Swaminathan J: AI in finance – what can change, what must never change

Text of the XX CUB Shri V Narayanan Memorial Lecture by Mr Swaminathan J, Deputy Governor of the Reserve Bank of India, at SASTRA University, Thanjavur, 11 April 2026.

* * *

Dr S. Vaidhyasubramaniam, Vice-Chancellor of SASTRA University, Shri G. Mahalingam, Chairman of the Board of City Union Bank, Dr. N. Kamakodi, MD & CEO, City Union Bank, distinguished guests, esteemed faculty members, staff and dear students, ladies, and gentlemen. A very good morning to all of you.

It is indeed an honour to deliver the Shri V. Narayanan Memorial Lecture at SASTRA University. This lecture series is special because it commemorates not merely an individual, but a rich tradition of banking exemplified by him.

Shri V. Narayanan is remembered as a transformational leader of City Union Bank, founded in 1904 in Kumbakonam, a town with which I, too, share a personal connection. Often described as a 'statesman banker', he combined institutional vision with personal warmth, prudence with progress, and ambition with rootedness.

Under his leadership, City Union Bank grew from a largely regional institution into one with a wider national presence. He invested in staff development, strengthened systems, strongly supported small and medium enterprises, and brought technology into banking, ahead of its time.

Yet, even while embracing change, he never allowed banking to become impersonal. That, to my mind, is what makes his legacy so relevant to our times.

We are living through another moment of profound change in finance. Artificial Intelligence is beginning to reshape how financial institutions serve customers, process documents, assess credit, monitor risks, and strengthen oversight. The speed of that change is remarkable. The real question before us is not whether finance will become more intelligent but whether it will remain fair, accountable, inclusive, and humane.

That is why I felt it appropriate to speak today on the subject: AI in Finance: What can change, what must never change. It is a fitting theme for this occasion.

It is fitting, first, because SASTRA has been consciously building capabilities in this space through collaboration, research, and practical engagement. That is both timely and important. As a country, we will need our own talent, our own institutional capacity, and our own ethical judgment to design, test and govern AI systems suited to our economy and society.

It is fitting, second, because Shri V. Narayanan, in whose memory we gather today, believed in the responsible use of technology and in ensuring that progress remained anchored in sound judgment.

The responsibility of institutions such as SASTRA, therefore, is not merely to produce engineers and professionals, but to help shape responsible builders of the future.

Opportunities in AI

Let me begin with the promise that AI holds for finance. At the outset, however, let me clarify that this lecture is not intended to be a technical exposition, for which this University undoubtedly has ample talent. I propose instead to reflect on the broader questions that AI raises from the perspective of a financial sector practitioner.

Finance, at its best, reduces uncertainty and expands opportunity. It helps households save, businesses grow, farmers invest, students pursue their aspirations, and entrepreneurs dream a little bigger.

Yet finance also has its barriers. It can be overwhelming, documentation-heavy, language-bound, and at times physically distant. In a country as large and diverse as India, technology can help reduce many of these frictions.

AI-enabled systems can make customer interaction simpler, more intuitive, and more responsive. Multilingual chatbots and voice-based interfaces can help customers who are not comfortable with formal paperwork or English-language interfaces. Routine queries can be answered faster. Complaints can be tracked better. Information can be delivered more clearly. For many people, that can make the difference between formal finance feeling accessible and alien.

AI can also help improve credit delivery. Traditional finance has relied on collateral, financial statements, and standardised credit templates. These remain important and will continue to matter. However, they do not always capture the full story of a borrower, especially for small businesses, informal enterprises, first-time borrowers, and others with thin formal credit histories.

Used responsibly, AI can supplement traditional methods by drawing insights from a wider set of patterns in transaction behaviour, repayment flows and business activity. This can help identify viable borrowers who might otherwise remain excluded. For a country committed to inclusive growth, this is a significant opportunity.

AI can contribute meaningfully to fraud detection and risk management as well. Modern financial systems generate vast quantities of data. AI can help identify unusual patterns, flag suspicious activity and support faster intervention. This is especially important in payments, where public confidence depends on both convenience and safety. In this sense, AI can contribute not just to speed, but to safety.

There is also a role for AI in compliance and supervision. Financial supervision today cannot rely only on periodic reporting and backward-looking assessments. Intelligent tools can assist in analysing large volumes of information, identifying patterns, drawing attention to anomalies and supporting early warning. Used well, such tools can help institutions manage risk more effectively and enable supervisors to focus more on emerging issues.

Concerns

So, the promise is real. But, as history has proven, every powerful technology is a double-edged instrument.

If AI is adopted without adequate safeguards, it can amplify existing weaknesses and create entirely new forms of harm. Therefore, the conversation about AI in finance must be balanced. We should neither be taken in by technological hype nor retreat into being defensive.

Let me briefly highlight five major concerns.

(i) Bias and unfair outcomes

The first is bias and unfair outcomes. AI systems learn from data. But data does not emerge from a vacuum. It carries the imprint of past behaviour, existing inequalities and structural exclusions. If these distortions are embedded in the data, they can be reproduced by the model, sometimes with even greater efficiency and scale.

In credit assessment, this can create outcomes that are difficult to justify and harder to detect. What appears objective on the surface may, in fact, nurture unfairness beneath the surface. In finance, this is not merely a technical concern. It is a question of consumer protection, inclusion, and equity.

(ii) Black box nature of some systems

The second is opacity. Many advanced systems operate like black boxes. They can produce an output, but not always in a way that is intelligible to a customer, a manager or even a regulator. But finance cannot become a black box. If a person is denied credit, an account is frozen, a transaction is wrongly flagged, or a product is incorrectly pushed to a customer, the institution must be able to explain the basis for that decision. A decision that materially impacts a citizen's economic life cannot be defended by saying, "the machine decided."

(iii) Data privacy and misuse

The third concern is data privacy and misuse. AI systems rely on large volumes of data, and financial data are among the most sensitive forms of personal information. Institutions must therefore think seriously about consent, storage, sharing, access controls and purpose limitation. Data governance cannot be treated as a side issue. In the age of AI, trust becomes central.

(iv) Model risk

The fourth concern is model risk and concentration risk. In an earlier era, a weak judgment in one office might affect a limited number of accounts. In the AI era, a flawed model can affect decisions across millions of customers. Further, if multiple institutions rely on similar models, common datasets, a small set of vendors or shared infrastructure, individual vulnerabilities can become correlated vulnerabilities. This is where even a local weakness can acquire broader systemic significance.

(v) Cyber risk

The fifth concern is cyber risk. AI can strengthen defences, but it can also equip attackers. Fraudsters and bad actors can use AI to craft more convincing phishing attempts, create deepfakes, probe systems more effectively and automate malicious activity. As finance becomes more digital and more interconnected, resilience becomes even more critical.

Guiding principles

What then should guide us, as we set course on the path towards a full-scale AI adoption? In my view, five broad principles should shape the responsible use of AI in finance:

First, human responsibility must remain central. AI may support decision-making, but accountability must remain with humans and institutions. A bank or NBFC cannot outsource responsibility to an algorithm, a vendor or a platform. Technology may help process information at speed and scale, but judgment and responsibility must continue to reside where they belong.

Second, fairness and explainability must be built into the system from the beginning. They cannot be treated as optional extras. Different stakeholders need different kinds of explanations.

A customer deserves a clear and understandable reason for an important decision. Management needs to understand how the model behaves, where its limitations lie and what assumptions drive it. Supervisors need confidence that systems are robust, auditable, and well-governed. The point is not to make every model simplistic. The point is to ensure that it remains understandable at the appropriate level.

Third, strong data governance is essential. Institutions must think carefully about the full lifecycle of data: how it is collected, on what basis it is used, how long it is retained, who can access it and how it is protected. Privacy and innovation should not be seen as mutually opposed. The institutions that endure will be those that learn to reconcile both.

Fourth, institutional capacity must be strengthened. AI in finance is not only a technology challenge. It is also a governance, capability, and cultural challenge. Boards and senior management need to understand enough to ask the right questions. Risk managers need to know what to validate. Supervisors need the capacity to examine AI-enabled systems intelligently. And universities need to produce graduates who are not only technically competent, but also alive to questions of ethics, regulation and public purpose.

Fifth, inclusion must be a design objective, not an accidental by-product. Scale by itself does not mean inclusion. We must ask a harder question: who is still left out? The best innovation is not the one that dazzles those already well served. The best innovation is that which makes formal finance simpler, safer and more useful for those who are at the margins, because of geography, language, literacy, age or income. If AI helps bridge those gaps, it advances inclusion. If it quietly deepens exclusion, we would have failed in its design. As I have said before, inclusion should be innovation's highest purpose[1].

In India, the true value of AI in finance should be judged by three tests.

i. Does it advance inclusion?

ii. Does it improve efficiency?

iii. Does it strengthen trust?

If it does these three things positively, then it serves a meaningful public purpose. If it does not, then its sophistication alone should not impress us.

Lesson from Shri Narayanan

At this point, let me come back to Shri V. Narayanan.

Those who knew him recall that he thought ahead of his time. He brought technology to banking earlier than many comparable institutions did. Yet he also retained a human touch, especially in lending relationships with small and medium enterprises. That combination is deeply instructive. He was not choosing between technology and relationships. He was showing how progress and human judgment must go together.

One line often associated with him captures this beautifully: *"Take care of the bank; the bank will take care of you."* It is a simple statement, but it contains a profound institutional ethic. It speaks of stewardship. It reminds us that institutions flourish when people treat them not merely as sites of transaction or employment, but as repositories of trust. That insight is just as relevant in the age of AI as it was in the age of ledgers and branch registers.

The lesson from Shri Narayanan's life is that technological change in finance must remain anchored in stewardship, trust and responsibility.

Banking, at its heart, is a business of trust. A financial institution can survive a difficult quarter, an operational mistake, or even a strategic setback. But it cannot easily survive the erosion of trust. That is why innovation in finance must always remain subordinate to integrity, fairness and accountability.

Role of Students

For the students in this hall, this is not a distant issue. By the time many of you are in mid-career, AI will be woven into almost every part of the financial world.

I would suggest you learn these tools deeply. Understand the technology seriously. Build technical competence with rigour and curiosity. But more importantly, also carry with you an equally deep commitment to ethics, transparency and public interest.

In a world shaped by AI, technical excellence without ethics can do great harm. The real test of your generation will not be whether you can build powerful systems. It will be whether you can build systems worthy of public trust.

If India can combine its digital strengths, entrepreneurial energy, scientific talent and institutional wisdom, then we can build a financial sector that is not only more efficient, but also more inclusive, more resilient and more trustworthy. That should be our aspiration. We should not pursue technology for its own sake, but rather use it in the service of people.

Intelligence without accountability does no good; it must be guided by sound and ethical judgment. Our endeavour, therefore, should be to foster innovation that strengthens institutions for the long term.

Conclusion

Let me conclude with this thought.

Every generation receives a few powerful tools. This generation has grown up with digital technology and artificial intelligence. History does not judge societies by the sophistication of the tools they possessed, but by the values that guided their use.

If AI helps widen opportunity, improve access, strengthen prudence, protect customers and deepen trust, then it will have served a noble purpose. On the other hand, if it weakens accountability, obscures decisions, excludes the vulnerable or turns finance into an impersonal black box, then it will have taken us away from the ideals that bankers like Shri V. Narayanan stood for.

The enduring task, therefore, is to make finance more intelligent, without making it less human; to make it more digital, without making it less accountable; and to make it more inclusive, without making it less prudent. This, in a nutshell, is what can change and what must never change.

My heartfelt gratitude to the organisers for this opportunity. I wish SASTRA, its faculty, and its students the very best in all their endeavours. May God guide you and bless your efforts in all that you seek to achieve. Thank you. Jai Hind.

---

[1] Swaminathan J, "Inclusion is Innovation's Highest Purpose: Lessons from India," Reserve Bank of India, October 15, 2025, https://rbi.org.in/scripts/BS_SpeechesView.aspx?Id=1526