**Written Testimony of Dr. Christian Lau**
Co-Founder and President, Dynamo AI

**Digital Assets, Financial Technology, and Artificial Intelligence Subcommittee of U.S. House of Representatives Financial Services Committee**

**Hearing on "Unlocking the Next Generation of AI in the U.S. Financial System for Consumers, Businesses, and Competitiveness"**

**September 18, 2025**

Chairman Steil, Ranking Member Lynch, members of the Subcommittee, and staff, I thank you for inviting me to testify today and am honored to participate in discussions focused on advancing AI across the financial services ecosystem.

## Introduction

We are at a pivotal moment in the history and development of technology. Artificial intelligence (AI) stands to reshape so much across our economy and daily lives, and this is particularly true within the financial services sector. With a vibrant and active market of players driving much of this change, it is an honor to be able to contribute to this conversation on behalf of Dynamo AI and the thriving ecosystem of AI and security start-ups in America.

In 2021, I founded Dynamo AI alongside my co-founder and CEO Vaikkunth Mugunthan during our PhDs at MIT. Our mission has always been to help enterprises navigate complex regulatory environments, particularly where compliance requirements pose open technology challenges that institutions struggle to solve. We found that, when faced with new technology regulations or internal compliance requirements around new technologies, enterprises are often left paralyzed, asking themselves not only "how do I comply with this new requirement?" but also "is it even *technically possible* for us to comply with this new requirement?"

Nowhere did we see this to be more prevalent than with the struggles of financial institutions striving to adopt AI. Since founding Dynamo AI, we've had the opportunity to work with some of the largest global financial institutions, the most cutting edge fintech companies, as well as regional banks across America to help them navigate compliance and governance challenges posed by AI. Dynamo AI itself is backed by over 40 of the top 100 US financial institutions and a consortium of community banks who often lean on Dynamo AI to navigate their governance of AI and securely deliver AI applications into production.

Every day, our team witnesses exciting new AI proof of concepts (POCs) of financial services providers that promise to transform the customer experience, enable more efficient and comprehensive compliance, or enable more data driven decisioning. But for nearly every exciting AI POC we encounter, we also see another AI POC fail to make it into production and deliver value. We commonly see these POCs fail not because the underlying AI technology can't deliver, but rather because financial institutions struggle to answer open questions about managing AI risk in heavily regulated environments. We routinely encounter technology teams at financial institutions paralyzed by questions surrounding AI risks from legal, compliance, and security stakeholders: How can institutions teams manage risks of giving every banking associate an AI assistant that can be prompted in infinitely different ways? How can such AI assistants encourage rather than mislead employees in following bank protocols and procedures? What happens when an AI assistant inevitably hallucinates, fabricating facts in its response?

While financial institutions often struggle to respond to these concerns, these are not intractable problems. At Dynamo, we've worked with a multitude of financial institutions to establish effective AI risk management that accelerates, rather than blocks AI transformation. This often first involves institution-wide education about AI, including AI's capabilities, function, and risks, followed by an initiative to align diverse risk stakeholders across the financial institution around a cross-functional governance framework. To layer in the necessary technical controls, financial institutions must also embrace technology solutions that can help risk teams scale.

Our Dynamo team spends just as much time educating legal, risk, compliance, security, technology teams, as well as regulators about best practices in AI governance as we spend on implementing technical controls around AI risks, including controls like AI guardrails, red-teaming evaluations, and observability over AI usage. For example, to date, our AI guardrails service checks over 1 million user interactions every day for security vulnerabilities and noncompliance with bank policies, giving nontechnical risk stakeholder auditability into AI applications across the bank. We're starting to see comprehensive AI risk management take shape, which we believe will be key to ushering in this exciting new era of advancement for financial services.

## Advancement of AI across Financial Services

### *A Proliferation of Economic Advancement and Business Opportunity*

Working across the financial services sector, our team has gained unique insight into AI opportunities through two key experiences: leading an AI company backed by a consortium of financial institutions—from credit unions to regional banks and systematically important financial institutions—and helping financial services organizations enable AI use cases within this highly regulated sector.

To date, financial institutions have looked to AI primarily to enhance productivity to derive return on investments, increase worker and operational efficiency, and learn how to establish effective AI governance foundations. This takes many forms, including AI-powered chatbots that help employees understand company policies and business line standards to better execute processes and engage with customers and colleagues. This also is evident in the rapid development of AI used to generate code for technology systems. I expect AI to further extend across the financial services value chain, including financial product operations, investment analysis, and customer experience over the near term. Additionally, financial regulators themselves are starting to look to AI as a tool to support market oversight and monitoring of potential financial misconduct in the marketplace.

Many institutions have identified hundreds of AI use cases in their internal queues, ready for exploration, assessment, and deployment. This is a result of the ingenuity and excitement of financial services personnel from all lines of business and operations, looking to deploy AI to strengthen their organizations and propel customer value.

The financial services vendor ecosystem is also rapidly introducing AI to enhance core technology and operational processes, from compliance management to procurement. This includes the introduction of AI capabilities within existing third party applications, such as a sales or document processing software that an institution may already use. While this opens up opportunity and efficiencies in existing embedded processes, it also introduces additional institutional risks that we will discuss further.

As this subcommittee learns about AI's benefits and risks in financial services, it is crucial to recognize how community and regional banks can leverage AI to better serve consumers and strengthen local economies. Some of the most forward-thinking community and regional banks I interact with view AI as a transformational, once in a generation opportunity to deepen customer relationships, enhance community engagement, prevent fraud, and achieve operational efficiencies that directly benefit consumers.

AI's ability to personalize services at scale enables community and regional banks to compete more effectively with larger institutions, while delivering tailored financial solutions that expand consumer access and financial wellbeing. This technology can help smaller banks offer sophisticated services previously available only at major institutions, creating a more competitive and consumer-friendly banking landscape.

With proper risk mitigation, AI deployment in financial services can deliver significant advantages for consumers across all institution sizes—from enhanced fraud protection and more scalable customer service to broader access to credit and banking services in underserved communities.

## Landscape of AI Risks for Financial Services Institutions and Regulators

I believe the financial services industry is in the midst of a pivotal moment in history with respect to risk management. Alongside the development and deployment of initial AI use cases across the financial services landscape, organizations are learning how to operationally govern AI use and integrate these activities into existing risk management frameworks. I see this in the cross-functional AI governance working groups that are being established, the policies being written and deployed, the training that is expanding across employees, and the active engagement with regulators, trade associations, and the marketplace. More broadly, financial services risk management practices often become the standard that other industries adopt. Therefore, it is critical that policymakers and regulators clearly signal their priorities—both the innovation they want to encourage and the risks they consider most important to address. This regulatory tone will shape how financial institutions design their AI governance frameworks today and how these systems operate for years to come.

There are risks that are unique to AI, as well as risks that are unique to AI and financial services. For the subcommittee, I will provide a brief overview of a number of these AI risks, of which I can answer questions about, but will also provide deeper insights into five risks that I believe are critical to mitigate to advance a vibrant, secure financial services ecosystem.

### *Key Cross-Sector AI Risks*

Generative AI introduces several unique risks that organizations often struggle to mitigate and monitor, which in turn can delay effective use and integration of the technology. These risks include:

- The risks of **hallucination**, a scenario in which AI systems generate false, misleading, or fabricated information, while presenting it as factual or accurate. This risk can degrade trust and lead to harmful or misinformed decision-making by whoever makes use of AI outputs.
- The risk of **adversarial security attacks**, including prompt injections that manipulate AI responses through malicious inputs and jailbreaking techniques that bypass AI guardrails, which can exploit model vulnerabilities and circumvent defenses at scale, require AI-specific threat evaluations and controls.
- **Data and Access Risks** including vulnerabilities where third party AI providers may consume volumes of enterprise data fed into models, resulting in the potential leakage of sensitive enterprise or client data to third parties.
- The risk of **misuse of AI systems** for unauthorized, high-impact use-cases, such as writing legal documents, obtaining unwarranted investment or healthcare advice, or making creditworthiness assessments without proper checks.

### *Key AI Risks Impeding Financial Services Innovation*

The financial services industry is home to a suite of rigorous risk management practices that aim to protect consumers, introduce effective challenge, independent thought, and advance financial opportunity for our citizens and the marketplace. The worthy challenge of introducing AI in the financial services ecosystem is to balance risk management expectations, while also promoting innovation for consumers, vendors, and our marketplace. Therefore, there are a few key risks I believe require attention from the subcommittee.

One risk is rooted in the **explainability** of a model's response. Simply put, it is impossible to determine exactly how or why a Generative AI model responded the way it did. This, in itself, makes Generative AI models different from conventional models used in financial services today, such as statistical or probabilistic models, for which a decision can be traced back to a specific component of the model. The reality of explainability and lack of traceability in AI models may warrant a revision of existing risk management guidance, evaluation methods, and subsequent controls.

An additional noteworthy risk is that of protecting **model sovereignty**, which refers to an organization's control over the AI models they deploy, including the data, infrastructure, and capabilities needed to maintain independence from external providers and ensure regulatory compliance. This is of particular importance for financial institutions and regulators to align on, as models are often developed and managed by third parties that may service other institutions and/or critical technology components of the same institution. Importantly, this risk is only amplified if US financial institutions utilize models that are developed or trained in nations with adversarial relationships with the United States, including China.  A discussion focused on acceptable third-party risk management protocols is a critical steppingstone to effective internal and external oversight.

A strength of the financial services regulatory system is that institutions can interpret principles-based rules and turn them into practical compliance processes—even deciding when to accept certain risks. With AI now entering processes once handled by people, the third major risk is **whether institutions can keep AI aligned with their own definitions, policies, and risk tolerance—their "ground truth."** AI tools often come with assumptions from vendors or global bodies, but it is imperative that financial institutions have the ability and controls to maintain autonomy over their "ground truth." For example, a regional bank may have a specific definition of "legal advice" and strict rules on when it can or cannot be given. An AI system might not share that definition. Without guardrails that let organizations overlay their own rules on third-party AI models, compliance gaps open and competitiveness suffers.

As alluded to earlier, AI is being actively enabled within a host of existing embedded technologies and processes already in place across financial institutions. For example, many financial institutions have human resource software or document processing software that may have AI components enabled, including document analysis and chatbots. However, market or organizational dependence on a single AI provider or tightly coupled ecosystem increases vulnerability to operational risk or systemic failure. The result is an increase in **third- and fourth-party risks,** in addition to risk surrounding **vendor concentration**. Encouraging diversity in infrastructure and support can improve market resilience, financial stability, and reduce consumer harm.

Finally, it is also worth briefly discussing the newest frontier of AI and AI Risk: **Agentic AI**. In short, AI agents are AI systems that can independently execute complex tasks, make decisions, and take actions on behalf of users or organizations. However, even simple AI agents require organizations to access sensitive information and make impactful decisions. And the value derived from replacing manual work with AI reduces human control and oversight over AI decision-making. As a result, Agentic AI amplifies many of the aforementioned risks. The more powerful the AI agent, the more risk there is in its deployment, as this requires organizations to transfer more decision-making authority from humans to AI systems. Therefore, for AI agents to truly provide value, organizations will need to mobilize novel technologies to rigorously test, sandbox, and embed safeguards into these tools.

Dynamo AI’s AI Compliance Strategy team has documented an inventory of key AI risks for the subcommittee, including key considerations and strategic implications for policymakers and industry leaders alike. For more, see Appendix A.

**AI as an Enabler for Risk Management and Compliance**

Despite the risks presented by AI, one of the most noteworthy benefits of this powerful technology is that **AI also serves as a critical protective function, or control, to mitigate many key AI and security risks** and an **effective enabler for organizations.**

Dynamo is proud to be a leader in this regard, developing technologies that product and risk management teams more comprehensively evaluate and guardrail both AI-specific threats, as well as compliance and security challenges. Within this context, we have seen novel innovations in the AI trust, security, and risk management ecosystem directly enable organizations to accelerate the deployment of AI applications, while maintaining compliance with complex regulatory requirements.

At Dynamo, our product suite uses AI to provide streamlined tests, evaluations, and real-time protections for AI applications launched in the financial services ecosystem.

*AI Used to Power Simulated Security Attacks and Evaluations*

With AI systems, the threat landscape of adversarial attacks is both constant and expansive. It is virtually impossible—and economically infeasible—for organizations to staff departments with sufficient personnel to effectively predict and block all possible adversarial threats that may target an AI system. Moreover, new attack methods emerge daily, requiring protection methods to be continuously updated and personnel to conduct ongoing alert review and escalation.
In my work at Dynamo, our team has **integrated AI into our test and evaluation suite**, using AI to attack existing models to identify threats proactively for human review and decision-making. In other words, the solution to many key challenges in this space lies in leveraging AI itself as a defensive tool.

*Real-time AI Guardrails*

A core challenge for the financial services marketplace is how to deploy AI while complying with existing laws and regulations. Dynamo has worked with institutions to develop what are known as **AI “guardrails”** – a specific category of technology we have helped to pioneer that can moderate how a model behaves, what data is permissible to enter or leave an AI system, and to do this in alignment with each institution's policies and definitions. In a highly regulated sector where compliance with regulatory guidance and requirements must be embedded into every process, AI guardrails enable financial institutions to scale their compliance interpretations and best practices and enable high-value AI use cases. I believe this to be one of the primary challenges of AI adoption across the financial services ecosystem; but, once guardrails expand across the ecosystem, organizations can fully realize the value and returns of these technologies, alongside a more sound and secure banking ecosystem.

*Innovative Tools for AI Monitoring*

A key component of effective AI risk management is to institute comprehensive human-in-the-loop observability to give humans the ability to monitor models, interactions, and demonstrate compliance with internal policy and regulations. Observability equips internal risk management and audit teams, as well as regulators with evidence needed to substantiate a level of operational assurance that can be measured and tested. At Dynamo, we've built a full observability suite to allow organizations to continually observe all internal and customer-facing AI interactions, so they can strengthen controls as AI is in use. While a complex technical solution in its own right, AI observability provides organizations with the necessary reporting and alerts for internal monitoring of powerful systems, and it will prove to be an essential ingredient for effective AI oversight in the sector in the near future.

These AI-powered approaches represent a necessary evolution in risk management. Together, they suggest that sustainable AI adoption in financial services will depend on institutions' ability to implement technology-based controls that can operate at the speed and scale of the systems they oversee.

## Key Considerations for Fostering a Competitive, Secure Ecosystem for AI

As the Subcommittee weighs policy and technical considerations for the continued promotion of AI innovation and a vibrant and fair financial services ecosystem, there are a few actions I would like to highlight.

*Sandboxes as a Vital Tool for Innovative Oversight and Information-Sharing*

Regulators, with the participation of financial services institutions, should continue to establish and utilize AI sandbox environments, a concept referenced in the current administration's recent AI Action Plan. These environments, either established by a financial regulatory and/or market consortia, allow parties on both sides of financial markets to explore AI use cases, risks, as well as acceptable controls to mitigate those risks. They also allow innovative American companies such as ours to participate easily alongside the broader vendor technology landscape and showcase advancements. In turn, sandboxes educate a wide variety of stakeholders across regulatory agencies and leaders within financial services institutions.

We applaud Committee Chairman Hill, Representative Torres, leaders of this Subcommittee, and colleagues in the Senate for introducing H.R. 4801, the Unleashing AI Innovation in Financial Services Act. This Bill strikes a strong balance between supporting innovation, fostering governance within enterprises, and educating regulators on emerging AI use cases and risks, allowing companies to experiment with the emerging technology and conduct necessary tests to assess key risks.

Governments across the globe, including in Singapore and the United Kingdom, have used sandboxes to drive technology forward and inform future policy decisions. We would welcome the opportunity to share best practices and insights from our current work in AI sandboxes.

### *Growing the AI Evaluation Ecosystem*

Continue to support a vibrant ecosystem of AI providers, including within the independent AI evaluation and guardrail market, whether through guidance, regulation, or support for effective challenge risk management methodologies. The administration's AI Action Plan underscores the importance of growing this ecosystem, and it is critical that this community includes not only federal agencies but also companies that are incentivized to accurately identify and protect against existing and emerging AI risks. We have seen similar requirements for independent evaluation in the areas of financial crimes and fraud, and I believe similar expectations may be warranted when applying AI security and compliance controls. As the AI industry continues to grow and new players emerge, it is essential to ensure that the fox is not guarding the henhouse.

Not only are AI evaluations important for driving secure AI adoption in the commercial space. They are also vital as federal agencies look to deploy AI. As mentioned in the AI Action Plan's section calling on agencies to "Accelerate AI Adoption in Government," the General Services Administration (GSA) is tasked with creating an "AI procurement toolbox" for federal agencies to easily acquire AI tools. As part of this agency-ready AI marketplace, we recommend the incorporation of adequate independent test and evaluation technologies into the federal procurement process for high-impact AI use cases, such that federal agencies can easily deploy AI in a compliant and secure way.

### *Considerations for the Future of Model Risk Management*

While mechanisms or applied controls for AI model evaluation may differ from historical model risk management, we are energized by ongoing discussions between private and public sector leaders on acceptable controls for AI use cases in financial services. We ask the Subcommittee to continue to promote an open dialogue between standard-setting institutions, financial regulators, and the financial services industry to arrive at best practices for mitigating risks found in common financial services use cases. This may, in turn, lead to new guidance that may be well received by the industry.

Finally, as the American AI industry continues to grow, we would like to underscore the importance of policy and regulatory initiatives that support the advancement of America's start-up ecosystem. This opportunity is a testament to the Subcommittee's commitment to supporting entrepreneurship, advancing America's technological competitiveness, and giving small businesses a platform. For that, I am very grateful.

## Conclusion

On behalf of the entire team at Dynamo AI, thank you again for the opportunity to testify. I welcome the opportunity to work with the Subcommittee to further create a competitive, secure, and compliant AI ecosystem within the financial services industry.

**Appendix A:** Key Frontier Findings: Current and Emerging AI Risks

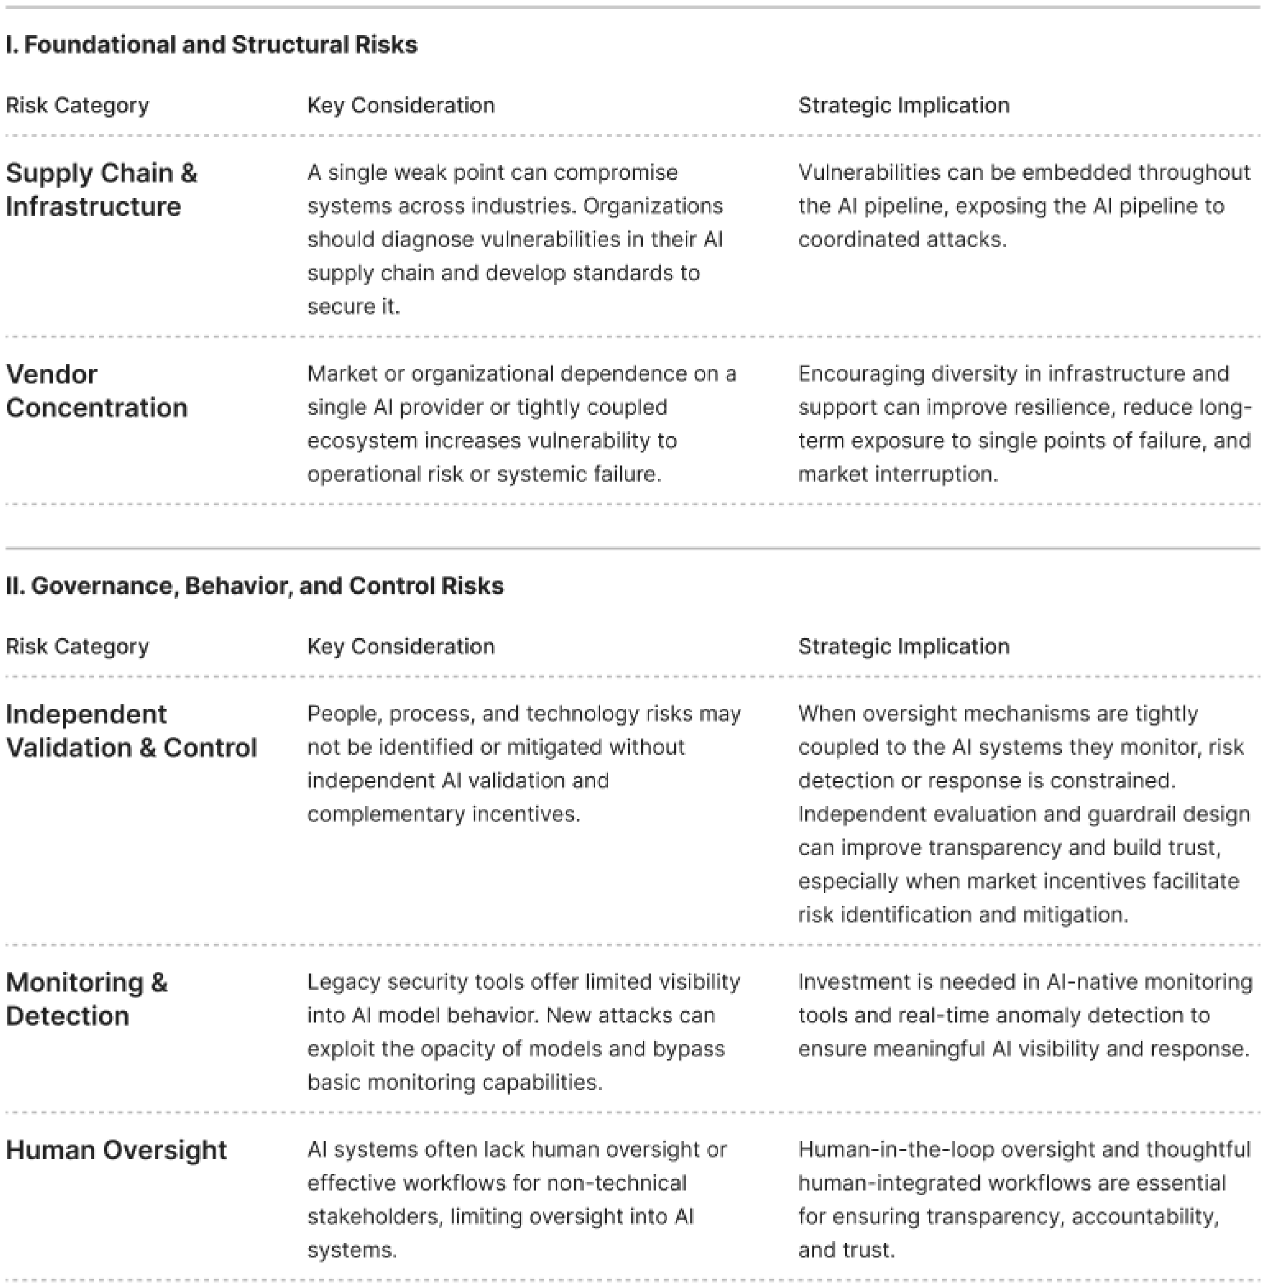

### I. Foundational and Structural Risks

| Risk Category | Key Consideration | Strategic Implication |
|---|---|---|
| **Supply Chain & Infrastructure** | A single weak point can compromise systems across industries. Organizations should diagnose vulnerabilities in their AI supply chain and develop standards to secure it. | Vulnerabilities can be embedded throughout the AI pipeline, exposing the AI pipeline to coordinated attacks. |
| **Vendor Concentration** | Market or organizational dependence on a single AI provider or tightly coupled ecosystem increases vulnerability to operational risk or systemic failure. | Encouraging diversity in infrastructure and support can improve resilience, reduce long-term exposure to single points of failure, and market interruption. |

### II. Governance, Behavior, and Control Risks

| Risk Category | Key Consideration | Strategic Implication |
|---|---|---|
| **Independent Validation & Control** | People, process, and technology risks may not be identified or mitigated without independent AI validation and complementary incentives. | When oversight mechanisms are tightly coupled to the AI systems they monitor, risk detection or response is constrained. Independent evaluation and guardrail design can improve transparency and build trust, especially when market incentives facilitate risk identification and mitigation. |
| **Monitoring & Detection** | Legacy security tools offer limited visibility into AI model behavior. New attacks can exploit the opacity of models and bypass basic monitoring capabilities. | Investment is needed in AI-native monitoring tools and real-time anomaly detection to ensure meaningful AI visibility and response. |
| **Human Oversight** | AI systems often lack human oversight or effective workflows for non-technical stakeholders, limiting oversight into AI systems. | Human-in-the-loop oversight and thoughtful human-integrated workflows are essential for ensuring transparency, accountability, and trust. |

**Appendix A Continued:** Key Frontier Findings: Current and Emerging AI Risks

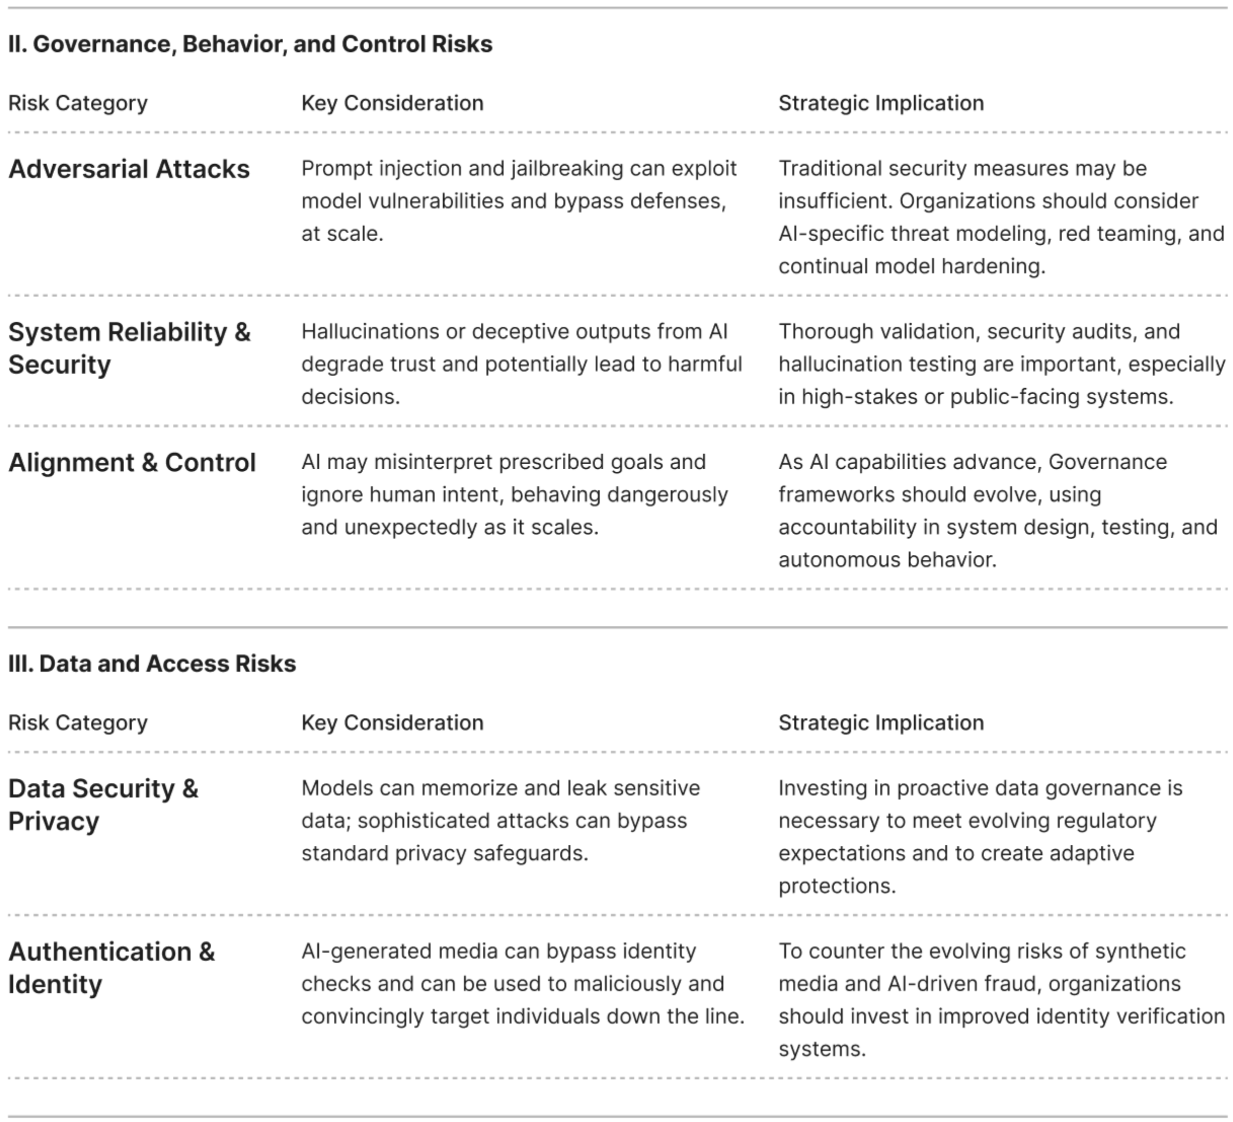

### II. Governance, Behavior, and Control Risks

| Risk Category | Key Consideration | Strategic Implication |
|---|---|---|
| **Adversarial Attacks** | Prompt injection and jailbreaking can exploit model vulnerabilities and bypass defenses, at scale. | Traditional security measures may be insufficient. Organizations should consider AI-specific threat modeling, red teaming, and continual model hardening. |
| **System Reliability & Security** | Hallucinations or deceptive outputs from AI degrade trust and potentially lead to harmful decisions. | Thorough validation, security audits, and hallucination testing are important, especially in high-stakes or public-facing systems. |
| **Alignment & Control** | AI may misinterpret prescribed goals and ignore human intent, behaving dangerously and unexpectedly as it scales. | As AI capabilities advance, Governance frameworks should evolve, using accountability in system design, testing, and autonomous behavior. |

### III. Data and Access Risks

| Risk Category | Key Consideration | Strategic Implication |
|---|---|---|
| **Data Security & Privacy** | Models can memorize and leak sensitive data; sophisticated attacks can bypass standard privacy safeguards. | Investing in proactive data governance is necessary to meet evolving regulatory expectations and to create adaptive protections. |
| **Authentication & Identity** | AI-generated media can bypass identity checks and can be used to maliciously and convincingly target individuals down the line. | To counter the evolving risks of synthetic media and AI-driven fraud, organizations should invest in improved identity verification systems. |