Working Paper Series

AUGUST.2025

Executive Concerns with AI Adoption: Identifying Business and Security Risks

August 25, 2025

Authors:

Keri Pearlson, MIT Sloan School of Management
Rajiv Dattani, AIUC

This research was supported by MIT Sloan School of Management and the Artificial Intelligence Underwriting Company (AIUC).

Executive Concerns with AI Adoption: Identifying Business and Security Risks[1]

Dr. Keri Pearlson, MIT Sloan School of Management
Rajiv Dattani, AIUC

August 25, 2025

Abstract

While AI capabilities are advancing at an unprecedented pace, many organizations lack a clear roadmap for mitigating the accompanying risks and vulnerabilities. To better understand the challenges faced by industry, we surveyed senior enterprise leaders from a variety of industries with experience implementing AI systems and facing resulting vulnerabilities. The results were unambiguous: data leaks, intellectual property infringement, adversarial attacks, and the risk of falling behind competitors due to slow AI adoption emerged as the most pressing concerns. These findings underscore the need for a targeted, practical framework to provide a roadmap for navigating AI risks in the enterprise.

Motivation and Background

A growing body of research has examined the business risks of AI adoption, underscoring that value capture depends on successful pairing of innovation with governance. Choudhary and Kar (2025) model how AI interacts with organizational privacy and cybersecurity practices, linking effective alignment to greater resilience against breaches (Choudhary & Kar, 2025). Humphreys et al. (2024) warn that "AI hype" can drive premature deployments that amplify cybersecurity exposure and ethical lapses, reframing generative AI rollouts as an enterprise risk and moral responsibility problem (Humphreys et al., 2024). And Wagman et al. (2025) has shown that reliability/hallucinations, data security concerns, and academic publishing/citation issues constrain AI use in scientific organizations (Wagman et al., 2025).

The impact of AI on cybersecurity has also been studied, finding that AI reshapes cyber-risk profiles in ways that are both enabling and hazardous. Jada and Mayayise (2024) found evidence that AI can improve automation, threat intelligence, and defensive effectiveness, while also expanding attack surfaces via adversarial attacks and data quality dependencies (Jada & Mayayise, 2024). Consistent with this, Humphreys et al. (2024) stress that rushing to deploy AI without adequate safeguards heightens the likelihood and organizational impact of cyber incidents (Humphreys et al., 2024).

AI also has implications for governance and leadership. Bello y Villarino and Bronitt (2024) examine AI through a regulatory lens and argue that AI-enabled compliance and monitoring tools alter corporate governance dynamics, increasing legal and oversight exposure and calling for stronger governance arrangements, including board-level attention (Bello y Villarino & Bronitt, 2024). Chowdhury et al. (2024) similarly position executive, cross-functional engagement as a prerequisite for responsible AI integration in core processes and workforce practices (Chowdhury et al., 2024). In practice, Wagman et al. (2025)

[1] The authors wish to thank Emil Lassen and Rune Kvist for their input and assistance with earlier drafts of this paper; ASFAI and Lena Smart for their support and contributions to this research; as well as the executives who shared their insights with us for this study. For more information on this research, please contact the lead author at kerip@mit.edu.

show that reliability and security concerns are salient to knowledge workers in high-stakes settings, reinforcing the role of executive oversight in setting trust, quality, and safety expectations for AI use (Wagman et al., 2025).

From our review of the literature, we found six key constructs to ground our investigation of the potential concerns to enterprise leaders. Table 1 summarizes these constructs and our definitions.

Construct	Definition
Data and Privacy	Protect user and enterprise data from AI leaks and unauthorized AI system access or usage
Security	Protect AI systems from adversarial attacks, data poisoning, and unauthorized access
Safety	Prevent AI systems from generating harmful outputs through internal and external evaluations, monitoring, and safeguards
Society	Prevent AI from enabling societal harm through cyberattacks or national security risks
Accountability	Enforce oversight of AI systems and develop emergency response protocols to mitigate issues
Reliability	Prevent hallucinations and system access that could cause harm

Table 1: Constructs and Definitions

Collectively, the literature review highlights both the new opportunities and risks that AI brings, but also the importance of addressing the new challenges at the board and executive levels. In walking the tightrope, executives must balance business opportunities with security concerns. The study reported in this paper addresses many of these issue but from an executive leadership perspective. We want to understand what enterprise leaders say are their top concerns around AI adoption, governance and risk management.

Methodology: Surveying Enterprise Leaders About Their AI Concerns

Since executives must balance business with security risk, we began our survey study by asking two open-ended questions:

1. What is the biggest BUSINESS risk you see for adopting AI in organizations?
2. What is the biggest SECURITY risk you see for adopting AI in organizations?

Diving deeper into the constructs we found in the literature review, we included questions about how concerned the executives were across six dimensions: Data and Privacy, Security, Safety, Society, Accountability and Reliability. In addition, survey respondents were asked to share their job title and industry to provide us with relevant demographics. Full survey questions are included as *Appendix 1.*

The survey was administered to members of the ASFAI in June 2025. The 28 respondents fell into 3 broad categories: CEOs/partners (36%), technical executives including CTOs, CISOs and CROs (32%), and other executives including Heads of AI and COOs (32%). The most prominent sectors represented were technology (32%), finance/financial services (25%), healthcare and pharma (21%), and media (7%).

Findings: Enterprise Leaders Are Concerned and Confused About AI

The survey findings indicate that not only are leaders concerned about the implementation of AI in their organizations, but they see both security and business risks as important areas to address.

When asked about the biggest ***business*** risk the respondents see from adopting AI in their organization, the most concerns were in data security and privacy, compliance and governance and reliability/model risk. Table 2 summarizes the responses by categorizing them into 12 categories.

Business Risks	Examples from Responses	Frequency
1. Data Security & Privacy	Availability/accuracy/security, data breach, data privacy liability, cyber incident liability, loss of customer data, exposing PII/IP, ransomware, stealing secret sauce	8
2. Compliance & Governance	Compliance with data retention policies, compliance with frameworks, lack of frameworks for high-risk systems	3
3. Competitive & Strategic Risk	Failing to keep up with competitors, devaluing products, not leveraging AI as strategic capability	3
4. Culture & Change Management	Change management, culture change, using AI just to replace workers	3
5. Leadership & Skills Gap	Leadership skill gap/lack of management buy-in, upskilling	2
6. Adoption & Scaling Challenges	Lack of adopting/becoming a laggard (3 mentions), failure to scale	4
7. Reliability & Model Risk	Rubbish in/rubbish out, hallucinations in decision making, misinterpreted insights/bad decisions	3
8. Accountability & Oversight	Unclear accountability, automation without monitoring	2
9. Safety & Control Risks	Unsafe control functions from generated code	1
10. Financial & Cost Concerns	Cost	1
11. Reputation Risk	Loss of reputation	1
12. IP-Specific Risks	IP risk for healthcare companies, exposing IP publicly	2

Table 2: Categorization of Business Risks

When asked about the biggest ***security*** risk the respondents see from adopting AI in their organization, the most concerns were in data security and breaches, IP leaks, malicious actors, and privacy risks. Table 3 summarizes the responses by categorizing them into ten categories:

Security Risks	Examples from Responses	Frequency
1. Data Security & Breaches	Data leaks, Data breach, Data breaches/ransomware, Data loss, Data leakage, Risk of loss/corruption of critical info, PII/trade secret exposure, Sharing patient data in public models	9
2. Privacy Risks	Privacy, Stealing PI, Impersonation	3
3. Intellectual Property (IP) Risks	Rights violations & IP infringement, IP and information loss, IP exposed via 3rd party apps, IP leakage	4
4. Malicious Actors & Threats	Bad actors, Gen AI fraud/phishing, Infiltration (in/out without trace, loss of operations)	3
5. Ransomware & Cyber attacks	Ransomware, Security ability to manage/respond to ransomware & attacks	2
6. Misuse of AI / Unsafe Deployment	Deploying without multi-layered frameworks, Uploading confidential data & poor review, Reverse prompt engineering / lack of visibility into agentic tools	3

7. Loss of Control & Oversight	Loss of control, Back drop to my data	2
8. Employee & Human Factor Risks	Limited employee understanding of AI/IP tools	1
9. Operational Risks	Infiltration leading to loss of control of operations (also overlaps with Malicious Actors)	1
10. Accountability & Governance	Gaps in ability to recognize/respond quickly to threats	1

Table 3: Categories of Security Risks

Enterprise leaders are clearly concerned about the way AI is used in their enterprises. For *all* areas, at least 80% of respondents expressed they have *concerns*, *significant concerns* or *strong concerns.* We were surprised to find that only in two categories do we see *any* respondents marking *no concern* and, in both instances, it is below 10% of respondents.

The six categories largely reflect the areas where enterprise leaders are concerned - when asked if there were any risks missing, only 7 executives responded, and responses were mainly concerned with internal employee use of AI. This is an area to explore in future studies.

Data & Privacy was the top area of concern with 82% of the respondents marking it as a *strong* or *significant* concern, and an average rating of 4.25/5.00. This finding aligns with the qualitative results where executives mentioned inadvertent data leaks, adversarial data breaches, and IP infringement risks as major concerns. Security concerns followed closely**.** 75% of the respondents rated it as a *strong* or *significant* concern with an average rating of 3.96/5.00.

None of the respondents rated Data and Privacy, Security, Safety or Accountability a 1, or *no concern*. This was interesting since at the same time, Safety and Accountability had the lowest number of respondents rate it as a *strong concern*. A summary of all the responses is shown in Table 4.

AI Risk Constructs	**1 - No Concern**	**2 - Some Concern**	**3 - Concern**	**4 - Significant Concern**	**5 - Strong Concern**	**Avg. score**
Data & Privacy: Protect user and enterprise data from AI leaks and unauthorized AI system access or usage	0%	7%	11%	32%	50%	**4.25**
Security: Protect AI systems from adversarial attacks, data poisoning, and unauthorized access	0%	14%	11%	39%	36%	**3.96**
Safety: Prevent AI systems from generating harmful outputs through internal and external evaluations, monitoring, and safeguards	0%	11%	18%	39%	32%	**3.93**
Society: Prevent AI from enabling societal harm through cyberattacks or national security risks	7%	11%	18%	11%	54%	**3.93**
Accountability: Enforce oversight of AI systems and develop emergency response protocols to mitigate issues	0%	18%	25%	21%	36%	**3.75**

Reliability: Prevent hallucinations and system access that could cause harm	4%	14%	21%	32%	29%		**3.68**

Table 4: Level of Concern for Risks Constructs (% of respondents)

Discussion: What we Learned From the Findings

In this section, we discuss the findings. We start with learnings from analysis of the business and security risks, noting that they were not distinct lists. Instead, there was significant overlap. Further, the risks listed in our open-ended questions closely resembled the list of risks we used when we asked respondents about how concerned they were for 6 categories of risk. We then draw some conclusions of additional issues and concerns that were supported by the data.

Business and Security Risks Overlap

Respondents to this survey provided insightful ideas about the business and security risks they foresee from the adoption of AI in their organizations. The open-ended questions sought to elicit ideas that could supplement the concerns we found in our initial literature review. Tables 2 and 3 above summarize all the responses received for business and security risks. However, we noted that the responses to these two questions had significant overlap. It's possible that respondents could not differentiate between a business risk and a security risk because when it comes to the impact AI will have on their organization, security risks create the business risks. They are not necessarily distinct risks or concerns. We did additional analysis[2] of the data seeking further consolidation of the responses by considering the overlap and found that the responses could further be combined into seven categories, shown in Table 5.

Risks and Concerns	Examples from Responses	Frequency
1. Data Security, Privacy & IP Risks	Data leaks & leakage, Data loss, Data breaches/ransomware, PII & trade secret exposure, Sharing patient data in public models, Exposing IP via 3rd parties, Rights violations & IP infringement, IP leakage, Stealing PI, Stealing business's "secret sauce"	20
2. Governance, Compliance & Oversight	Compliance with data retention policies, Compliance with frameworks, Lack of frameworks for high-risk systems, Unclear accountability, Deploying without multi-layered evaluation frameworks, Gaps in recognizing/responding to threats, Automation without monitoring, "Backdrop to my data"	9
3. Model Reliability & Safe Deployment	Rubbish in/rubbish out, Hallucinations in decision-making, Misinterpreted insights, Unsafe control functions from generated code, Uploading confidential data without review, Reverse prompt engineering, AI making uncontrolled production changes	7
4. People, Skills & Culture	Change management, Culture change, Using AI to just replace workers, Leadership skill gap, Lack of buy-in from management, Upskilling needs, Employees not understanding how to use AI/IP tools properly	7
5. Adoption, Scaling & Competitive Risk	Lack of adoption / laggard (3 mentions), Failure to scale, Failing to keep up with competitors, Devaluing products, Not realizing AI is strategic, Not leveraging AI properly	7
6. Malicious Actors & Threats	Bad actors, GenAI fraud/phishing, Infiltration (in/out without trace, loss of operations), Cyber incident liability, Ransomware attacks	5

[2] Further analysis of results from Tables 2 and 3 resulted in these seven categories shown in Table 5. The additional analysis included thematic clustering by grouping semantically similar items referring to the same type of risk, followed by a frequency analysis for each cluster of risks, and ultimately merging overlapping and redundant clusters. Generative AI tools were used to complete this analysis.

7. Financial & Reputation Risk	Cost, Loss of reputation	2

Table 5: Categorized List of Consolidated Business and Security Risks From Respondents

These seven areas are closely aligned with the six categories of risk used in survey question 4 where we asked respondents to quantitively rate their concern for each risk. Table 6 displays the two list of risks for easy comparison.

Risk categories Identified by this Study	Risk Construct from Our Survey
(source: Table 5 above)	(source: Table 1 above)
1. Data Security, Privacy & IP Risks	1. Data and Privacy
2. Malicious Actors & Threats	2. Security
3. Adoption, Scaling & Competitive Risk	3. Safety
4. People, Skills & Culture	4. Society
5. Governance, Compliance & Oversight	5. Accountability
6. Model Reliability & Safe Deployment	6. Reliability
7. Financial & Reputation Risk	

Table 6: Comparing Risk Categories

Additional Concerns

The open-ended responses provided additional insights that can loosely be described by these five concerns:

Concern #1: Data leaks
Data leaks came up several times in the open-ended responses as a key concern. The qualitative responses highlighted concerns with both inadvertent leaks from, e.g., misconfigured AI agents or access control settings, but also concerns with data leaks resulting from new adversarial threat vectors that could lead to data breaches. A recent example of the latter is the McDonald's case which exposed 64 million job applicant records, including names, email addresses and phone numbers, on its AI hiring platform (Identity Theft Resource Center, 2025). Executives are justified in worrying about data leaks given even a single leak can severely damage customer trust and trigger regulatory penalties. Proactive monitoring, strong access controls, and contractual safeguards with AI vendors are therefore becoming board-level priorities.

Concern #2: Adversarial vulnerabilities
Adversarial vulnerabilities more generally surfaced as another key concern. In the responses, leaders named concerns including prompt injection, jailbreaking attempts, training data poisoning, data scraping, and attack types that leave no trace. For example, one cyber tech executive named his top concern as: *"Reverse prompt engineering, not knowing the number of AI agentic tools internal to the tenant or way of determining the number".* The challenge for leaders is that these attacks are novel and evolving quickly, often outpacing standard security playbooks. As a result, many organizations are investing in red-teaming and adversarial testing specifically tailored to AI systems – it is critical that these practices are reviewed regularly in the coming time as new threat vectors, attack patterns and mitigation strategies are changing.

Concern #3: IP infringement
Several respondents named IP infringements as their number key concern. For example, the Head of AI of a major healthcare company named their top concern as *"IP being exposed through 3rd party applica-*

tions" while the CEO of a large tech company named their top concern as *"Exposing PII and company private or IP that should not be in the public domain."*. For enterprise leaders, IP leakage is not only a legal and financial risk but also a competitive one – losing proprietary data could undermine years of R&D. Executives are therefore emphasizing stricter policies around how employees use third-party AI tools and are exploring contractual indemnities from vendors.

Concern #4: Worst-case scenarios

The uncertainty about the novel risks that AI was reflected in the responses that named potential worst case scenarios as a top concern. For example, one CEO from the financial sector named their top concern as *"Infiltration, in and out in seconds and leaving no trace. And loss of control of operations".* With a growing body of recent examples highlighting AI-specific risks, such as the zero-click vulnerability identified in Microsoft Copilot which could allow hackers to access data without any specific user interaction (Cybersecurity Dive, 2025), it is not surprising that executives tasked with keeping their organizations secure highlight concerns like these. These scenarios highlight how traditional risk management models may fall short when applied to AI. Leaders increasingly see the need for specialized incident response planning and scenario exercises to prepare for low-probability but high-impact failures.

Concern #5: Losing out to competition

Finally, several executives addressed the tightrope that enterprises are facing with competitive pressures to adopt AI on one side and thoughtful risk mitigation on the other. For example, a CIO in the financial services sector named their top concern simply as *"Failing to keep up with competitors".* Beyond the challenge of navigating the tightrope, the survey responses highlighted the internal AI skill gaps and steep learning curves that many organizations are facing as they move to adopt AI. This dynamic means executives cannot afford to delay adoption indefinitely, but they also recognize that rushing in without proper governance could be just as damaging. Many are approaching this by piloting AI in low-risk domains first, while building internal expertise and risk frameworks before scaling.

Actionable Insights, Future Research and Conclusions

This study makes it clear that new solutions are needed to equip enterprise leaders with the tools and frameworks they need to navigate the tightrope. We highlight three areas where executives can take immediate action:

First, strengthen data governance and access controls: Conduct a comprehensive audit of where sensitive data can enter or leave AI systems, implement tiered access permissions for AI agents, and require regular validation of configuration settings to reduce inadvertent and adversarial leaks.

Second, institutionalize adversarial resilience testing: Integrate red-teaming, prompt-injection simulations, and model output audits into your AI deployment lifecycle, ensuring findings are reviewed at the executive level and remediation steps are tracked to closure.

Third, balance competitive adoption with capability building: Establish an AI adoption roadmap that sequences high-impact, lower-risk applications first while building internal skills, governance processes, and monitoring capabilities before scaling into more sensitive or high-stakes use cases

Additional research must be done to better understand the risks and concerns that AI generates for executives and their organizations. This study showed several concerns raised by executives, but future stud-

ies might compare these concerns with vulnerabilities from AI technology implementations, track the frequency and severity of the business impact these risks create, monitor how many of (and how) these risks become actual events for organizations, and quantify the costs to remedy the resulting business impact. In addition, while this study utilized a community of AI-knowledgeable executives, a future study might differentiate between roles the executives have: CEOs, Board Members, Technical Executives, Cybersecurity Executives, Financial Executives, Operations Executives, and Legal/Compliance Executives are likely to approach the question of risks and concerns from AI from very different perspectives. Building "AI-proofed organizations" requires the input, insights, expertise, and leadership from all these perspectives.

AI is rapidly changing the business and the security landscape, and executives are concerned. The pace of introduction of new technologies combined with the promises the new technologies make for new sources of revenue, new opportunities for cutting costs, new innovations for product and service offerings, and new ways employees can streamline their work mean that the risks and concerns of executives leading enterprises have a difficult job ahead of them. Keeping up with the opportunities while managing the accompanying vulnerabilities and risks requires new ways of thinking. Are our leaders prepared for this challenge?

Appendix 1: Survey design

This appendix contains the questions used in the survey for this study. The survey was administered using Qualtrics. Respondents took between 5-10 minutes to respond to these questions.

Q1: Please provide a few details to help us understand the context of your responses:

Q1a: - In what industry do you work (e.g. financial services, healthcare, high tech, education, media, transportation, etc.)?

Q1b: - What is your primary role (Board member, CISO/CTO/CIO, business executive, etc)?

Q2: What is the biggest BUSINESS risk you see for adopting AI in organizations?

Q3: What is the biggest SECURITY risk you see for adopting AI in organizations?

Q4: How concerned are you about each of these areas of AI risk? Please rate each risk below on a scale from 1 - 5; where 1 - No Concern and 5 - Strong Concern

Q4a: - Data & Privacy: Protect user and enterprise data from AI leaks and unauthorized AI system access or usage

Q4b: - Safety: Prevent AI systems from generating harmful outputs through internal and external evaluations, monitoring, and safeguards

Q4c: - Security: Protect AI systems from adversarial attacks, data poisoning, and unauthorized access

Q4d: - Reliability: Prevent hallucinations and system access that could cause harm

Q4e- Accountability: Enforce oversight of AI systems and develop emergency response protocols to mitigate issues

Q4f- Society: Prevent AI from enabling societal harm through cyberattacks or national security risks

Q5: What other risks are missing from this list?

Bibliography

Bello y Villarino, J. M., & Bronitt, S. (2024). AI-driven corporate governance: a regulatory perspective. Griffith Law Review, 33(4), 355–374. https://doi.org/10.1080/10383441.2024.2405752

Choudhary, S., & Kar, A. K. (2025). Modeling the role of generative AI in organizational privacy and security. Decision Support Systems, 196, 114500. https://doi.org/10.1016/j.dss.2025.114500

Chowdhury, S., Budhwar, P., Dey, P. K., Joel-Edgar, S., & Abadie, A. (2024). Generative artificial intelligence in business: Towards a strategic human resource management framework. British Journal of Management, Volume 35, Issue 4, 1680-1691. https://doi.org/10.1111/1467-8551.12824

Humphreys, D., Koay, A., Desmond, D., & Mealy, E. (2024). AI hype as a cyber security risk: The moral responsibility of implementing generative AI in business. AI and Ethics, 4, 791–804. https://doi.org/10.1007/s43681-024-00443-4

Identity Theft Resource Center. (2025, February 21). Weekly breach breakdown: Weak password, McDonald's data breach. https://www.idtheftcenter.org/podcast/weekly-breach-breakdown-weak-password-mcdonalds-data-breach/

Jada, Irshaad & Mayayise, Thembekile. (2024). The impact of artificial intelligence on organisational cyber security: An outcome of a systematic literature review. Data and Information Management. Data and Information Management, Volume 8, Issue 2,. https://doi.org/10.1016/j.dim.2023.100063

Neto, N. N., & Pearlson, K. (2025). Understanding the cyber risks of artificial intelligence: An ongoing, comprehensive, multi-faceted approach for CIOs, CTOs, and CSOs. SSRN. https://doi.org/10.2139/ssrn.5251525

Pearlson, K. (2025). AI-proofing the board and the c-suite: Managing AI business and security risk. MIT Sloan School of Management. https://mitmgmtfaculty.mit.edu/kpearlson/wp-content/uploads/sites/178/2025/06/AI-Proofing-the-BoardCSuite-Pearlson-Position-Paper.pdf

Wagman, K. B. (2025). Generative AI uses and risks for knowledge workers in a scientific-organization field study. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25), Article 1199, 1–17. https://doi.org/10.1145/3706598.3713827

Weaver, C. (2025, February 13). Microsoft Copilot flaw enables zero-click attack. Cybersecurity Dive. https://www.cybersecuritydive.com/news/flaw-microsoft-copilot-zero-click-attack/750456/