Artificial Intelligence, Algorithmic Governance and the Future of Securities Markets

A compliance officer stays late one Friday and catches what an algorithm missed. That single moment holds this article's whole argument. Artificial Intelligence is quietly becoming the infrastructure beneath securities markets, drafting disclosures, flagging trades, and shaping decisions before anyone sees them. Yet markets have never run on cleverness alone, only trust. Drawing on Kautilya, Toffler, Drucker, Kahneman, and a steadying verse from the Bhagavad Gita, this article asks one question. When the algorithm is fast, confident, and almost always right, who still stays behind to ask one more question.

CS Rahul Sahasrabuddhe, FCS
Independent Director & Designated Partner
SPRS & Co LLP, Company Secretaries
Mumbai
rahul@sprsllp.com

INTRODUCTION

Artificial Intelligence has moved from experimental use to routine infrastructure inside listed companies and market intermediaries faster than most governance frameworks have managed to track. Surveillance systems now flag anomalies before compliance teams do. Drafting tools produce the first version of disclosure notes. None of this arrived through a single law or announcement. It arrived gradually, one tool at a time, until it simply became how markets function.[1]

On current trends, this will look less like a technology add-on within five years and more like standard market infrastructure. Five years, at today's pace, is starting to resemble a full generational shift rather than a modest stretch of time.

This article asks whether governance frameworks, and the people responsible for them, are ready for that shift. It works through first principles rather than cataloguing tools, since tools keep changing while questions of trust, accountability, and judgment do not. Artificial intelligence does not change why markets exist. It changes who gathers information, who interprets it, and who must answer for the decisions built on it.

1. *Securities and Exchange Board of India, circulars on AI and machine learning reporting for stock brokers and depository participants (January 4, 2019), market infrastructure institutions (January 31, 2019), and mutual funds (May 9, 2019).*

WHAT ALGORITHMIC GOVERNANCE ACTUALLY MEANS?

It helps to say plainly what algorithmic governance means, since the phrase gets used more than it gets explained. It is not another word for Artificial Intelligence. It is the set of rules and human duties built around a system that makes or shapes decisions on its own. A model that flags a suspicious trade is AI. The rule that a human must check every flag before acting on it is algorithmic governance. A robo-adviser that recommends a mutual fund is AI. The duty on the firm to test that tool for bias, document how it works, and take the blame when it fails, is algorithmic governance. AI supplies the engine. Governance supplies the seatbelt, the wheel, and the licence to drive.

In practice, this rests on a few habits, whichever regulator is writing the rulebook. Test the system before it goes live, and keep testing it as markets change. Be able to explain, in plain words, why the model reached its answer. Keep a written record of what it was told and who approved using it. Keep a genuine, tested switch to turn it off. Strip away the jargon, and these are an old question asked in a new way. Who is watching, and can they actually stop it if something goes wrong.

WHY MARKETS EXIST?

Markets exist to solve one very old problem, namely how strangers trade something valuable without knowing or trusting each other. A farmer sells grain to a trader he has never met. A retail investor buys shares in a company run by people she will never speak to. A pension fund commits billions based on numbers it cannot personally check. Every one of these trades works only because someone stands behind the promise.

Kautilya understood this two thousand years before the word regulator existed. In the Arthashastra, he describes a state that appoints officers to check weights and measures and keep traders honest, not because trade was suspicious but because it was too valuable to leave unwatched.[2] A market with no referee does not stay a market for long. It turns into fraud and force. That is the oldest rule in this field, and AI has not changed it.

2. *Kautilya, Arthashastra, Book II, "The Duties of Government Superintendents," R. Shamasastry translation.*

Modern securities regulation describes the same idea in more technical language. A market is a web connecting issuers who raise money, intermediaries who move it, and investors who supply it. That web holds together on trust and transparency. Rules enforced by bodies such as the Securities and Exchange Board of India protect this web, mainly by stopping insider trading and fraud, and by forcing companies to disclose what investors need to know, on time and in plain terms.[3] AI does not change this underlying goal. It changes the tools used to reach it, and raises fresh questions about who is watching those tools.

THE LONG ARC OF POWER

Alvin Toffler gave this larger story a name before most people had touched a computer. In his 1990 book Powershift, he argued that real power had moved, over centuries, from violence, to wealth, to knowledge, each one quieter and harder to spot than the last.[4] Markets have always followed whichever force was ascendant. Money then steered them for a long while. Today, markets are shaped by whoever understands information first, sometimes by no more than a few seconds. AI is simply the sharpest tool yet built for exactly the kind of power Toffler described. A few of his specific predictions look dated now, but his larger claim about where power moves has aged well.

A WEALTH OF INFORMATION, A POVERTY OF ATTENTION

Herbert Simon, an economist studying how organisations decide things, put his finger on the cost of an information-rich world. In a 1971 essay, he wrote a line that has aged unusually well, that too much information leaves us with too little attention to spend on any of it.[5] Markets today are drowning in exactly that flood, as machine-written research, instant commentary, and endless disclosure feeds pour in faster than any mind can absorb. Genuine signal does not get easier to find inside that flood. It may get harder, since everything now arrives dressed in the same confident, machine-polished tone, whether or not it has earned that confidence.

DISRUPTION FROM BELOW

Clayton Christensen spent his career asking why strong, well-run companies keep losing to small newcomers nobody took seriously at first. His 1997 book The Innovator's Dilemma set out the answer, known as disruptive innovation.[6] New technology usually enters at the cheap, unimpressive bottom of a market, serving customers the big players never cared about, until it has quietly improved enough to eat the whole market from below.

A decade ago, AI-driven robo-advisers and unlicensed social-media finfluencers looked like a cheap joke, fit only for customers too small for a wealth manager to bother with. Today they shape the choices of millions of retail investors, often reaching further than any licensed research analyst. Regulators, understandably, built their rulebooks around the players already inside the tent, registered analysts and licensed advisers, exactly the group Christensen would call least disruptive. The real threat grew up outside the fence regulators were watching, and by the time it mattered, it was already too large to easily contain.

THE DISCIPLINE OF NOT DELEGATING

Peter Drucker wrote about institutions rather than markets, but his idea belongs alongside Christensen's. In his 1967 book The Effective Executive, he argued that every institution exists to make one specific contribution to someone outside itself, and should be judged on that, not its own internal rules.[7] A regulator that forgets this becomes obsessed with process for its own sake. A regulator that remembers it keeps asking whether a rule helps a stranger trust another stranger with money. Drucker also insisted that one task can never be fully delegated, the decision itself, since everything else is preparation for it. Machines can prepare a decision beautifully. They are simply not able to own it.

CONFIDENT MACHINES, QUIET DOUBTS

That distinction matters more, not less, as machines take on more of the thinking. Daniel Kahneman, in his 2011 book Thinking, Fast and Slow, laid out two ways human minds decide things, one fast and pattern-driven, the other slow, careful, and doubtful.[8] His research, done with Amos Tversky, showed how easily the fast mode takes charge whenever a confident answer sits in front of us.

Algorithms behave like an extreme version of that fast mode. They never hesitate, and never look unsure. A director who would question a junior analyst's confident guess may not think to question an algorithm's equally confident answer, because the machine never visibly flinches. Good governance needs deliberate friction built in on purpose, a named person whose job is to disagree with the model before it acts. Skip that step, and the slow, careful mode of thinking gets quietly handed to a system that was never built to have one.

A DIFFERENT VERSE, A DIFFERENT LESSON

This is where a particular verse from the Bhagavad Gita earns its place, read a little differently than usual. Most readers go first to the famous line about acting without attachment to results. There is another passage, a little later in the second chapter, that speaks more directly to what a fiduciary needs when the world around her is loud, fast, and full of confident machines.[9]

duhkheshv anudvigna-manah sukheshu vigata-sprihah
vita-raga-bhaya-krodhah sthita-dhir munir ucyate

3. *Securities and Exchange Board of India Act, 1992, and the SEBI (Prohibition of Insider Trading) Regulations, 2015.*
4. *Alvin Toffler, Powershift, Bantam Books, 1990.*
5. *Herbert A. Simon, "Designing Organizations for an Information-Rich World," in Martin Greenberger (ed.), Computers, Communications, and the Public Interest, Johns Hopkins University Press, 1971, pages 37 to 52.*
6. *Clayton Christensen, The Innovator's Dilemma, Harvard Business School Press, 1997.*
7. *Peter Drucker, The Effective Executive, Harper and Row, 1967.*
8. *Daniel Kahneman, Thinking, Fast and Slow, Farrar, Straus and Giroux, 2011, drawing on research conducted with Amos Tversky.*
9. *Bhagavad Gita, Chapter 2, verses 56 and 58, on the person of steady wisdom, sthitaprajna.*

In plain words, the person of steady wisdom stays calm through both pain and pleasure, free of the fear and anger that usually follow either. A few lines later comes an image most people remember even after they forget the rest of the chapter.

yada samharate cayam kurmo 'nganiva sarvasah
indriyani indriyarthebhyas tasya prajna pratishthita

When a person pulls the senses back from their objects, the way a tortoise draws its limbs into its shell, that person's wisdom stands firm. Krishna is describing the ability to choose, on purpose, what deserves attention, instead of chasing whatever is loudest. That image fits a modern director or fund manager rather well. Good governance, in the age of AI, looks a lot like that tortoise, not slow for its own sake, but unhurried on purpose about which signals deserve a real response.

FRAGILITY IN NUMBERS

Nassim Nicholas Taleb offers a useful warning about what happens when many decision-makers lose that steadiness together. In his 2012 book Antifragile, he points out that systems built hard for efficiency often turn quietly brittle, stable in normal times but prone to sudden, joint collapse the moment conditions shift.[10] When many market players license similar AI models trained on similar data, their behaviour can drift together in ways no insider-trading rule was built to catch. Sensible choices made separately by independent firms can still add up to a fragile market as a whole. The Bank for International Settlements now treats this kind of correlated, model-driven risk as a real threat to financial stability, not just a theory.[11]

> Surveillance is where AI helps the most and draws the least argument, since no human team could scan millions of daily trades for a pattern a well-trained model spots with ease, provided the model itself stays open to inspection.

HOW REGULATORS ARE RESPONDING?

None of this means AI should be feared or worshipped. It is a powerful tool that changes who gathers information, who ranks it, and who shapes a decision before any person sees it. The AI Principles of Organisation for Economic Co-operation and Development (OECD), agreed in 2019 and used since by more than forty countries, ask for the same four things in every industry. Systems should serve human values, stay explainable, remain safe and steady, and sit inside a clear chain of responsibility.[12]

Bodies such as SEBI have asked market infrastructure firms and mutual funds to report their AI use since 2019, and newer rules, including Regulation 16C and a detailed 2025 paper on responsible AI use, push that oversight further into the industry.[13] International Organization of Securities Commissions (IOSCO) 2026 Supervisory Toolkit sets out detailed tools on governance, third-party risk, disclosure, recordkeeping, and monitoring, building on its earlier 2021 and 2025 reports.[14] European Securities and Markets Authority (ESMA) has taken a similar path under its investment-services rules, and the American Securities and Exchange Commission (SEC) dropped one specific proposal in 2025, its predictive-data-analytics rule.[15] The question these regulators keep circling back to is the same one Kautilya's officers were asking two thousand years ago. Who is responsible when something goes wrong, and how do we find out before it does.

WHERE THE IMPACT SHOWS UP?

Disclosure rules, built for a world of slow, human-written statements, now have to decide whether a company's internal AI forecasting tool needs its own disclosure, the way a change in accounting policy would. Surveillance is where AI helps the most and draws the least argument, since no human team could scan millions of daily trades for a pattern a well-trained model spots with ease, provided the model itself stays open to inspection.

Market integrity has its own warning story in India's National Stock Exchange co-location matter, where certain trading members were found to have faster, preferential access to trading infrastructure, showing how fast an efficiency tool turns into an unfair edge once oversight falls behind.[16] Insider-trading law, built around a named person holding secret information, sits awkwardly with an algorithm that draws a valuable conclusion purely from public data, an advantage that looks much like the old kind without technically breaking the old rule.[17] Enforcement orders, which have

10. *Nassim Nicholas Taleb, Antifragile, Random House, 2012.*
11. *Bank for International Settlements, published commentary on artificial intelligence, model correlation, and financial stability, available at bis.org.*
12. *Organisation for Economic Co-operation and Development, OECD AI Principles, adopted May 2019, updated 2024, available at oecd.ai.*
13. *Securities and Exchange Board of India, "Guidelines for Responsible Usage of AI/ML in Indian Securities Markets," Consultation Paper dated June 20, 2025; and SEBI (Intermediaries) (Amendment) Regulations, 2025, inserting Regulation 16C, notified February 10, 2025, available at sebi.gov.in.*
14. *International Organization of Securities Commissions, "Supervisory Toolkit for AI Use in Capital Markets," Final Report FR/02/2026, May 2026; "The Use of Artificial Intelligence and Machine Learning by Market Intermediaries and Asset Managers," Final Report FR06/2021, September 7, 2021; and "Artificial Intelligence in Capital Markets, Use Cases, Risks, and Challenges," Report FR/04/2025, 2025, all available at iosco.org.*
15. *European Securities and Markets Authority, "Public Statement on the Use of Artificial Intelligence in the Provision of Retail Investment Services," ESMA35-335435667-5924, May 30, 2024; and United States Securities and Exchange Commission, "Conflicts of Interest Associated with the Use of Predictive Data Analytics by Broker-Dealers and Investment Advisers," proposed August 9, 2023, formally withdrawn June 12, 2025.*
16. *Securities and Exchange Board of India, adjudication and appellate record in the National Stock Exchange co-location matter, including the Securities Appellate Tribunal order of January 23, 2023, and subsequent SEBI orders.*
17. *SEBI (Prohibition of Insider Trading) Regulations, 2015, Regulation 2(1)(n) and Regulation 3.*

always needed to be explainable to survive appeal, get harder to defend when the evidence is a probability score nobody in the room can fully walk through. Capital formation, often the quiet good-news story here, is genuinely being helped by AI tools that cut the cost of raising money for smaller companies, even as that same efficiency risks concentrating real power in a few dominant vendors.

SUSTAINABLE GROWTH AND THE SAME OLD QUESTION

Rules folding Environmental, Social and Governance (ESG) factors into regulatory mandates are, at bottom, an attempt to protect the system from slow-building, systemic risks rather than single bad trades. AI touches this thread too. It can genuinely help, by processing huge volumes of climate and governance data which no human team could read in full and catching early warning signs buried in disclosures, nobody has time to reread. It can also quietly hurt, if a model trained mostly on the past keeps steering capital toward familiar industries and away from the harder-to-score bets long-term growth actually needs. Sustainable growth depends on the same thing, this whole article keeps returning to, a human being willing to ask a confident model why it reached its answer.

WHAT THIS MEANS IN PRACTICE?

None of this argues for one sweeping AI statute, and the pattern across jurisdictions bears that out. Nearly everyone agrees the old duties, loyalty, prudence, fair execution, and honest, timely disclosure, need to be read carefully enough to survive contact with a decision-maker that is not, strictly speaking, a person.

For a board, this means treating algorithmic oversight as a genuine fiduciary matter, not something handed to a technology sub-committee, and judging management on the diligence of its process rather than whether last quarter's model guessed right. Compliance Officers should insist on a testable explanation for any AI-driven flag before acting on it, and actually test the kill switch. Investors can simply ask any adviser, human or robotic, what data trained the recommendation, and whether its interests align with theirs. A regulator's task is to anchor rules to who owns an algorithm and what it optimises for, since that kind of rule tends to outlast whatever model is fashionable this year. Intermediaries would do well to extend existing vendor discipline to every third-party AI tool inside the client relationship.

THE ROAD AHEAD

Expect existing frameworks to be stretched and amended rather than torn up. Expect explainability to slowly become its own disclosure category, sitting beside financial and sustainability reports rather than buried in a technical appendix. And expect supervisory technology to grow quickly too, with regulators using AI to watch AI, since algorithmic activity has already outpaced what human examiners can review by hand.

CONCLUSION

Consider a scenario that plays out inside compliance departments without ever making a headline. A Compliance Officer at a mid-sized brokerage was closing out a Friday evening when the exchange's surveillance system, freshly upgraded that quarter, sent its usual quiet report. Nothing had been flagged, not once, all day. He/she could easily have gone home.

Instead, he/she noticed something the system had no reason to notice, because it was never built to notice it. A dozen small accounts were each trading a modest amount of a thinly held stock, each just under the size that would trigger review, each opened within the same fortnight, each funded from addresses on the same street. No single account on its own looked strange, but together they told a very different story, the kind a machine built to watch one account at a time was never going to piece together alone.

He/she flagged the pattern anyway, against a system that had told him/her that there was nothing there, and set out his/her reasoning in three short paragraphs, a colleague could challenge if he/she disagreed. It took the exchange eleven days to confirm what he/she had suspected, a coordinated push to walk up the price of a stock nobody had been watching closely enough. The algorithm had not really failed. It had simply done what it was built to do, watch single accounts for single anomalies, since nobody had ever told it to notice a pattern among neighbours.

No new rule needed to be written because of that Friday evening. The principle that mattered was already in place, the same one Kautilya had in mind when he/she told his/her officers that a market needs more than one pair of eyes watching it. It is the same principle every fiduciary duty in every jurisdiction still quietly assumes, that a machine prepares the picture and a person decides what it actually means.

So here is the question worth sitting with as these tools keep improving. The algorithm will be fast, confident, and almost always right. Who among us will still be the one willing to stay a little later, ask one more question, and risk being wrong out loud rather than quietly right and unaccountable. Markets, like anything worth trusting, have never really been stronger than that last person at the desk.

❐