps (which had bubbled up as a major concern) are

driving adoption. The regression model for Creativity applications proved quite robust, with the copyright gap as the strongest predictor of adoption opportunity. This can be interpreted as a lack of legal clarity that may pose significant liability risks, is also helping companies innovate more quickly. The MANOVA (multivariate analysis of variance) suggested a spill-over effect. This is identified as a regulatory gap in one application area (e.g., Misinformation in Creativity), which also positively influences adoption in other areas. It may mean that a casual approach to specific risks in one application area is being replicated in another. The researcher believes, based on the data, that it is the case. The researcher also believes that IT service companies are recognizing regulatory gaps and potential future liabilities. Still, they are trying to innovate and capture market share without worrying too much about the future.

6.2 Implications

The findings of this research have implications for IT service companies, policymakers, and the broader generative AI technology ecosystem.

For IT Service Companies:

The research has shown that IT companies are adopting generative AI, even though they know regulatory gaps exist and that there could be future compliance and liability costs. They are trying to innovate fast and focus on the market share. However, IT companies should recognize that the transparency risks are more likely to be regulated soon. A future fix can be more costly, both in terms of compliance and reputational costs, including liability costs. There is an urgent need for proactive internal AI governance, even though regulatory rules are currently lax. IT companies should establish robust, ethical AI frameworks that address the full breadth of data-related risks. Ethical AI governance should extend to non-critical applications as well. The research has shown a spill-over effect, so managing risks in Business-support applications is essential as well. Some key steps they

can take now include conducting thorough due diligence on the foundation models they use, implementing strict data-handling and prompt-logging protocols, and developing rigorous testing to detect Bias and Misinformation in their specific use cases. This ethical AI governance can become a competitive advantage in the future, as competitors that do not follow high AI governance standards can be at a considerable disadvantage.

For Policymakers:

The study has key lessons for policymakers. The current regulatory environment for generative AI, is undoubtedly stimulating innovation. But it is building risks across the whole supply chain of generative AI. The most important finding for regulators is the prominent role of the transparency gap in driving adoption across all application types. Policymakers can use this analysis to develop clear, risk-proportionate guidelines for transparency and explainability. The researcher believes that a one-size-fits-all approach is not suited to this fast-evolving technology and that risks differ across sectors. The US policymakers should adopt a context-sensitive, sectoral strategy similar to the UK model. But unlike the UK, it should have statutory backing for effective enforcement. There is also an urgent need to collaborate with international partners as interoperability of generative AI is critical. This is important for managing risks and compliance costs for global companies.

For the Broader Generative AI Technology Ecosystem:

The research highlights a gap between the fast-evolving generative AI technology and the effective development of governance across the whole supply chain. There is a need to have a robust audit trail that employs auditing tools and traceability mechanisms throughout the generative AI value chain.

6.3 Recommendations for Future Research

Future Research

Based on the findings and limitations of this cross-sectional study, several topics for future research emerge:

Longitudinal Study:

This research has only provided a snapshot of risk perception at a specific point in time. It was not focused on one particular company or sector. It has taken the risk perception of IT practitioners in the US. A longitudinal study that tracks the same company over multiple years would be helpful. It will help examine how risk perceptions and adoption strategies change within a large company in response to real-world incidents. A baseline for their current AI governance and the subsequent changes in response to market events, such as copyright lawsuits, data breaches, and new generative AI regulations, would be valuable.

How IT Companies are Mitigating Risks:

This study has not examined how some large IT service companies are mitigating the risks associated with generative AI in the absence of clear regulatory guidance. There is a need to identify and document the leading internal AI governance practices, lessons learned, and technical tools used to audit and manage key risks related to Copyright, Bias, and Misinformation.

Sector-Specific Research for Risks and Mitigation:

This research was sector-agnostic and treated IT service companies as a broad category. Clearly, the risks of generative AI differ across sectors and applications. Future work should focus on specific sectors, such as healthcare, finance, and legal services, and on the use of generative AI across different application types. Each sector has unique risks and regulatory bodies that manage risks. Research could delve into sector-specific risks

and the effectiveness of existing regulations and regulators within that sector. A cross-sector analysis of some specific risks and mitigation can be helpful as well.

Cost of Compliance:

As the US government begins implementing AI regulations, an interesting research question is how to measure their financial impact on companies of different sizes and sectors. A baseline study with early regulatory initiatives and late adoption can help in assessing the financial implications. What is the cost of compliance compared to the cost of unmitigated risks? Such cost-benefit analyses will help in evidence-based policymaking.

Research on Specific Risks and Audit Tool:

Future research may focus on developing transparency without compromising the secret sauce of IT companies as they fine-tune the models. Also, there is a need to audit and mitigate biases in fine-tuned models. Given the broad adoption of generative AI, any audit or bias-mitigation solution must be cost-effective and scalable, as it is used by both corporate and private users worldwide.

6.4 Conclusion

This research has examined the relationship between the lack of strict regulation of generative AI and its impact on the adoption in the US IT services sector. A comparative analysis of the US, the UK, and China has covered their regulatory approaches. The US has a voluntary, innovation-focused approach, whereas China has a state-controlled model. The UK follows a principles-based, sectoral framework. Despite their differences in approach, they all seem to struggle to assign accountability and transparency to the opaque supply chain of generative AI. Based on the analysis, it's clear that there is a significant risk assurance gap with the current frameworks and regulations. Interestingly, these gaps are acting as a catalyst in the adoption of generative AI. Another important finding is that the absence of strict rules, particularly regarding Transparency and Copyright, is helping

in the adoption. IT service companies are innovating and advancing across three application types: Business-critical, Business-support, and Creativity. The researcher calls it a first-mover advantage, as these IT companies are not waiting for regulations and compliance but are focusing on gaining a competitive advantage by moving quickly.

The study has also identified input data as the root of key risks (Copyright, Privacy, Bias, Toxicity, and Misinformation) that are not fully mitigated through the current regulations. The gist of these findings is that IT companies are aware of risks. The very regulatory gaps that act as catalysts for rapid innovation and productivity gains are also building a significant risk bubble for the future.

The study highlights that IT service companies need to build a robust ethical AI governance framework and processes. They should not wait for known risks to go out of control and later fix them at high liability and reputational cost. For policymakers, the research underscores the need for targeted, risk-proportionate sector-specific regulations. Transparency and explainability should be top priorities before harmful practices become too entrenched.

REFERENCES

- Anthis, J. R., Lum, K., Ekstrand, M., Feller, A., & Tan, C. (2025) *The Impossibility of Fair LLMs*, Available at: [2025.acl-long.5.pdf](https://arxiv.org/abs/2025.acl-long.5.pdf) (Accessed: 5th September 2025).
- Babbie, E. (2015) *The Practice of Social Research*, 14th edn., Boston, MA: Cengage Learning, ISBN 9781305104945.
- Baker, M. (2025) *The Spread of AI-Generated Misinformation*, Available at: <https://doi.org/10.22541/au.174405164.43017701/v1> (Accessed: 15th May 2025).
- Bhattacharjee, A. (2012) *Social science research: principles, methods, and practices*, 2nd edn., South Carolina: CreateSpace Independent Publishing Platform, ISBN: 9781475146127.
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., ... Liang, P. (2022) *On the Opportunities and Risks of Foundation Models*, Available at: <http://arxiv.org/abs/2108.07258> (Accessed: 12th June 2023).
- Bryman, A., Bell, E. and Kleinknecht, S. (2022) *Social Research Methods*, 6th edn., Ontario, Canada: Oxford University Press, ISBN: 9780190165796.
- Chen, C., Fu, J., & Lyu, L. (2023) *A Pathway Towards Responsible AI Generated Content*, Available at: <http://arxiv.org/abs/2303.01325> (Accessed: 18th July 2024).
- Cheng, J., & Zeng, J. (2023). 'Shaping AI's Future? China in Global AI Governance', *Journal of Contemporary China*, 32(143), pp. 794-810 [Online]. Available at: <https://doi.org/10.1080/10670564.2022.2107391> (Accessed: 13th August 2023).
- Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., & Zimmel, R. (2023) *The economic potential of generative AI: The next productivity frontier*, Available at: <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier#/> (Accessed: 22nd November 2024).
- Creemers, R. (2021) *CHINA'S CONCEPTION OF CYBER SOVEREIGNTY: RHETORIC AND REALIZATION*, Available at: <https://ssrn.com/abstract=3532421> (Accessed: 22nd November 2024).
- Creswell, J.W. and Creswell, J.D. (2022) *Research Design: Qualitative, Quantitative, and Mixed Methods Approaches*, 6th edn., Thousand Oaks, CA: Sage Publications, ISBN 9781071817940.

- Daugherty, P., Ghosh, B., Narain, K., Guan, L., Wilson, J. (2025) *A new era of generative AI for everyone*, Available at: <https://www.accenture.com/content/dam/accenture/final/accenture-com/document/Accenture-A-New-Era-of-Generative-AI-for-Everyone.pdf> (Accessed: 2nd October 2025).
- Department for Science, Innovation and Technology (2023a) *A pro-innovation approach to AI regulation*, Available at: <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper> (Accessed: 22nd January 2025).
- Department for Science, Innovation and Technology (2023b) *Pro-innovation regulation of technologies review: Digital technologies*, Available at: <https://www.gov.uk/government/publications/pro-innovation-regulation-of-technologies-review-digital-technologies> (Accessed: 23rd January 2025).
- Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., ... Wright, R. (2023) ‘Opinion Paper: “So what if ChatGPT wrote it?” Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy’, *International Journal of Information Management*, 71, Available at: <https://doi.org/10.1016/j.ijinfomgt.2023.102642> (Accessed: 20th November 2023).
- Felten, E., Raj, M. & Seamans, R. (2023) *How will Language Modelers like ChatGPT Affect Occupations and Industries?*, Available at: <https://arxiv.org/abs/2303.01157> (Accessed: 2nd January 2024).
- Fung, B. (2023) *Biden administration unveils an AI plan ahead of meeting with tech CEOs*, Available at: <https://www.cnn.com/2023/05/04/tech/white-house-ai-plan/index.html> (Accessed: 12th May 2024).
- G’sell, F. (2024) *Florence G’sell Regulating under Uncertainty: Governance Options for Generative AI*, Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4918704 (Accessed: 10th January 2025).
- Gallegos, I. O., Rossi, R. A., Barrow, J., Tanjim, M. M., Kim, S., Dernoncourt, F., Yu, T., Zhang, R., & Ahmed, N. K. (2024) ‘Bias and Fairness in Large Language Models: A Survey’, *Computational Linguistics*, 50(3), pp. 1097-1179 [Online]. Available at: https://doi.org/10.1162/COLI_A_00524 (Accessed: 15th January 2025).

- Green, S. B. (2010) 'How Many Subjects Does It Take To Do A Regression Analysis', *Multivariate Behavioral Research*, 26(3), pp. 499-510 [Online]. Available at: https://doi.org/10.1207/S15327906MBR2603_7 (Accessed: 17th May 2023).
- Gu, J. (2024) *A Survey on Responsible Generative AI: What to Generate and What Not*, Available at: <https://arxiv.org/pdf/2404.05783> (Accessed: 1st December 2024).
- Gutierrez, C. I., Aguirre, A., Uuk, R., Boine, C. C., & Franklin, M. (2022) *A Proposal for a Definition of General Purpose Artificial Intelligence Systems*, Available at: <https://ssrn.com/abstract=4238951> (Accessed: 18th May 2023).
- Hagendorff, T. (2024) 'Mapping the Ethics of Generative AI: A Comprehensive Scoping Review', *Minds & Machines*, 34(39), Available at: <https://doi.org/10.1007/s11023-024-09694-w> (Accessed: 15th December 2024).
- Hamilton, D. (2023) 'FTC investigating ChatGPT creator OpenAI over consumer protection issues', *AP News*, Available at: <https://apnews.com/article/openai-chatgpt-investigation-federal-ftc-76c6218c506996942282d7f5d608088e> (Accessed: 13th April 2024).
- Harris, L. (2025) 'Regulating Artificial Intelligence: U.S. and International Approaches and Considerations for Congress', *CRS Report R48555*, Washington, D.C.: Congressional Research Service, Available at: <https://www.congress.gov/crs-product/R48555> (Accessed: 15th June 2025).
- IBISWorld (2025) *IT Consulting in the US - Number of Businesses*, Available at: <https://www.ibisworld.com/united-states/number-of-businesses/it-consulting/1415/> (Accessed: 23rd August 2025).
- Ignatius, A. & Bernstein, A. (2023a) 'How Generative AI Changes Productivity', *Harvard Business Review Ideacast*, 2 May, Available at: <https://hbr.org/podcast/2023/05/how-generative-ai-changes-productivity> (Accessed: 3rd August 2023).
- Ignatius, A. & Bernstein, A. (2023b) 'How Generative AI Changes Creativity', *Harvard Business Review Ideacast*, 30 May, Available at: https://hbr.org/podcast/2023/05/how-generative-ai-changes-creativity?ab=at_pod_art_h3x1_s02 (Accessed: 23rd August 2023).
- Interesse, G., (2022) *China to Regulate Deep Synthesis (Deepfake) Technology from 2023*, Available at: <https://www.china-briefing.com/news/china-to-regulate-deep-synthesis-deep-fake-technology-starting-january-2023/> (Accessed: 3rd March 2024).
- Jaidka, K., Yue, A., Hsu, W., Kan, Y., Kankan-halli, M., Li Lee, M., Chen, T., Chesterman, S., Kan, M.-Y., Kankanhalli, M., Seres, G., Sim, T., Taihagh, A., Tung, A., & Xiao, X.

- (2025) 'Misinformation, Disinformation, and Generative AI: Implications for Perception and Policy', *Digital Government: Research and Practice*, 6(1). Article No. 11, pp. 1-15 [Online]. Available at: <https://doi.org/10.1145/3689372> (Accessed: 25th February 2025).
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y., Chen, D., Dai, W., Chan, S., Madotto, A., & Fung, P. (2023) 'Survey of Hallucination in Natural Language Generation', *ACM Computing Surveys*, 55(12), Article No. 248, pp. 1-38 [Online]. Available at: <https://dl.acm.org/doi/abs/10.1145/3571730> (Accessed: 3rd April 2024).
- Kandpal, N., Wallace, E., & Raffel, C. (2022) *Deduplicating Training Data Mitigates Privacy Risks in Language Models*, Available at: <http://arxiv.org/abs/2202.06539> (Accessed: 4th July 2024).
- Kelly, J. (2023) *Goldman Sachs Predicts 300 Million Jobs Will Be Lost Or Degraded By Artificial Intelligence*, Available at: <https://www.forbes.com/sites/jackkelly/2023/03/31/goldman-sachs-predicts-300-million-jobs-will-be-lost-or-degraded-by-artificial-intelligence/> (Accessed: 7th April 2024).
- Kolt, N., (2024) *Algorithmic Black Swans*, Available at: <https://wustllawreview.org/2024/04/18/algorithmic-black-swans/> (Accessed: 11th January 2025).
- Lee, K., Feder Cooper, A., Grimmelmann, J. (2024) *Talkin' 'Bout AI Generation: Copyright and the Generative-AI Supply Chain*, Available at: <https://afedercooper.info/paper/lee2023talkin.pdf> (Accessed: 9th December 2024).
- Markov, T., Zhang, C., Agarwal, S., Eloundou Nekoul, F., Lee, T., Adler, S., Jiang, A. and Weng, L. (2023) 'A Holistic Approach to Undesired Content Detection in the Real World', *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(12), pp. 15009-15018 [Online]. Available at: <https://ojs.aaai.org/index.php/AAAI/article/view/26752> (Accessed: 8th April 2024).
- Migliorini, S. (2024) 'China's Interim Measures on generative AI: Origin, content and significance', *Computer Law & Security Review*, 53(105985), Available at: <https://doi.org/10.1016/J.CLSR.2024.105985> (Accessed: 17th December 2024).
- NIST (2020) *NIST PRIVACY FRAMEWORK: A TOOL FOR IMPROVING PRIVACY THROUGH ENTERPRISE RISK MANAGEMENT, VERSION 1.0*, Available at: <https://doi.org/10.6028/NIST.CSWP.01162020> (Accessed: 17th May 2023).

- NIST (2023) ‘Artificial Intelligence Risk Management Framework (AI RMF 1.0)’, *NIST AI 100-1*, Gaithersburg, MD: U.S. Department of Commerce, Available at: <https://doi.org/10.6028/NIST.AI.100-1> (Accessed: 12th April 2024).
- OECD (2022) ‘OECD framework for the classification of AI systems’, *OECD Digital Economy Papers*, Paris: OECD Publishing, Available at: https://www.oecd.org/en/publications/oecd-framework-for-the-classification-of-ai-systems_cb6d9eca-en.html (Accessed: 16th July 2023).
- Office of the Governor, State of Alabama (2024) *Executive Order No. 738: Artificial Intelligence*, Available at: <https://governor.alabama.gov/assets/2024/02/EO-738-Artificial-Intelligence.pdf> (Accessed: 4th February 2025).
- OpenAI (2023) *Our approach to AI safety*, Available at: <https://openai.com/index/our-approach-to-ai-safety/> (Accessed: 12th January 2024).
- O’Shaughnessy, M., & Sheehan, M. (2023) *Lessons From the World’s Two Experiments in AI Governance*, Available at: <https://carnegieendowment.org/posts/2023/02/lessons-from-the-worlds-two-experiments-in-ai-governance?lang=en> (Accessed: 10th March 2024).
- Parrish, A., Chen, A., Nangia, N., Padmakumar, V., Phang, J., Thompson, J., Htut, P. M., & Bowman, S. R. (2022) ‘BBQ: A Hand-Built Bias Benchmark for Question Answering’, *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 2086-2105 [Online]. Available at: <https://aclanthology.org/2022.findings-acl.165/> (Accessed: 12th April 2023).
- Paul, R. K., & Sarkar, B. (2023) ‘Generative AI and Ethical Considerations for Trustworthy AI Implementation’, *INTERNATIONAL JOURNAL OF ARTIFICIAL INTELLIGENCE & MACHINE LEARNING (IJAIML)*, 2(1), Article No. IJAIML_02_01_010, pp. 95-102 [Online]. Available at: https://iaeme.com/Home/article_id/IJAIML_02_01_010 (Accessed: 6th January 2024).
- Perez, F., & Ribeiro, I. (2022) *Ignore Previous Prompt: Attack Techniques For Language Models*, Available at: <https://arxiv.org/abs/2211.09527> (Accessed: 10th August 2023).
- Perrigo, B. (2023) *How to Spot an AI-Generated Image Like the “Balenciaga Pope.”*, Available at: <https://time.com/6266606/how-to-spot-deepfake-pope/> (Accessed: 4th March 2024).
- Rekabsaz, N., & Schedl, M. (2020) *Do Neural Ranking Models Intensify Gender Bias?*, Available at: <https://dl.acm.org/doi/abs/10.1145/3397271.3401280> (Accessed: 18th June 2023).

- Roberts, M.E. (2020) *Censored: Distraction and Diversion Inside China's Great Firewall*. 1st edn., Princeton, NJ: Princeton University Press, ISBN 9780691178868.
- Rough, E., & Sutherland, N. (2024) *Artificial intelligence: A reading list*, Available at: <https://researchbriefings.files.parliament.uk/documents/CBP-10003/CBP-10003.pdf> (Accessed: 7th May 2025).
- Rui, F., Liu, Y. (2023) *Wang Liming: The biggest risk of generative AI is technological backwardness and being "held hostage" by technological bottlenecks.*, Available at: <https://news.caijingmobile.com/article/detail/496460> (Accessed: 15th August 2024).
- Russell, S., Perset, K., & Grobelnik, M. (2023) *Updates to the OECD's definition of an AI system explained*, Available at: <https://oecd.ai/en/wonk/ai-system-definition-update> (Accessed: 13th July 2024).
- Saunders, M., Lewis, P. and Thornhill, A. (2023) *Research methods for business students*, 9th edn., Harlow: Pearson Education Limited, ISBN 9781292402727, Available at: https://www.researchgate.net/publication/240218229_Research_Methods_for_Business_Students (Accessed: 10th January 2024).
- Schumpeter, J.A. (1994) *Capitalism, Socialism and Democracy*, 1st edn., London: Routledge, ISBN 9780415107624.
- Stanford CRFM (2024) *Foundation Model Transparency Index*, Available at: <https://crfm.stanford.edu/fmti/May-2024/index.html> (Accessed: 10th December 2024).
- Sutton, R., & Hargadon, A. (1996) *Brainstorming Groups in Context: Effectiveness in a Product Design Firm*, Available at: <https://web.mit.edu/~mcyang/www/papers/suttonHargadon96.pdf> (Accessed: 19th April 2023).
- Tamkin, A., Brundage, M., Clark, J., & Ganguli, D. (2021) *Understanding the Capabilities, Limitations, and Societal Impact of Large Language Models*, Available at: <http://arxiv.org/abs/2102.02503> (Accessed: 11th November 2023).
- The President of the United States (2019) 'Executive Order 13859 of February 11, 2019: Maintaining American Leadership in Artificial Intelligence.' *Federal Register*, 84(31), pp. 3967-3972 [Online]. Available at: <https://www.govinfo.gov/content/pkg/DCPD-201900073/pdf/DCPD-201900073.pdf> (Accessed: 18th July 2023).
- The White House (2025a) *Winning the Race America's AI Action Plan*, Available at: <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf> (Accessed: 10th August 2025).

- The White House (2025b) *Removing Barriers to American Leadership in Artificial Intelligence*, Available at: <https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/> (Accessed: 10th August 2025).
- UK Parliament (2025) ‘Artificial Intelligence (Regulation) Bill [HL]’, *Parliamentary Bills*, Available at: <https://bills.parliament.uk/bills/3942> (Accessed: 16th April 2025).
- U.S. Congress House (2022) ‘Algorithmic Accountability Act of 2022’, *H.R. 6580, 117th Congress*, Available at: <https://www.congress.gov/117/bills/hr6580/BILLS-117hr6580ih.pdf> (Accessed: 3rd August 2023).
- U.S. Congress. Senate (2024a) ‘S.4714 - Fraudulent Artificial Intelligence Regulations (FAIR) Elections Act of 2024’, *118th Congress*, Available at: <https://www.congress.gov/bill/118th-congress/senate-bill/4714/text/is> (Accessed: 3rd December 2024).
- U.S. Congress Senate (2024b) ‘S.4862 - A bill to ensure that new advances in artificial intelligence are ethically adopted to improve the health of all individuals, and for other purposes.’, *118th Congress*, Available at: <https://www.congress.gov/bill/118th-congress/senate-bill/4862/text> (Accessed: 3rd December 2024).
- U.S. Congress Senate (2025) ‘S.2117 - Preventing Deep Fake Scams Act’, *119th Congress*, Available at: <https://www.congress.gov/bill/119th-congress/senate-bill/2117/text/is> (Accessed: 3rd July 2025).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. and Polosukhin, I. (2017) *Attention is all you need*, Available at: <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html> (Accessed: 23rd February 2023).
- Veltman, C. (2023) *Thousands of authors urge AI companies to stop using work without permission*, Available at: <https://www.wbur.org/npr/1187523435/thousands-of-authors-urge-ai-companies-to-stop-using-work-without-permission> (Accessed: 5th August 2024).
- Wang, X., & Wu, Y. C. (2024) ‘Balancing Innovation and Regulation in the Age of Generative Artificial Intelligence’, *Journal of Information Policy*, 14, pp. 385–416 [Online]. Available at: <https://doi.org/10.5325/jinfopoli.14.2024.0012> (Accessed: 5th December 2024).
- Webster, G., Creemers, R., Triolo, P., & Kania, E. (2017) *Full Translation: China's 'New Generation Artificial Intelligence Development Plan' (2017)*, Available at:

<https://www.newamerica.org/cybersecurity-initiative/digichina/blog/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/> (Accessed: 15th December 2023).

Xiang Chloe. (2022) *GitHub Users File a Class-Action Lawsuit Against Microsoft for Training an AI Tool With Their Code*, Available at: <https://www.vice.com/en/article/bvm3k5/github-users-file-a-class-action-lawsuit-against-microsoft-for-training-an-ai-tool-with-their-code> (Accessed: 5th December 2023).

Yan, J., Yadav, V., Li, S., Chen, L., Tang, Z., Wang, H., Srinivasan, V., Ren, X., & Jin, H. (2023) *Backdooring Instruction-Tuned Large Language Models with Virtual Prompt Injection*, Available at: <https://arxiv.org/abs/2307.16888> (Accessed: 5th July 2024).

Zou, M., & Zhang, L. (2025) *Navigating China's regulatory approach to generative artificial intelligence and large language models*, Available at: <https://doi.org/10.1017/cfl.2024.4> (Accessed: 5th March 2025).

APPENDIX A
SURVEY COVER LETTER

Dear [Name / Research Colleague],

I am Rohit Kumar, a Doctor of Business Administration candidate at the Swiss School of Business Management. I am currently conducting a dissertation study titled **“Responsible Generative AI: Global Regulatory Gaps, US Innovation Catalysts,”** which explores how the **absence of strict regulations** influences the *adoption and ethical deployment* of generative AI technologies.

This research is conducted strictly for academic purposes as part of my DBA dissertation. Your insights will inform theoretical frameworks and best-practice recommendations for this emerging regulatory landscape.

What’s Involved

- A 10-15-minute **anonymous** and **voluntary** online survey.
- **No identifying information (PII, company you work for) will be asked for or collected.**
- All responses will be kept **confidential**, and **access to individual data will be restricted to me, the researcher.**

How Your Input Will Be Used

Results will be reported only in aggregate form. You will *not* be identified; your data will be used solely to contribute to academic knowledge and inform policy considerations.

Spread the Word

To support a diverse and robust sample, I would deeply appreciate it if you could **forward this invitation (with the survey link below)** to any colleagues, peers, or professional connections who meet the criteria (e.g., Experience in implementing

generative AI solutions). This “snowball” approach will significantly strengthen the validity of the findings.

Survey Link:

<https://forms.gle/nPGrVTQirsckqfnb7>

If you have any questions about the study’s purpose, methodology, or your rights as a participant, please do not hesitate to contact me at rohit@ssbm.ch.

Thank you for considering this invitation. Your participation-and any referrals you make-will play a crucial role in shaping ethical guidance for generative AI.

Warm regards,

Rohit

DBA Candidate, Swiss School of Business Management

APPENDIX B
SURVEY QUESTIONS

Eligibility Questions

Q1	Have you implemented any generative AI projects or solutions?	Yes	No
Q2	Are you familiar with the risks associated with generative AI?	Yes	No

Section A: Personal Demographic and Project-Related Questions

Q3	Which industry do you identify with your work domain	IT-Service	IT-Product	Other	
Q4	How many years of AI experience do you have?	Less than 1 year	1 to 5 years	6 to 10 years	more than 10 years
Q5	How many AI projects experience you have?	Less than 3	3 to 5	above 5	
Q6	Have you implemented a generative AI project for productivity enhancement, innovation acceleration, or both	Productivity enhancement	Innovation acceleration	Both	
Q7	Have you implemented a generative AI project for a business-critical application (e.g., Finance, Compliance, HR, supply chain) or business-support application (e.g., customer support, employee benefits) or creativity application (e.g., new business offering, new product development, assistant tool for Subject Matter Expert, etc)	Bussiness critical	Business support	Creativity	
Q8	Have you implemented a generative AI project for internal purposes or for your clients?	Internal	External		
Q9	Your role in the implementation of the generative AI project	Business	Technical		

Section B: Perceived risk not fully covered by existing AI regulations and frameworks

The risk below is not fully covered for the generative AI solution you have implemented or plan to implement under the existing AI regulations and	1 = Strongly Disagree	2 = Disagree	3 = Neutral	4 = Agree	5 = Strongly Agree
--	--------------------------	-----------------	----------------	--------------	-----------------------

	Risk Management Frameworks (RMFs).					
Q10	<i>Accountability risk</i>					
Q11	<i>Fairness risk</i>					
Q12	<i>Transparency risk</i>					
Q13	<i>Explainability risk</i>					
Q14	<i>Resilience risk</i>					
Q15	<i>Privacy risk</i>					
Q16	<i>Safety risk</i>					
Q17	<i>Security risk</i>					
Q18	<i>Sustainability risk</i>					

Section C: Business-critical application risk coverage and impact of regulations

	The risk below is not fully covered for the generative AI solution for business-critical applications	1 = Strongly Disagree	2 = Disagree	3 = Neutral or Not Applicable	4 = Agree	5 = Strongly Agree
Q19	<i>Copyright risk</i>					
Q20	<i>Privacy risk</i>					
Q21	<i>Bias risk</i>					
Q22	<i>Toxicity risk</i>					
Q23	<i>Misinformation risk</i>					
	Lack of well-defined, regulated risk positively affects the adoption of generative AI solutions for business-critical applications	1 = Strongly Disagree	2 = Disagree	3 = Neutral	4 = Agree	5 = Strongly Agree
Q24	<i>Accountability risk</i>					
Q25	<i>Fairness risk</i>					
Q26	<i>Transparency risk</i>					
Q27	<i>Explainability risk</i>					
Q28	<i>Resilience risk</i>					
Q29	<i>Privacy risk</i>					
Q30	<i>Safety risk</i>					

Q31	<i>Security risk</i>					
Q32	<i>Sustainability risk</i>					
Q33	<i>Copyright risk</i>					
Q34	<i>Bias risk</i>					
Q35	<i>Toxicity risk</i>					
Q36	<i>Misinformation risk</i>					

Section D: Business-support application risk coverage and impact of regulations

	The risk below is not fully covered for the generative AI solution for a business-support application	1 = Strongly Disagree	2 = Disagree	3 = Neutral or Not Applicable	4 = Agree	5 = Strongly Agree
Q37	<i>Copyright risk</i>					
Q38	<i>Privacy risk</i>					
Q39	<i>Bias risk</i>					
Q40	<i>Toxicity risk</i>					
Q41	<i>Misinformation risk</i>					
	Lack of well-defined, regulated risk positively affects the adoption of generative AI solutions for business-support applications	1 = Strongly Disagree	2 = Disagree	3 = Neutral or Not Applicable	4 = Agree	5 = Strongly Agree
Q42	<i>Accountability risk</i>					
Q43	<i>Fairness risk</i>					
Q44	<i>Transparency risk</i>					
Q45	<i>Explainability risk</i>					
Q46	<i>Resilience risk</i>					
Q47	<i>Privacy risk</i>					
Q48	<i>Safety risk</i>					
Q49	<i>Security risk</i>					
Q50	<i>Sustainability risk</i>					
Q51	<i>Copyright risk</i>					
Q52	<i>Bias risk</i>					
Q53	<i>Toxicity risk</i>					

Q54	<i>Misinformation risk</i>					
-----	----------------------------	--	--	--	--	--

Section E: Creativity application risk coverage and impact of regulations

	The risk below is not fully covered for the generative AI solution for the creativity application	1 = Strongly Disagree	2 = Disagree	3 = Neutral or Not Applicable	4 = Agree	5 = Strongly Agree
Q55	<i>Copyright risk</i>					
Q56	<i>Privacy risk</i>					
Q57	<i>Bias risk</i>					
Q58	<i>Toxicity risk</i>					
Q59	<i>Misinformation risk</i>					
	Lack of well-defined, regulated risk positively affects the adoption of generative AI solutions for creativity applications	1 = Strongly Disagree	2 = Disagree	3 = Neutral or Not Applicable	4 = Agree	5 = Strongly Agree
Q60	<i>Accountability risk</i>					
Q61	<i>Fairness risk</i>					
Q62	<i>Transparency risk</i>					
Q63	<i>Explainability risk</i>					
Q64	<i>Resilience risk</i>					
Q65	<i>Privacy risk</i>					
Q66	<i>Safety risk</i>					
Q67	<i>Security risk</i>					
Q68	<i>Sustainability risk</i>					
Q69	<i>Copyright risk</i>					
Q70	<i>Bias risk</i>					
Q71	<i>Toxicity risk</i>					
Q72	<i>Misinformation risk</i>					