Testimony of David Cox
Vice President for Artificial Intelligence Models at IBM Research and Director of the MIT-IBM Watson AI Lab

Before the Senate Committee on Banking, Housing, and Urban Affairs
Subcommittee on Securities, Insurance, and Investment

Hearing on "Guardrails and Growth: AI's Role in Capital and Insurance Markets"
Wednesday, July 30th, 2025

Chairman Rounds, Ranking Member Warner, and distinguished members of the Subcommittee, thank you for the opportunity to appear before you today. My name is David Cox, and I serve as Vice President for Artificial Intelligence Models at IBM Research and Director of the MIT-IBM Watson AI Lab.

IBM has a long history of helping enterprises, including those in the financial sector, use artificial intelligence (AI) technologies to unlock business value, and doing so in ways that are responsible and engender trust. Our work spans a broad spectrum, from helping identify appropriate use cases, to providing tooling to govern both the development and deployment of AI systems, to inventing new technologies for trustworthy AI.

At IBM, we believe AI represents a generational opportunity, particularly in the financial services sector. From automating routine processes and enhancing compliance to powering customer service and investment research, AI can unlock untapped productivity, increase resilience, and deliver meaningful value to consumers. But as we capitalize on these advances, it is imperative that we proceed thoughtfully, with guardrails in place to manage risk, enhance transparency, and promote trust.

AI *can* and *should* be developed responsibly, and open innovation will be key to that future. We know this firsthand because at IBM we are not only an AI developer, we are also AI's first and most demanding client. We refer to this internally as "Client Zero." Every AI model, framework, and tool we create is tested in our own operations and with some of the world's most complex enterprise environments, before we deploy them with clients.

For example, by leveraging IBM's watsonx© AI and automation tools, we help to augment the skills of our own workforce by eliminating repetitive tasks. This has enabled our employees to focus on more challenging, rewarding, and impactful work with the time they save – which was an estimated 3.9 million hours in 2024 alone. Utilizing the financial transparency enabled by Apptio, we were able to pinpoint the total cost of our IT operations, allowing us to make better, more informed tradeoffs and increase productivity. This has contributed to IBM's efforts to drive approximately $600 million in enterprise IT cost savings since 2022. These are but two examples of the many benefits our "Client Zero" strategy has helped enable, and which in aggregate have helped IBM achieve productivity improvements of approximately $3.5 billion over the past two years, applied across more than 70 business arenas, including finance.

Experiences like these give us unique insight into both the promise and the perils of generative AI at scale.

Opportunities for Financial Services

For the financial industry, the implications of AI – particularly generative AI and large language models (LLMs) – are transformative. LLMs are not merely chatbots; they are multifunctional tools that can be adapted across a wide range of financial services use-cases, including:

- Regulatory compliance and audit trail generation;
- Summarization of lengthy prospectuses and client contracts;
- Faster, more personalized client service through natural language interfaces;
- Automated reporting, complaint resolution, and workflow optimization; and
- Enhanced developer productivity and IT resilience.

What makes LLMs so powerful is their adaptability. Unlike previous AI systems, they don't require extensive retraining for every use case. This means faster deployment, broader use, and a lower barrier to entry, even for smaller firms and institutions. This democratization of capability can benefit both firms and consumers through improved access, faster services, and lower costs.

LLMs are like an AI Swiss Army knife, useful for a wide variety of purposes. This is in contrast to previous generations of AI technology, which required substantial additional effort for each new use case. In previous generations of AI technology, fielding an AI solution would require assembling and curating large amounts of data for each specific use case, and undertaking painstaking and complex work to train AI models focused on each task. For each new task, the process would be repeated, usually from scratch, amounting to a fully recurring cost. Only those use cases for which enough data was available, and a high enough return on investment was expected, were good candidates for application of AI.

With LLMs, an application developer is now able to “stand on the shoulders of giants,” leveraging the massive amounts of data that were used to train the LLM. Then, with only a small amount of data to train with, or even no data – just a natural language “prompt” describing what you want the model to do – you can build sophisticated applications. Because the data and cost/time thresholds are so much lower, LLMs allow a much wider range of use cases to be targeted.

While called “LLM”, bigger isn’t always better. Financial institutions don’t need models that can solve physics problems; they need domain-specialized tools that are right-sized for the task. Smaller, more efficient models may not only reduce costs, but also can outperform larger general-purpose models in key financial applications where risk-based investment and fraud detection decisions must be made in a matter of microseconds.

Regardless of size, the conversational abilities of LLMs also enable a broader range of the workforce to interact with AI systems, further amplifying their impact. This is the power of innovation: it compounds over time.

Risks and Challenges

At the same time, as with any new technology, there are concerns we must address and keep in mind as we harness the technology’s power. We believe that many of the fundamental considerations remain constant, such as the need to focus risk assessments on use cases, rather than just technology per se. However, LLMs bring a number of new issues into focus.

One concern is their size, and the largest LLMs have been growing steadily over time. Larger models require more computational resources to both train and deploy, leading to greater power consumption and greater cost. The energy consumption of LLMs has significant potential to strain our infrastructure. In fact, by some estimates, year-over-year energy usage for computing is on track to exceed the planet’s worldwide energy budget by the year 2040. Cost of deployment also grows as the size of the model increases, and many leaders in the financial services sector with whom I have spoken are already beginning to express concerns about the cost sustainability of generative AI.

Another cause for growing concern is the level of transparency in how a given LLM was created. LLMs require enormous quantities of training data, and few LLM builders are fully forthcoming about the data used to train their models. This lack of transparency can pose problems for any enterprise but especially in regulated industries where companies and regulators have a legitimate need to know what data went into any AI system that they deploy.

Finally, many enterprises are concerned about the security of their own data, and the data of their customers – especially in contexts where LLMs are offered as a cloud service. A number of companies have banned the use of some as-a-service LLMs for fear that their workers might inadvertently lose control of sensitive data with a vendor that has not been properly vetted by their chief information security office.

But one of the biggest risks is that industry won't deploy AI fast enough to gain the early benefits of adoption – both for the economy and for consumers. In order to help expedite AI adoption, responsible AI governance is key.

Responsible AI Governance: Open, Trusted, and Secure

As the financial industry knows better than most, trust takes years to build and seconds to lose. In regulated environments, firms must understand exactly what data underpins their models and be able to audit those systems over time.

That's why IBM's enterprise AI efforts are grounded in three principles: Open, Trusted, and Targeted.

Open – We strongly believe in the value of open source AI. As demonstrated by the rise of Linux, which powers much of the modern Internet, open technology fosters security, community participation, and innovation. Open models better enable regulators, academics, and enterprises to inspect, improve, and adapt models to specific needs. Open source also reduces reliance on a handful of opaque, proprietary providers – an especially important consideration for national security, innovation competitiveness, and economic resilience.

In fact, open source AI is one of our strongest levers for managing the very risks that concern this Committee: transparency, auditability, security, and cost control. IBM's contributions to the open-source LLM ecosystem – along with our commitment to open data governance frameworks – are grounded in this philosophy.

Open innovation also benefits the broader AI ecosystem, driving economic growth, enhancing security, and strengthening democratic values. It enhances security, trust, and collaboration through transparency, enables smaller firms and research organizations to compete without prohibitive upfront capital investments, and expands the pipeline of the future talent that our country will undoubtedly need to ensure we remain competitive in the ongoing global AI race. Open source accomplishes all of this by crowdsourcing the collective wisdom of all willing participants, allowing everyone to play a role in building the AI future from which we all stand to benefit.

That future should be controlled by the many, not the few.

Trusted – Enterprises deserve transparency. Curating data that goes into an LLM for use by an enterprise is an important undertaking, and one that should not be taken likely. At IBM, we take data curation seriously, applying AI-based content filters to help remove objectionable material, along with rigorous filtering procedures and blocklists designed to avoid problematic data finding its way into models. We also make substantial efforts to proactively *add* trusted data sources, for instance, in the domain of finance, to ensure that models are trained with high quality, authoritative information.

Transparency is a cornerstone of trust, so we disclose details of how we train our flagship Granite models and what data goes into them, and we have open-sourced the software framework we created to curate our data. We have implemented a custom data management application to maintain full lineage between models and the data that went into them in order to enable auditability of all of the inputs into each model we produce, and we have integrated this application directly with IBM's data clearance process. Stanford University researchers produce an annual report with a transparency index covering a variety of different dimensions of transparency, and IBM's models have emerged at the top of the rankings among peer model providers.

The institutions that use models should know what's in them – especially in industries that deal in sensitive information.

Secure – It is crucial to follow best practices in securing AI at each stage of the AI pipeline, including during data collection and handling, model development and training, and model inference and use. Organizations should apply a risk-based approach to secure the data, the model and the model's usage, and the infrastructure on which the AI models are being built and run. Finally, they need to establish AI governance and monitor for drift over time.

IBM's framework emphasizes that securing each layer (e.g., data, model, and infrastructure) requires a combination of access controls, encryption, anomaly detection, and tailored machine learning detection and response capabilities to help mitigate emerging threats across the AI lifecycle.

IBM has implemented an integrated governance program (IGP) that moves from a more reactive to a continuous compliance model for data, privacy, security, and AI. This continuous compliance approach empowers us to bring AI products and services to market with speed and trust, for ourselves and for our clients. IGP is underpinned by IBM's own technology, serving as a living lab for our trusted solutions.

Policy Recommendations: Balancing Innovation with Guardrails

To promote an environment that appropriately balances the value of innovation against the need for reasonable guardrails, policymakers should prioritize legislation and policies that emphasize innovation and open technology. A healthy, competitive AI market ecosystem requires open source AI to ensure America's AI future is controlled by the many, not only a select few. As the recent AI Action Plan correctly identified: "We need to ensure America has leading open models founded on American values. Opensource and open-weight models could become global standards in some areas of business and in academic research worldwide. For that reason, they also have geostrategic value. While the decision of whether and how to release an open or closed model is fundamentally up to the developer, the Federal government should create a supportive environment for open models."

Additionally, policies should also avoid unduly burdensome regulation, and policymakers should instead focus on leveraging existing regulatory capabilities to address specific AI use-cases. Where clear, risk-based guardrails are needed, they should be tailored to the roles and capabilities individual organizations play in the broader AI developmental lifecycle.

In order to help cultivate that environment, we recommend:

- **Support for open AI ecosystems**: Encourage open source contributions and academic-industry-government collaborations. These partnerships drive safety, transparency, and innovation – just as they did during the internet's rise.

- **Use-case-based risk frameworks**: Regulate the *application* of AI, not just the technology. An AI model used to auto-summarize SEC filings should be governed differently than one used for autonomous financial trading.

- **Transparency requirements**: Enterprises should be able to apply best practices to answer: *What data was used to train this model? What decisions did it influence? What safeguards were in place*?

Conclusion

Generative AI is here. The question is not whether we use it, but how. Will we treat AI like a black box, or will we open the lid and shape it to serve our institutions, our consumers, and our values?

At IBM, we believe the financial industry has a once-in-a-generation opportunity to lead: to show that AI can be powerful and principled, fast-moving and trustworthy. Through thoughtful deployment, responsible governance, and, most importantly, a commitment to open technology, we can make that vision real.

Thank you, and I look forward to your questions.