Testimony of Dr. David Cox
Before the HOUSE FINANCIAL SERVICES COMMITTEE SUBCOMMITTEE ON DIGITAL ASSETS, FINANCIAL TECHNOLOGY, and ARTIFICIAL INTELLIGENCE

Hearing on "Unlocking the Next Generation of AI in the U.S. Financial System for Consumers, Businesses, and Competitiveness"
Thursday, September 18th, 2025

Chairman Steil, Ranking Member Lynch, and distinguished members of the Committee, 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.

A Brief History of AI: From Rules to Watson

For decades, artificial intelligence has captured the human imagination. Ever since the term was first coined by computer scientist John McCarthy in a 1955 proposal for the 1956 Dartmouth Summer Research Project on Artificial Intelligence, IBM has played an integral role in this technology's evolution. Indeed, it was an IBM researcher, Nathaniel Rochester, who co-authored the Dartmouth proposal and co-hosted the historic conference. But it wasn't until the 1970s, however, that much of the technology's practical evolution began unfolding.

Early AI systems were largely rules-based "expert systems." These programs relied on manually encoded knowledge, painstakingly crafted by engineers, to make decisions. While useful in narrow contexts — like medical diagnosis support or industrial automation — these systems were brittle, unable to adapt beyond their programmed rules.

By the 1980s and 1990s, machine learning emerged. Instead of encoding all the rules by hand, researchers trained systems using statistical methods and data. This shift gave rise to algorithms for fraud detection, credit scoring, and basic speech recognition — tools that began to find homes in finance, healthcare, and telecommunications. It was during this time, in 1998, that IBM's "Deep Blue" made headlines by defeating Garry Kasparov, the world chess champion. But chess was only a milestone. It demonstrated how computing power, when combined with intelligent design, could solve problems long thought unsolvable.

The 2000s ushered in the era of scale. With exponential increases in computational power and data availability, machine learning accelerated dramatically. In 2011, IBM Watson became a household name by winning *Jeopardy!*, defeating two of the game's greatest champions. This was not just a media spectacle; it was proof that computers could understand natural language, navigate ambiguity, and reason across vast stores of information. Watson symbolized the leap from AI as an experimental tool to AI as a practical enterprise technology.

Watson's success laid the foundation for IBM's modern AI efforts. What began as a demonstration on a television stage evolved into a full suite of enterprise-ready tools, now unified under watsonx©, which combines foundation models, generative AI, and trusted governance.

IBM's Watson Legacy and the "Client Zero" Approach

From the very beginning, IBM has treated Watson not as a product in isolation, but as a living ecosystem. We have applied Watson technologies inside IBM's own operations, making ourselves the first and most demanding client — what we call "Client Zero."
AI tools, models, and governance frameworks are tested in IBM's own systems before reaching clients. This philosophy means that when we speak to our clients in banking, insurance, and capital markets about AI, we speak from firsthand experience, not just theoretical rumination.

For example, in 2024 alone, AI-powered automation tools deployed internally saved IBM employees an estimated 3.9 million hours of routine work. That time was reinvested into higher-value activities, from research and product development to client services. By using AI-powered financial transparency tools, IBM also identified cost-saving opportunities that contributed to nearly $600 million in IT savings since 2022**.** By 2025, our Client Zero initiatives are on track to generate approximately $4.5 billion in productivity improvements, spread across more than 70 business functions, including finance.

This internal proving ground allows us to deliver AI solutions to clients with confidence: if they can scale inside IBM, they can scale anywhere.

Partnerships and Open Innovation

IBM has also recognized that no one company can solve AI's challenges alone. That is why our AI story is also a story of partnerships, with academia, others in industry, and the larger open source community.

Academia. Through the **MIT-IBM Watson AI Lab**, of which I am the director, we collaborate with leading researchers to advance fundamental AI science. This lab is one of the largest university-industry AI research collaborations in the world, and it has produced breakthroughs in natural language processing, trustworthy AI, and domain-specific model development.

Industry. We partner with financial institutions to co-develop AI models optimized for compliance, customer engagement, and fraud detection. By embedding client expertise into model training, we ensure that Watson is not a "one-size-fits-all" solution, but a toolkit tailored to highly regulated, high-stakes environments.

Open Source. IBM has been a champion of open innovation**.** Just as Linux became a backbone of modern computing, open source AI models can become a foundation for trustworthy, transparent, and secure enterprise AI. By contributing to open source communities, IBM ensures that regulators, academics, and startups have the same opportunity to inspect, adapt, and improve upon foundational models.

These partnerships reflect IBM's longstanding ethos: that AI should be a collaborative, democratizing force, not a black box controlled by a small handful of providers.

Financial Services: Watson in Action

Nowhere is IBM's AI impact clearer than in financial services — an industry that thrives on data, trust, and efficiency. IBM and our partners are applying Watson-powered AI solutions across a wide array of domains and applications.

Regulatory Compliance and Auditability. In the regulatory compliance and auditability realm, AI can automatically generate and maintain auditable trails of financial activities, helping compliance officers navigate complex and evolving regulatory requirements. Long prospectuses and contracts can be summarized in minutes, giving decision-makers clarity without sacrificing detail.

Customer Service and Personalization. Financial firms are deploying IBM-enabled chatbots and natural language systems to deliver faster, more personalized customer service. Clients can interact with their institutions using non-technical language, resolving routine queries instantly while freeing human representatives for more complex cases.

Fraud Detection and Risk Management. AI helps financial institutions detect anomalies in real time, flagging fraudulent activity before it spreads. By analyzing vast amounts of transaction data, domain-specific models can make credit, investment, and risk-based decisions in microseconds.

Investment Research. AI tools are accelerating investment research by scanning global news, financial statements, and regulatory filings, producing concise and actionable insights. This not only saves time but levels the playing field by giving smaller firms access to capabilities once reserved for the largest institutions.

IT Resilience and Developer Productivity. Financial institutions use watsonx to optimize IT operations, identify cost inefficiencies, and boost developer productivity. This leads to more resilient infrastructure, faster innovation cycles, and improved consumer trust in the reliability of financial systems.

Each of these applications builds on IBM's central promise: to augment human expertise, not replace it. By working alongside financial professionals, AI enhances decision-making, accelerates workflows, and deepens trust with customers.

Risks and Challenges

As with any transformative technology, there are challenges we must anticipate and manage even as we take advantage of AI's opportunities. Many of the core principles remain unchanged, such as the importance of evaluating risks in relation to specific use cases rather than the technology alone. Yet large language models (LLMs) introduce new dimensions that demand careful attention.

One area of concern is scale. LLMs have been growing rapidly in size, and larger models require significantly more computational resources to train and run. This drives up both costs and power consumption. Energy demand in particular could put real pressure on infrastructure. Some projections suggest that by 2040, the global computing sector's energy use may surpass the world's entire energy budget.

Another pressing issue is transparency. Training LLMs requires massive amounts of data, but opacity can create serious challenges for enterprises, especially those operating under strict regulatory oversight, where both companies and regulators need confidence in the provenance and quality of the data underlying deployed systems.

Security is also top of mind. Organizations worry about safeguarding proprietary and customer data, particularly when using LLMs delivered through cloud-based services. Some firms have prohibited employees from using certain as-a-service models out of concern that sensitive information could be exposed to third-party vendors without sufficient security guarantees.

Perhaps the most consequential risk, however, lies not in overuse but in underuse. If industry hesitates too long to deploy AI, the broader economy and consumers may miss out on the early advantages of adoption. The path forward requires accelerating adoption while ensuring strong governance frameworks. Responsible AI is not a brake on innovation, but a catalyst for realizing its benefits securely and sustainably.

Responsible AI: Open, Trusted, Secure

The financial sector understands better than most that trust is painstaking to earn yet can vanish in an instant. In highly regulated industries, firms should have confidence in their models and maintain the ability to verify those systems over time. This is why IBM's enterprise AI strategy rests on three core principles: **Open, Trusted, and Secure.**

Open. We are firm believers in the importance of open source AI. Just as Linux became a backbone of the modern digital economy, open technologies drive security, collaboration, and innovation. Open models give regulators, researchers, and businesses the ability to examine, refine, and tailor AI to their particular needs. They also reduce dependency on a small group of closed, proprietary vendors — a key factor for safeguarding national security, competitiveness, and long-term resilience.

Open source AI is not a liability but one of the strongest tools we have to manage risk. It advances transparency, auditability, security, and cost-efficiency. IBM contributes actively to the open source large language model ecosystem and supports open data governance practices with this philosophy in mind.

Moreover, open innovation strengthens the entire AI landscape. It promotes economic dynamism, bolsters security, and aligns with democratic values by making AI development more transparent and collaborative. It enables smaller firms and academic institutions to participate without prohibitive upfront costs and helps cultivate the workforce America will need to stay ahead in the global AI race. By pooling collective expertise through open contribution, we can ensure the AI future benefits everyone — not just a select few.

Trusted. Transparency is the foundation of trust. Preparing the data that powers enterprise-grade language models is a serious responsibility. IBM applies AI-based content filters, extensive blocklists, and curated datasets to reduce the likelihood of harmful material contaminating our models. We also deliberately add authoritative sources to ensure our models are built on high-quality, reliable data.

We make the training of our Granite models transparent, disclosing the data processes involved and releasing the frameworks we use to curate inputs. Our custom data management system maintains full lineage between models and their training data. This system ties directly into IBM's internal clearance process. Stanford University's annual transparency index ranks IBM's models among the most open of any peer provider.

Users of AI systems — especially those handling sensitive information — deserve to know what is inside the models they rely on.

Secure. Security must be embedded throughout the AI lifecycle — from data gathering and preprocessing, to model training and development, to deployment and usage. A robust, risk-based approach secures not only the data and the models, but also the infrastructure they depend upon. This includes access controls, encryption, anomaly detection, and machine learning detection and response (MLDR) capabilities designed to defend against evolving threats.

IBM's integrated governance program (IGP) puts this into practice. It shifts away from reactive compliance toward continuous oversight across data, privacy, security, and AI systems. By embedding compliance into day-to-day operations, IBM is able to deliver AI solutions faster while maintaining the highest levels of trust and accountability. The IGP itself operates on IBM's own technology, serving as a proving ground for the secure, trusted AI practices we bring to clients.

Policy Recommendations

To foster an environment that balances innovation with appropriate safeguards, policymakers should prioritize policies that emphasize openness and technological progress. A healthy, competitive AI ecosystem depends on open source AI to ensure America's future in this domain is shaped by the many, not controlled by a select few. As the recent AI Action Plan noted:

> "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."

At the same time, policies should avoid unnecessary or overly burdensome regulation. Instead, regulators should leverage existing authorities to address specific AI use cases. Where guardrails are warranted, they should be risk-based and tailored to the distinct roles and responsibilities of organizations across the AI development lifecycle.

To help AI thrive responsibly in financial services, we recommend that policymakers:

- **Support Open Ecosystems**. Avoid unduly burdening the open source AI ecosystem with legislative and regulatory policy proposals that hamper contributions to, and the flourishing of, open innovation. Open source AI models should not be disadvantaged relative to proprietary AI models

- **Adopt Use-Case-Based Regulation**. Regulate applications, not just technology. An AI model used for summarizing SEC filings should be governed differently than one used for autonomous trading.

- **Promote Transparency Requirements**. Enterprises and regulators must be able to answer: Can I test my model for accuracy? What decisions did it influence? What safeguards are in place?

Conclusion

Generative AI is no longer a future technology — it is here, and it is reshaping financial services today. IBM's journey with Watson, from expert systems to *Jeopardy!* to watsonx, illustrates that AI's future is not about machines replacing people. It is about machines augmenting people — empowering professionals, enriching consumers, and expanding opportunity.

The financial industry has a once-in-a-generation opportunity to lead by example: to show that AI can deliver value while upholding trust. By embracing openness, demanding transparency, and insisting on security, we can ensure that AI strengthens, not undermines, the financial system.

Thank you, and I look forward to your questions.