

**Artificial Intelligence in Investing:  
Adoption, Systemic Risk, and the  
Question of Who Sets the Price  
India Has Built Its AI Rulebook. The  
Risks Are Moving Faster.<sup>1</sup>**

*July 2026*

**Shalini Talwar**

Professor, Finance and Accounting, SPJIMR, Mumbai

**Vidhu Shekhar**

Associate Professor, Finance and Accounting, SPJIMR, Mumbai

**Anubhav Roy**

Student, PGDM (Batch 2024-2026), SPJIMR, Mumbai

# **Artificial Intelligence in Investing: Adoption, Systemic Risk, and the Question of Who Sets the Price**

**India Has Built Its AI Rulebook. The Risks Are Moving Faster.<sup>1</sup>**

**Shalini Talwar**

Professor, Finance and Accounting, SPJIMR, Mumbai

**Vidhu Shekhar**

Associate Professor, Finance and Accounting, SPJIMR, Mumbai

**Anubhav Roy**

Student, PGDM (Batch 2024-2026), SPJIMR, Mumbai

## **Executive Summary**

Artificial intelligence (AI) is increasingly permeating the operating core of how money is allocated and being deployed in markets that are already heavily automated. Industry estimates put algorithmic and high-frequency strategies at roughly 60-70% of equity volume on the world's leading exchanges<sup>i</sup>, with algorithmic trading accounting for around 53% of cash-market volume on India's NSE in 2024.<sup>ii</sup> SEBI itself has reported that in FY 2024, 97% of foreign portfolio investors' profits and 96% of proprietary traders' profits in Indian equity derivatives were generated by algorithmic trading.<sup>iii</sup> These figures measure automation rather than AI specifically. They encompass rule-based execution, high-frequency strategies, and conventional quantitative models alongside machine-learning systems, and do not separate the two. What they establish is the scale of automation into which AI is now being deployed across investing. This paper examines the deployment: its visible failures, hidden fragilities, and systemic risks, then builds toward the core question of what it could do to price discovery and value.

AI in investing is a genuine opportunity, both to expand reach and to enhance information-processing capabilities at a speed and scale no human can match. For example, it

---

<sup>1</sup> Prepared for consideration by policymakers, market regulators, firms in the BFSI sector, and institutional investors

is argued that AI could generate efficiencies equivalent to 25-40% of asset managers' cost base.<sup>iv</sup> But the risks arising from its deployment are equally real. It has been documented to produce errors and enable malicious use: a factual error in a chatbot demonstration was followed by a roughly USD 100 billion drop in Alphabet's market value,<sup>v</sup> while a finance employee authorized approximately USD 25 million in transfers after an AI-generated deepfake video call.<sup>vi</sup> Risks also arise from structural exposure stemming from dependence on a small number of foundation models, cloud, and specialized hardware providers.<sup>vii</sup> Historical precedents, evident in events such as LTCM, the 2007 Quant Quake, the 2010 Flash Crash, and Knight Capital, make concerns about AI deployment more serious. While these events were not AI-related, together they show how model dependence, leverage, crowded positions, automated execution, software failures, and deteriorating liquidity can interact to amplify losses. AI does not mitigate these vulnerabilities. Instead, when institutions rely on the same models, data, and vendors, it may intensify some of them.

This may also make organizations cautious about how visibly they position their use of AI. JioBlackRock's retail advisory platform is a case in point. It extends genuine reach to Indian retail investors and uses BlackRock's Aladdin technology for daily portfolio monitoring.<sup>viii</sup> Interestingly, although Aladdin includes AI and machine-learning capabilities, the service is presented as being underpinned by institutional-quality risk analytics rather than explicitly positioned as an AI application. That restraint could be read as caution about placing AI at the front of a mass-market retail product.

Indian markets sit in an unusually exposed position. Retail participation has expanded rapidly, with demat accounts rising from roughly 36 million in 2019 to 194 million by mid-2025. By the second quarter of 2024, India's equity index options market had become the world's largest by number of contracts traded.<sup>ix</sup> Scale of that kind, combined with activity spread across affiliated entities and across the cash, futures, and options segments, creates a demanding surveillance problem. SEBI's July 2025 interim order against four Jane Street entities, alleging coordinated manipulation of index levels, illustrates the difficulty.<sup>xi</sup> Jane Street contests the allegations, and its appeal remains before the Securities Appellate Tribunal.

India has moved early by international standards. Since 2024, it has approved the IndiaAI Mission, issued SEBI's framework for retail participation in algorithmic trading, published the RBI's FREE-AI report, released national AI Governance Guidelines, and articulated the M.A.N.A.V. vision. This is a real architecture, and in the RBI's case, an

operationally specific one. But roughly INR 400 crore has been released against the Mission's INR 10,372 crore five-year envelope,<sup>xii</sup> and publishing frameworks is not the same as building the capacity to apply them. None of these instruments establishes that regulators have the specialist staff, tooling, and data access needed to inspect a model whose developers cannot fully explain it. The rulebook exists. The question is whether it can keep pace with the machine.

While these operational risks cannot be ignored, this paper focuses on something deeper, albeit narrower: how the expansion of AI in investing may affect price discovery and, through it, market efficiency. Market efficiency is not self-sustaining. It depends on a price-discovery process in which enough investors are willing to gather information, study businesses, and form independent judgments. Our argument is that as AI systems optimized for pattern recognition rather than business fundamentals become more widely used in investing, they could change both the incentive to do that work and the range of views that ultimately shape prices. We do not claim that AI has already weakened this process. Rather, we ask what may happen as large pools of capital increasingly rely on similar data, signals, models, vendors, and assumptions. If price discovery starts reflecting an algorithmic and model-driven consensus as much as, or more than, independent analysis of underlying value, the question of who or what is setting the price becomes unclear.

This paper draws on public regulatory filings, industry research, academic work, including working papers and preprints, and major firm disclosures. It walks through the visible failures, hidden fragilities, and finance-specific systemic risks that AI creates or amplifies; sets them against the governance framework India has built since 2024; and closes with the authors' view of what all this may be doing to price discovery. A companion paper takes up the policy response.

## **1. Background: AI Use Cases in Investing**

AI adoption in finance is no longer experimental at the global frontier; it is advancing rapidly worldwide. Globally, AI is reshaping every layer of the investment value chain from data ingestion and signal generation to trade execution and risk monitoring. In India, adoption is accelerating across institutional asset managers, retail platforms, and regulatory infrastructure. AI in today's era is driven by a combination of market scale, digital penetration, and competitive pressure from globally deployed AI systems.

The following maturity tiers are an analytical framework we developed for this paper. They describe how deeply data-driven systems are embedded in decision-making, determined by data structure, decision frequency, and feedback-loop speed. Conventional algorithms and quantitative models are included where relevant, but should not be read as AI:

- **Tier 1 (systems at the core).** Quantitative hedge funds, electronic market making, high-frequency execution, and automated portfolio allocation. Renaissance, Two Sigma, Numerai, Virtu, Jane Street, Betterment, Zerodha. Software is integral to processing data or executing decisions, though the systems range from deterministic rules to machine learning, and the mix is rarely disclosed.
- **Tier 2 (AI is scaling).** Asset management, sell-side research, ESG investing. Aladdin, BloombergGPT, Kensho, Alpha Sense, MSCI's AI ESG. AI compresses research cycles and extends coverage; alpha generation remains a work in progress.
- **Tier 3 (AI is early).** Private equity, venture capital, and investment banking. AI is being used for sourcing, screening, document preparation, and workflow support, while negotiation, due diligence, transaction judgment, and accountability remain predominantly human. EQT's Motherbrain uses AI for sourcing.

A key observation is that AI adoption is fastest where data is structured, decisions are frequent, and feedback loops are short. This has direct implications for the places where supervisory attention should first concentrate.

## 1.1 Global asset managers and trading firms Use Cases

| Institution           | Documented AI and automation deployments                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>BlackRock</b>      | Aladdin platform: centralized investment operating system serving the end-to-end investment process for BlackRock and external clients across public and private markets; 360,000+ users across Aladdin, Preqin, and eFront; technology revenue of USD 2 billion in 2025, up 24% year on year. 'Asimov' virtual analyst continuously scans research, filings, and communications. <sup>xiii</sup> USD 14.0 trillion AUM at year-end 2025; USD 698 billion in net inflows. <sup>xiv</sup> |
| <b>JPMorgan Chase</b> | LLM Suite for firm-wide use; COIN platform processes a large number of commercial loan agreements annually, work that previously required 360,000 hours of legal review. <sup>xv</sup>                                                                                                                                                                                                                                                                                                   |
| <b>Goldman Sachs</b>  | Generative AI assistant for document summarization; piloting 'Devin,' <sup>xvi</sup> an autonomous software engineer. Goldman Sachs runs the GS AI Platform firm-wide, hosting GPT-4, Gemini, LLaMA, and Claude behind its own firewall. <sup>xvii</sup>                                                                                                                                                                                                                                 |

| Institution                     | Documented AI and automation deployments                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Vanguard</b>                 | Vanguard Digital Advisor is an automated advisory service that builds a personalized plan from a client's risk assessment, age, and goals, drawing on over 300 glide paths and running 10,000 simulated scenarios to project retirement outcomes. As of 30 June 2024, it had more than USD 19 billion in assets under management, at a net advisory fee of roughly USD 15-16 per USD 10,000 annually. <sup>xviii</sup> Vanguard describes this as an automated advisory service driven by an algorithm; the product's public materials do not characterize the allocation process as AI. |
| <b>HSBC</b>                     | HSBC has implemented a Google Cloud-based "Dynamic Risk Assessment" system for anti-money laundering (AML) compliance. The system analyses billions of transactions across millions of accounts, has achieved a 60% reduction in false positives, and detected 2-4 times more financial crimes. <sup>xix</sup>                                                                                                                                                                                                                                                                           |
| <b>Renaissance Technologies</b> | Renaissance Technologies' Medallion Fund returned an average of 39.9% net annually between 1988 and 2022. However, the fund is capped at roughly USD 15 billion with gains distributed annually, so the compounded figure is not achievable at scale. The fund has been closed to outside investors since 1993. <sup>xx</sup> Medallion's record predates modern machine learning and is better read as a precedent for large-scale systematic investing than as an AI deployment.                                                                                                       |
| <b>Two Sigma</b>                | Two Sigma has moved beyond applying AI solely to market signals. According to industry reporting, the firm has adopted an "AI-first" internal mandate requiring employees to integrate frontier AI models into daily workflows, positioning AI as a universal operational layer across research, data engineering, compliance, and infrastructure. The stated aim is operational efficiency rather than a replacement for its existing quantitative trading models. <sup>xxi</sup>                                                                                                       |
| <b>Acadian Asset Mgmt</b>       | Acadian selects from over 40,000 forecasted stocks in its investment universe using systematic, data-driven signals derived from fundamental and technical market inefficiencies, illustrating how automated screening extends analytical coverage well beyond what human teams can sustain. <sup>xxii</sup> Acadian's published materials describe this as systematic investing and do not characterize the process as machine learning.                                                                                                                                                |

## 1.2 Indian Institutions Use Cases

Indian financial institutions are expanding the use of AI across fraud detection, customer service, credit assessment, and document processing. Public evidence varies in quality: several widely repeated performance figures trace to vendor blogs and secondary analyses rather than to institutional disclosure.

The most important development of 2025 was JioBlackRock, a 50:50 joint venture between Jio Financial Services and BlackRock with an initial capital of USD 300 million and a digital distribution footprint of approximately 480 million customers within Reliance's parent ecosystem.<sup>xxiii</sup> Its first New Fund Offer collected around INR 17,800 crore (~ USD 2.1 billion) across three schemes. The JioBlackRock Flexi Cap Fund, whose New Fund Offer ran

from 23 September to 7 October 2025 with inception on 13 October 2025, applies BlackRock's Systematic Active Equity (SAE) framework. Signal research scores supplied by BlackRock are derived from traditional and alternative data using machine learning and advanced analytics, grouped under valuation, quality, sentiment, and fundamental momentum, and employing over 400 India-specific signals. Portfolio construction runs on Aladdin, and fund managers apply judgment and risk oversight to finalize positions.<sup>xxiv</sup>

Retail-facing platforms such as Trendlyne, Tickertape, Screener.in, and Tijori Finance provide stock screening, financial data analysis, and earnings call or document analysis tools, some of which use NLP or AI.<sup>xxv xxvi</sup> Zerodha's Nudge flags risky trades, illiquid or speculative (T2T) stocks, and stocks promoted by unauthorized entities in real time before execution,<sup>xxvii</sup> and its Kill Switch lets traders disable a trading segment for 12 hours to curb over-trading and revenge trading.<sup>xxviii</sup> Groww's GR1, an AI investing co-pilot in beta, provides context-aware portfolio insights and news sentiment analysis but does not execute trades and requires explicit user approval for any action.<sup>xxix</sup> In our view, oversight of these tools should be proportionate to the degree of personalization, influence, and execution authority each carries, rather than to whether a platform describes a product as AI.

### **1.3 Investment Decision-making: Three Overlapping Phases**

For analytical purposes, the development of investment decision-making can be understood through three overlapping phases. The first emphasized fundamental analysis, human judgment, and proximity to information. The second saw the expansion of electronic financial data, factor models such as the Fama-French factors,<sup>xxx</sup> statistical arbitrage, and algorithmic execution. The third adds machine learning, natural-language processing, generative AI, and emerging agentic systems. These phases coexist: fundamental investors increasingly use AI tools, conventional quantitative models remain widespread, and most AI-assisted investment decisions continue to involve human oversight. The policy significance of agentic AI lies not in its having already displaced earlier systems, but in the possibility that some systems may increasingly generate, revise, and execute actions with reduced human intervention.<sup>xxxi</sup> Existing regulatory tools, including disclosure, registration, and post-event review, may therefore need to operate more quickly. They may also need to address longer and more complex chains of responsibility.

## 1.4 Snapshot of AI in Finance

| Metric                                                                   | Value                                                                                                                                                                                                                              |
|--------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Global AI spending in finance (forecast)                                 | 2023 → 2027 (est.): USD 35 billion → USD 97 billion <sup>xxxii</sup>                                                                                                                                                               |
| AI's projected contribution to global GDP by 2030 (PwC estimate)         | up to USD 15.7 trillion <sup>xxxiii</sup>                                                                                                                                                                                          |
| AI in asset management market projection (vendor forecast, 2024 → 2034)  | USD 4.62B (2024) → USD 38.94B (2034); CAGR 23.76% over 2025–2034 <sup>xxxiv</sup>                                                                                                                                                  |
| AI infrastructure spending (US)                                          | Capital expenditure by major US hyperscalers: Microsoft, Amazon, Meta, Alphabet, company-wide and not finance-specific: ~USD 410B (2025) → ~USD 725B (2026) <sup>xxxv</sup>                                                        |
| Algorithmic share of NSE equity trades                                   | ~53% in cash market segment (2024); ~62% in equity derivatives segment (2024) <sup>xxxvi</sup>                                                                                                                                     |
| Share of FY24 gross profits accruing to algo-classified entities (India) | 97% foreign / 96% proprietary <sup>iii</sup>                                                                                                                                                                                       |
| Indian demat accounts                                                    | ~194 million (mid-2025), up from ~36 million in 2019 <sup>ix</sup>                                                                                                                                                                 |
| Mutual-fund AUM-to-GDP (India vs global)                                 | ~20% vs. ~64% (as of March 2025) <sup>xxxvii,xxxviii</sup>                                                                                                                                                                         |
| ESMA AI adoption survey (EU financial firms)                             | Summer 2025 survey of 728 EU entities: 55% report AI use cases in production or in development; 70% expect AI investment to rise 2025–2027; adoption “partial and uneven,” with smaller firms lagging larger ones <sup>xxxix</sup> |

The 20% AUM-to-GDP ratio for Indian mutual funds, compared with a global average of 64%, clearly captures the policy stake. Pooled-investment penetration remains low on this measure, even though India's derivatives markets are already among the world's largest. The AI infrastructure being built now will shape how household participation develops from here.

The same capabilities that make AI transformative in finance also introduce a class of risks that existing regulatory frameworks were not designed to handle. Some are already on record, documented in enforcement actions, regulatory filings, and post-mortems of market events. Others are harder to detect, such as hallucinations, fabricated outputs, and model errors that surface only after capital has been deployed based on flawed analysis. A third category predates AI entirely but takes on a significantly more dangerous dimension in an AI-driven market: historical failures that were serious enough in a human-speed world and whose consequences could be amplified today by greater speed, interconnectedness, automation, and dependence on common technologies.

## **2. AI in Investing: Visible Failures and Concerns**

Concerns about AI and automated financial systems have surfaced in documented cases of losses, compliance failures, fraud, and operational near-misses across the global financial system. Each underscores the urgency of framing and implementing regulations that protect retail investors from manipulation and from the exploitation of systems in the absence of guardrails.

### **2.1 The black-box problem**

Complex machine-learning models used in trading and credit decisioning can be difficult to explain consistently to regulators, clients, and sometimes the teams that built them, particularly when outputs arise from high-dimensional, non-linear relationships. A 2024 CFA Institute survey of 200 investment-firm employers found that 85% see a need for industry-wide AI standards, and 82% say the lack of standards is slowing adoption.<sup>xi</sup> For a regulator like SEBI, limited explainability complicates model validation and makes it harder to connect system design, human decisions, and resulting market conduct.

### **2.2 Overfitting**

Overfitting occurs when a model learns incidental noise from its training sample rather than durable relationships, producing results that look impressive in backtests but fail in live markets. The canonical illustration is David Leinweber's demonstration that butter production in Bangladesh explained roughly 75% of the in-sample variation in S&P 500 returns, a result that is patently spurious yet statistically significant.<sup>xli</sup> The risk is not confined to contrived examples: a systematic review of 167 reinforcement-learning studies from 2017 to 2025 identifies explainability, robustness in non-stationary environments, deployment feasibility, and the lack of standardized benchmarking as persistent obstacles, and concludes that implementation quality and domain knowledge often outweigh algorithmic complexity.<sup>xlii</sup>

### **2.3 Regime blindness**

Overfitting is a failure of validation: the model fits noise in the sample. Regime blindness is a failure of scope: the model fits the sample correctly, but the world changes. Both are live risks in AI-driven strategies, and they require different controls: out-of-sample and walk-

forward testing for the first, regime-aware modeling and stress testing for the second. Models trained on the 2010-2023 period, a long bull market with episodic shocks, may behave unpredictably in a sustained high-inflation, deglobalizing, or rate-shock regime. The 2022 Treasury selloff was the largest in 40 years, exceeding those seen in 1994, 2003, and 2013, with cumulative returns on a ten-year zero-coupon bond reaching -26.9% by June 2022.<sup>xliii</sup> A regime shift of that magnitude is precisely the condition under which models calibrated on a decade of falling rates would be expected to underperform.

## **2.4 Algorithmic monoculture and herding**

When many funds train on similar datasets and run similar architectures, they may react to the same signals at the same time. The literature now refers to this as ‘algorithmic monoculture’. Kleinberg and Raghavan show, in the context of algorithmic screening decisions, that convergence on a common algorithm can reduce the collective quality of decisions even under normal operations, and even when that algorithm is more accurate for any single user in isolation.<sup>xliv</sup> This leads to homogenization of outcomes and reduced system robustness. A 2026 preprint using 99.5 million SEC Form 13F holdings across 10,957 institutional managers (2013-2024) estimates that AI adoption increases tail-loss amplification by 18-54%, with risk rising super-linearly rather than linearly as adoption rises.<sup>xlv</sup> The Financial Stability Board notes the same mechanism, while also observing that AI could reduce correlations where it enables genuinely differentiated strategies.<sup>vii</sup> The concern is conditional on convergence, not automatic.

The August 2007 ‘Quant Quake’ incident was a major example of a sudden, violent period of deleveraging that hit computer-driven hedge funds (quant funds) just before the broader Great Recession. It is often cited as the first major warning shot that highly automated, mathematical trading could create systemic instability. During the summer of 2007, volatility in the subprime mortgage market was rising. As losses mounted in mortgage-related markets, some large multi-strategy funds are widely believed to have raised liquidity by selling their more liquid holdings, including positions held by their quantitative equity market-neutral businesses, though the specific trigger has never been established. Research by Khandani and Lo<sup>xlvi</sup> indicates that crowded positions, forced deleveraging, liquidity pressure, and common risk-management responses interacted to produce the losses. Many quant funds held similar factor-based positions, long undervalued companies with strong earnings and short overvalued companies with weak earnings. As long positions fell and

short positions rose, risk limits were triggered, further unwinding and reinforcing the price moves. A simulated market-neutral contrarian strategy, run at 8:1 leverage, which the authors argue was typical by then, lost roughly 25% between 6 and 9 August 2007 before recovering 23.67% on 10 August; unleveraged, the same strategy lost 6.85% over 7-9 August. The identity of the initiating trades and the precise causal sequence cannot be established conclusively from public evidence. Still, the episode demonstrates how common exposures and forced deleveraging can generate feedback loops.

## **2.5 Vendor concentration and the ‘Aladdin problem’**

BlackRock’s Aladdin (Asset, Liability, Debt, and Derivative Investment Network) is a widely deployed investment operating system worldwide. It serves the end-to-end investment process across public and private markets for BlackRock’s own USD14.0 trillion of assets under management and for external institutional clients simultaneously, with more than 360,000 users across the Aladdin, Preqin, and eFront platforms.<sup>xlvi</sup> This constitutes a discrete category of systemic risk: a model error or cyberattack in a single vendor system could simultaneously affect many large portfolios, though the scale would depend on clients’ configurations, controls, and redundancies. Third-party concentration of this kind is a failure mode that prudential frameworks, built around the resilience of individual institutions, were not originally designed to address.

The 2024 LockBit ransomware claim against Motilal Oswal, which serves approximately six million clients, illustrates the broader cybersecurity exposure of digitally integrated financial institutions. Motilal Oswal stated that the incident did not affect its business operations or IT environment.<sup>xlvi</sup> The incident was not an AI-system failure, but it demonstrates the importance of cybersecurity, operational segregation, and recovery arrangements as financial institutions become increasingly dependent on interconnected data, cloud, trading, and AI systems.<sup>xli</sup>

## **2.6 Regulatory and accountability gaps**

Existing financial law assigns liability to identifiable persons or institutions. When an autonomous agent built on a foreign foundation model, fine-tuned by a third-party vendor, deployed by a broker, and used by a retail investor results in a catastrophic trade, the liability chain is unclear. SEBI’s February 2025 framework addresses part of this chain for retail algorithmic trading. It treats the stockbroker as principal and the algorithm provider as agent,

assigns responsibilities to brokers, requires algorithms to be identified through exchange-provided tags, and distinguishes white-box algorithms, whose logic is disclosed and replicable, from black-box algorithms, whose logic is not disclosed, thereby requiring black-box providers to register as Research Analysts.<sup>i</sup>

Gaps remain upstream of that framework. They include the allocation of responsibility when a broker or adviser relies on a third-party foundation model; the documentation of model provenance and subsequent modifications; the treatment of AI-generated research that does not itself place orders; and responsibility for failures involving multiple vendors in sequence. Establishing traceable human responsibility for AI-driven decisions should therefore precede any technical standard: without it, no disclosure or audit requirement has an addressee. The upstream question of dependence on foreign foundation models, and the surveillance and security regime that such dependence demands, remains incompletely addressed.

## **2.7 Talent concentration and displacement**

The advent of AI in workplaces introduces a classic structural ‘squeeze’ in the labor market caused by rapid technological adoption. The most sophisticated AI requires data scientists who understand both machine learning and financial markets, a talent pool that is small and globally mobile. If entry-level analytical tasks are automated while demand for specialists rises, the result is a squeeze at both ends. The scale and distribution of any displacement remain uncertain and will vary by firm and function; a further risk is that fewer junior roles weaken the pathway by which future senior analysts acquire domain expertise. These are prospective risks rather than observed industry-wide outcomes, and they need labor-market and skilling responses, not just capital-market policy. The IndiaAI FutureSkills pillar of the IndiaAI Mission<sup>li</sup> addresses part of this; whether it scales fast enough is an empirical question.

## **2.8 Model collapse and synthetic-data pollution**

A prospective concern is synthetic-data contamination, or ‘model collapse’. Research published in *Nature* shows that repeatedly training generative models on data produced by earlier models causes output distributions to narrow and quality to deteriorate.<sup>lii</sup> In finance, a comparable risk arises if future systems are trained or updated on large volumes of unverified

AI-generated commentary, forecasts, or synthetic records. The mechanism is established; its current prevalence in financial applications is not.

## **2.9 Adversarial information manipulation**

A distinct threat, unrelated in mechanism to model collapse, is adversarial information manipulation. Malicious actors can inject false sentiment into social media or issue fake press releases to mislead systems that ingest machine-readable information, potentially triggering automated trading responses. Deliberate data poisoning is a related but narrower attack in which an adversary contaminates the data used to train or update a model. A related risk is already documented: in April 2013, a hacked Associated Press account falsely reported explosions at the White House, and the S&P 500 lost roughly USD 136 billion of its value within about two minutes before recovering once the hack was confirmed.<sup>liii</sup> That episode was an account compromise rather than a poisoned model, but it demonstrates the market sensitivity of machine-readable information. Generative AI lowers the cost and raises the realism of such manipulation. While model collapse passively degrades systems during training, adversarial manipulation actively targets them at chosen moments. The line between data quality and information warfare is thinner than it used to be.

Deepfake fraud represents a related form of deliberate information manipulation. In January 2024, a finance worker at the Hong Kong office of a multinational firm was tricked into transferring USD 25 million after a video call in which the apparent CFO and other colleagues were deepfake recreations.<sup>vi</sup> In India, scammers have used deepfake videos of Sudha Murty and N.R. Narayana Murthy to promote fraudulent AI-powered stock-trading platforms.<sup>liv</sup> A 2023 McAfee study found that 47% of Indian adults had experienced an AI voice-cloning scam or knew someone who had, nearly double the global average of 25%.<sup>lv</sup> This is the inverse problem: AI functioning not as a failing analytical system, but as a deliberately weaponized fraud.

## **3. AI in Investing: Hidden Fragilities**

### **3.1 Hallucination**

‘Hallucination’ is a misleading description of what’s really a specific technical behavior: AI models generate text that is a convincing, coherent response that may be factually incorrect. These systems do not deliberately fabricate information, but their outputs are not inherently

constrained to be true. They generate responses from learned probability distributions conditioned on the prompt and available context. In finance, the consequences are not abstract. The following are some of the classes of hallucination, with examples:

- Regulatory misinterpretation is a credible threat. Large language models can misread regulatory requirements and generate false mandates or exemptions, creating compliance risk and forcing firms to verify non-existent rules. FINOS's AI governance framework gives the illustrative case of a compliance chatbot stating that a low-risk transaction type is exempt from reporting, citing a non-existent clause in the Bank Secrecy Act.<sup>lvi</sup> Fabricated biotech data is a credible investment risk. Generative AI can produce plausible yet false clinical trial data and biomedical claims.<sup>lvii lviii</sup> In pharmaceutical and life sciences due diligence, such outputs could distort investment judgment unless independently verified.
- AI systems can miscalculate financial ratios and misinterpret financial statements, particularly where the analysis requires extracting figures and performing multiple calculations.<sup>lix</sup> This is a class of hallucination in which models could miss critical footnote disclosures or misinterpret balance sheet line items, potentially generating research outputs that claim, for example, that a company has 'zero debt' when it does not. In the FailSafeQA benchmark, which deliberately perturbs queries and supplies degraded, irrelevant, or empty documents, OpenAI's o3-mini ranked most robust of the 24 tested; nonetheless, it fabricated information in 41% of cases.<sup>lx</sup> At the institutional scale, even a lower error rate could lead to flawed analysis if outputs are not independently verified.
- A critical failure pattern is evident in AI-generated quantitative code: large language models can produce scripts that execute successfully while containing look-ahead bias, logic errors, and stateful implementation bugs. Across twelve models tested on a diagnostic benchmark, roughly half produced runnable code that violated temporal causality, and one frontier model's verified accuracy fell from 96% on intermediate tasks to 64% on the most complex ones.<sup>lxi</sup> A related risk is that Python scripts for calculations such as Value-at-Risk may invoke non-existent or inappropriate functions from financial libraries.<sup>lvi</sup> The risk is that some errors will halt execution, while others may go undetected as the code runs and produces plausible but incorrect results. Executable code may therefore be mistaken for valid risk analysis.

- AI systems have generated confident factual errors of the kind that move markets. In February 2023, Google's Bard chatbot, in a promotional demonstration, incorrectly claimed the James Webb Space Telescope had taken the first images of a planet outside our solar system, a feat achieved by the European Southern Observatory's Very Large Telescope in 2004. Alphabet shares fell close to 8% on the day, erasing roughly USD100 billion in market value.<sup>v</sup> The same failure mode applied to a listed company's filings or earnings could be similar in form and potentially harder to detect.

### 3.2 Factors that drive hallucination

- Three common factors increase the risk of AI hallucination:
- First, incomplete, inaccurate, or outdated training data can leave models without reliable factual grounding.<sup>lvi lxii</sup>
- Second, vague or ambiguous prompts increase the likelihood of unsupported interpretations and fabricated details.<sup>lvi</sup>
- Third, large language models generate probable and fluent continuations rather than independently verifying the truth of each claim.<sup>lvi lxiii</sup> A December 2025 information-theoretic detection method, ECLIPSE, reduced the hallucination rate among accepted outputs by 92% relative to an entropy-only detector at 30% coverage on a controlled financial benchmark of 200 samples with synthetic hallucinations, but did not eliminate hallucinations; at 90% coverage, the rate remained 44.4%.<sup>lxiv</sup>

### 3.3 Mitigation approaches

Three protective measures can reduce hallucination risk:

- First, grounding outputs in verified documents. Retrieval-augmented generation, or RAG, supplies the model with external source material and can improve factual reliability, although it does not eliminate hallucinations.<sup>lvi lxv</sup>
- Second, human review. Outputs intended for clients or regulators should be independently reviewed before use.<sup>lxvi</sup>
- Third, cross-checking. Repeated sampling of the same model can reveal inconsistent claims and identify possible hallucinations. However, consistency across responses does not establish that an answer is correct.<sup>lxvii</sup>

These safeguards are taken seriously. In 2023, JPMorgan Chase, Wells Fargo, and Goldman Sachs restricted employee use of consumer AI tools such as ChatGPT due to data security and third-party software concerns.<sup>lxviii</sup> JPMorgan subsequently introduced the LLM Suite, its proprietary generative AI platform that operates in a secure environment.<sup>lxix</sup> Goldman Sachs launched its AI assistant firmwide,<sup>lxx</sup> while Wells Fargo announced the deployment of agentic AI tools across the bank.<sup>lxxi</sup> This reflects a broader shift from unrestricted consumer tools to governed institutional deployment.

**Observation:** Treating LLM output as equivalent to human-authored research is a category error. The current SEBI framework for research analysts requires disclosure of the extent of AI use and holds research analysts responsible for research services based on AI outputs.<sup>lxxii</sup> It does not, however, specifically require disclosure of model limitations, source grounding, or an escalation path when outputs cannot be verified.

## 4. AI in Investing: Finance-specific Systemic Risk

Capital markets are not just another sector in which AI is being deployed. Three properties make AI-related risks unusually acute here: speed (microsecond execution), feedback (prices respond to predictions, then predictions respond to prices), and contagion (losses in one institution propagate through funding and counterparty networks). These properties turn local model errors into system-level events.

### 4.1 Historical precedents for AI-related systemic risk

Four events anchor the policy conversation. They span pre-AI quantitative strategies and classical algorithmic trading, and together they describe what the failure mode looks like across decades.

| Event              | What happened                                                                                                                                                                                                                                                                                         | AI / algorithmic dimension                                                                                                                                                                        |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>LTCM (1998)</b> | Long-Term Capital Management suffered severe losses after market spreads moved against its highly leveraged arbitrage positions. Fourteen banks and brokerage firms received a USD 3.6 billion private-sector recapitalization facilitated by the Federal Reserve Bank of New York. <sup>lxxiii</sup> | Pre-AI cautionary precedent: model assumptions broke down as market relationships diverged, and leverage amplified losses, raising concerns about disorderly liquidation and systemic spillovers. |

| Event                         | What happened                                                                                                                                                                                                                                                                                                                                           | AI / algorithmic dimension                                                                                                                                                                                           |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Quant Quake (Aug 2007)</b> | Multiple quantitative hedge funds simultaneously unwound similar positions over 6-9 August 2007, resulting in days of severe losses for strategies that had been profitable for years, as discussed in detail in 2.4.                                                                                                                                   | The first clear evidence of monoculture risk: when many firms run similar data and models, their trades become correlated, and the exit is too small for all of them to exit at once. <sup>lxxiv</sup>               |
| <b>2010 Flash Crash</b>       | Major equity indices, already down more than 4%, fell a further 5-6% within minutes before recovering almost as rapidly. <sup>lxxv</sup>                                                                                                                                                                                                                | A large automated sell program interacted with high-frequency trading, cross-market feedback, and liquidity withdrawal. The SEC subsequently introduced circuit breakers for individual securities. <sup>lxxvi</sup> |
| <b>Knight Capital (2012)</b>  | A software deployment failure triggered obsolete “Power Peg” code. In approximately 45 minutes, the system generated 4 million executions across 154 stocks and caused a USD 460 million loss <sup>lxxvii</sup> . Knight subsequently secured a USD 400 million capital infusion and later combined with GETCO to form KCG Holdings. <sup>lxxviii</sup> | The event demonstrates how defective code, inadequate deployment controls, and the absence of effective automated risk limits can rapidly threaten a firm and disrupt market prices.                                 |

AI did not cause any of these events in the modern machine-learning sense. They are precedents because they show what can happen when fast, automated, and correlated decision-making meets a market under stress. Modern AI could amplify these dynamics through greater speed, shared technologies, increased automation, and reduced interpretability.

## 4.2 The domino effect: AI as financial contagion

Section 2.4 discussed algorithmic monoculture and herding. The concern here is the next step: how common technologies and correlated decisions can transmit a local failure across institutions and markets. Three channels matter:

- **Performative prediction.** Models do not merely observe markets; their predictions can move prices, and those price movements then enter subsequent decisions and model updates. As more capital follows correlated AI-generated signals, systems may

increasingly respond to the behavior of other systems rather than to underlying fundamentals.<sup>lxxix xlv</sup>

- **Shared-dependency amplification.** When institutions rely on common foundation models, datasets, cloud providers, and technology vendors, an error or outage in one shared component can affect many firms simultaneously. The Financial Stability Board warns that concentration in models, data, and infrastructure could produce correlated outputs, amplify market stress, and exacerbate liquidity pressures.<sup>vii</sup>
- **Cognitive dependency.** Junior analysts, traders, and risk managers who rely heavily on AI-augmented tools may gradually lose some of the underlying skills needed to challenge or override them in unfamiliar conditions. This remains a forward-looking concern rather than an empirically established industry-wide outcome. Emerging work suggests that the systemic risk multiplier may grow superlinearly with increasing AI penetration.<sup>xlv</sup> A non-financial precedent is instructive. In July 2024, a faulty CrowdStrike update disrupted airlines, hospitals, banks, and other critical services worldwide because many organizations depended on the same technology provider.<sup>lxxx lxxxi</sup> AI adoption in finance could create an analogous concentration risk. Existing frameworks address some of these vulnerabilities, but the FSB concludes that they may need to be strengthened as dependence on common models, cloud services, and technology providers increases.<sup>vii</sup>

### **4.3 The Jane Street / Bank Nifty case: coordinated trading and the surveillance challenge**

The Jane Street matter is among the most policy-relevant recent cases in Indian markets. In an interim order dated July 3, 2025, SEBI alleged that four Jane Street entities acted as a group and used coordinated positions across Bank Nifty constituent stocks, futures, and index options to influence index levels and profit from much larger options positions.<sup>xi</sup> Through the order, SEBI directed the impounding of INR 4,843.57 crore in alleged unlawful gains associated with 21 impugned trading days, including three Nifty expiry days in May 2025.

#### ***Why This Case Matters for Policy Makers?***

The order describes activity spread across affiliated entities and linked cash, futures, and options markets. The mechanism SEBI described, involving multi-entity synchronization, and the order patterns it found were structured to evade exchange-level surveillance and

operated in a zone that the existing surveillance regime was not built to detect. SEBI treated the four entities collectively and, on a prima facie basis, found alleged violations of Section 12A of the SEBI Act and of Regulations 3 and 4 of the PFUTP Regulations (Prohibition of Fraudulent and Unfair Trade Practices). Jane Street has denied wrongdoing. After it deposited INR 4,843.57 crore into an escrow account, SEBI lifted the conditional trading restrictions imposed by the interim order.<sup>lxxxii</sup> However, Jane Street told the Securities Appellate Tribunal in April 2026 that it had not resumed trading because of regulatory uncertainty. Its appeal concerning access to documents remained pending, with the next hearing listed for June 23, 2026.<sup>lxxxiii</sup>

The policy point is not to prejudge the legality of any specific trade. It is that coordinated activity across affiliated entities and linked market segments can be difficult to identify and reconstruct quickly. The case concerned complex, high-frequency trading rather than machine learning. It matters here because AI could increase the speed, complexity, and opacity of the same supervisory problem.

## **5. India's AI Policy Posture and Current Regulatory Oversight**

Four major instruments are shaping India's AI governance architecture for finance.

- First, the IndiaAI Mission, approved by the Union Cabinet in March 2024, has an outlay of ~INR 10,372 crore over five years. Its seven pillars include compute capacity, indigenous foundation models, datasets, FutureSkills, startup financing, and Safe & Trusted AI.<sup>ii</sup>
- Second, SEBI issued its circular on “Safer Participation of Retail Investors in Algorithmic Trading” on February 4, 2025.<sup>i</sup> The framework assigns responsibilities to brokers, algo providers, and exchanges; requires unique identifiers for algorithmic orders; and requires providers of black-box algorithms to register as Research Analysts. A subsequent glide path made the framework applicable to all stockbrokers from April 1, 2026.<sup>lxxxiv</sup>
- Third, MeitY (Ministry of Electronics and Information Technology) released the India AI Governance Guidelines on November 5, 2025. The guidelines establish seven Sutras and recommendations across six pillars under a balanced, agile, flexible, and pro-innovation approach.<sup>lxxxv</sup>
- Fourth, on August 13, 2025, the Reserve Bank of India released the report of its Committee on the Framework for Responsible and Ethical Enablement of Artificial

Intelligence in the Financial Sector, chaired by Dr. Pushpak Bhattacharyya of IIT Bombay. The FREE-AI report sets out seven Sutras and 26 actionable recommendations across six strategic pillars.<sup>lxxxvi</sup> It provides a finance-sector governance blueprint alongside SEBI's trading-specific framework, although it remains a committee report rather than a binding regulation. MeitY subsequently adapted the seven Sutras for application across sectors.

Budget execution has lagged behind the initial estimates. The FY2025-26 budget estimate was reduced from INR 2,000 crore to INR 800 crore, while the FY2026-27 budget estimate was reduced to INR 1,000 crore.<sup>lxxxvii</sup> As of February 9, 2026, INR 21.79 crore had been released in FY2024-25 and INR 379.15 crore in FY2025-26, amounting to ~INR 400 crore against the five-year outlay.<sup>lxxxviii</sup> These figures indicate a material gap between the Mission's overall outlay, annual budget estimates, and funds released. They are not, however, a complete measure of implementation: the same official response reports that more than 38,000 GPUs had been onboarded, 12 teams had been shortlisted to develop foundation models, and 30 India-specific applications had been approved.

India's broader AI vision is encapsulated in M.A.N.A.V.: Moral and Ethical Systems, Accountable Governance, National Sovereignty, Accessible and Inclusive, and Valid and Legitimate. Prime Minister Narendra Modi presented this vision at the inauguration of the India AI Impact Summit in New Delhi on February 19, 2026.<sup>lxxxix</sup> The Accountable Governance pillar, combined with SEBI's requirement that providers of black-box algorithms register as Research Analysts, begins to create a framework of traceable responsibility. Enforcement readiness, however, remains an open question.

Taken together, these instruments establish a real but still developing AI governance architecture for finance. They do not yet demonstrate that regulators possess sufficient technical capacity to audit opaque models, detect coordinated cross-entity activity in real time, or allocate responsibility across a technology stack that includes foreign foundation models. The architecture exists. The rulebook exists. Whether its implementation can keep pace with the systems it is intended to govern is the subject of a companion paper.

## 6. The Deeper AI Risk: A Financial Theory Perspective <sup>2</sup>

The Efficient Market Hypothesis holds that asset prices reflect available information, making consistent risk-adjusted outperformance through analysis alone difficult to achieve. In its weak form, prices reflect historical market information; in its semi-strong form, all publicly available information; and in its strong form, private information as well.<sup>xc</sup> On this view, trading by market participants helps move prices toward their informationally efficient level. This premise became an important intellectual foundation for one of the most consequential developments in modern finance: the rise of index funds. The practical implication of this for investors was straightforward. If markets are informationally efficient and active managers cannot systematically outperform, the rational response for most investors is to stop trying.

The case for indexing does not rest on market efficiency alone. Sharpe's "arithmetic of active management" makes a simpler point: because active investors collectively hold the market portfolio, their average return before costs must equal the market return, and their average return after costs must be lower. This accounting identity holds whether or not prices are efficient.<sup>xc</sup> Pedersen qualifies this argument by noting that the market portfolio is not static. Firms continually issue and repurchase shares, and indexes are reconstituted, so even passive investors must trade to maintain market-cap weights. Active managers can therefore, in principle, add value in aggregate, justify positive fees, and play a genuine role in allocating capital.<sup>xcii</sup> This does not overturn Sharpe's cost argument. It suggests that markets still require some active, price-setting capital. If that pool becomes too small to incorporate new information into prices and respond to changes in the market portfolio, the efficiency on which indexing relies may no longer be assured.

Index funds operationalize that logic by offering broad market exposure at minimal cost, without requiring investors to identify mispriced securities themselves. Importantly, the case for indexing does not require the belief that every price is the best available estimate of intrinsic value. Rather, it rests on the more limited proposition that most investors cannot reliably identify and exploit mispricing after costs, with the case further strengthened by diversification, low fees, tax efficiency, and Sharpe's arithmetic. Markets may therefore contain mispricing, while indexing remains rational for the individual investor, provided there is a sufficient number of other participants who continue to perform the price discovery on which that rationality depends. Decades of empirical evidence showing active managers

---

<sup>2</sup> Authors' views and analysis

underperform after fees gave the case for indexing considerable force, and capital migrated accordingly.

Decades of empirical evidence showing active managers underperform after fees<sup>xciii</sup> gave that argument considerable force, and capital migrated accordingly. In the US, passive funds overtook active funds in total net assets in 2023 and, by year-end 2025, accounted for more than 55% of net assets.<sup>xciv</sup> A similar shift is visible in Europe, although passive investing remains less dominant there.<sup>xcv</sup>

The irony, however, is that this migration may gradually weaken the very condition on which it depends. Market efficiency is not self-sustaining. Rather, it depends on a sufficient number of active, fundamentals-driven participants whose research and trading contribute to the price-discovery mechanism by incorporating information into market prices. As capital shifts away from active fundamental analysis toward passive and algorithmically driven strategies, that critical mass may be eroded.

This tension has been examined in the past: the Grossman-Stiglitz Paradox (1980) theoretically established that perfectly efficient markets are self-defeating<sup>xcvi</sup> since rational investors will abandon costly research precisely when prices appear to reflect it fully. Grossman and Stiglitz's framework also contains a corrective mechanism: as prices become less informative, the returns to acquiring information through fundamental research rise, and capital should, in principle, flow back toward informed active management. The mechanism, however, assumes that the resulting informational advantage remains available long enough to justify the cost of producing it. It is less clear whether this still holds at scale when algorithmic and rule-based strategies can rapidly identify and trade on emerging inefficiencies. Crowding and the speed at which algorithmic capital identifies and competes away any emergent inefficiency may blunt the corrective: faster systems may capture the value of new information before fundamental research has time to pay off.

Quantitative, rule-based investing introduces a further distortion. Factor strategies, whether targeting value, momentum, quality, or low volatility, were originally constructed by identifying persistent empirical regularities in historical price and fundamental data. At a modest scale, these strategies can be seen as a systematic way of information processing, broadly consistent with the logic of market efficiency. However, as factor investing has become widely adopted by institutions, collective replication has itself begun to distort the very signals on which these strategies rely. When sufficiently large pools of capital flow

simultaneously into or out of factor exposures, prices can move partly because of flows generated by the strategy itself, rather than any new information about underlying businesses.

Evidence that publication itself changes return patterns supports this concern. McLean and Pontiff find that returns to published cross-sectional predictors decline by roughly 58% after publication. Here, returns are about 26% lower out-of-sample, which they treat as an upper bound estimate of data-mining effects. They are a further 32 % lower after publication, calculated as the 58% post-publication decline less the 26% out-of-sample decline, which the authors estimate reflects investors learning about and trading on the published predictors.<sup>xcvii</sup> At the same time, assets in factor-based strategies have risen by half over the past decade, from about USD 4 trillion to over USD 6 trillion.<sup>xcviii</sup> Crowding can therefore affect both the returns and risks of factor strategies, even when the underlying factor retains a sound economic rationale.

The proliferation of AI-driven trading further compounds this dynamic. Unlike rule-based quant strategies, which operate on defined, interpretable criteria, AI systems optimized for pattern recognition and cross-asset correlation may respond to signals that have no meaningful relationship to business fundamentals. Algorithmic trading, of which machine-learning-based strategies are a growing share, now accounts for an estimated 60-70 percent of equity transaction volumes in the United States and other major global markets.<sup>xcix</sup> At this scale, automated strategies can give model-generated signals and correlated trading greater influence over short-term price formation. However, this does not imply that prices have become detached from business fundamentals.<sup>c</sup> But the risk that prices could increasingly follow model outputs rather than business fundamentals, and momentum signals rather than long-term earnings, is real. As a result, prices may settle not at informationally efficient levels, but at the point where dominant algorithms converge: a consensus generated by the market's own machinery that need not reflect underlying value.

To summarize the preceding discussion, institutional portfolio construction was once anchored more firmly in earnings quality, balance-sheet strength, and long-term business assessment. These disciplines are now increasingly being supplemented, and at times displaced, by passive vehicles, factor replication, and AI-driven systems. Although these approaches differ substantially, their growing scale may reduce the relative role of independent fundamental research and judgment in price discovery.

## 7. Conclusion and Emerging Questions

This paper has argued that the risks arising from AI in investing extend beyond inaccurate outputs, flawed models, hallucination, and operational failure. As AI becomes more prevalent in investment research, portfolio construction, trading, and risk management, the more consequential question is how it affects the process by which information is impounded into asset prices, and therefore how those prices are formed in the first place.

The concern is not that passive, quantitative, or AI-driven strategies ignore economic value. They process large volumes of information and, in many settings, improve efficiency. The concern we highlight is narrower. When large pools of capital rely on similar data, signals, models, and assumptions, the resulting agreement among investment systems can appear to reflect market efficiency, even when it partly reflects the similarity of the systems themselves.

The cumulative effect could be a market in which price discovery increasingly reflects an algorithmic and model-driven consensus rather than independent assessments of value. One reason is that, although passive and active quantitative investing rest on different premises, both can reduce the relative role of independent fundamental research and judgment. ***This possibility raises a hard question for investment organizations and asset owners: if the growing scale of these strategies, together with wider AI use based on similar data and models, begins to alter the conditions on which they depend, do existing frameworks for portfolio construction, manager selection, and governance still hold?***

The paper does not claim that AI has already weakened fundamental price discovery; that evidence is not yet available. What it argues is that the conditions on which price discovery depends are changing, and that the consequences warrant closer examination. Some of the questions that such an examination will need to address are set out below, grouped under three headings.

- *Price formation.* Who sets the marginal price once fundamental investors, passive flows, factor strategies, and AI-driven systems all operate together? And when does algorithmic trading genuinely improve price discovery, as opposed to producing apparent agreement on price simply because different systems are reacting to the same signals?
- *Stability and hidden exposure.* What happens to liquidity and market stability when many systems respond to the same signal, or attempt to reduce risk, at the same

moment? And can portfolios that appear diversified still carry common exposures through shared models, datasets, vendors, or trading assumptions?

- *Accountability.* How should responsibility be allocated when a single investment decision runs through an asset manager, an external model, a data provider, and an automated execution system? And do current governance and regulatory frameworks give supervisors sufficient visibility into model concentration and correlated strategies to detect the problem as it forms?

Several of the above questions connect to the existing literature on market efficiency, price informativeness, and the effects of passive and quantitative flows. What largely remains open is how those findings hold once the strategies in question are AI-driven. Taking these in turn, Gabaix and Koijen examined the question of marginal price. Under their “inelastic markets hypothesis,” marginal prices are highly sensitive to new money entering or leaving the equity market from other assets, such as bonds or cash. Large institutional investors cannot easily absorb these flows because their demand is price-inelastic; strict mandates, benchmarks, and fixed allocations mean they do not usually buy more when prices fall or sell when they rise. The study’s estimates suggest that every USD 1 invested or withdrawn shifts the total market value by about USD 5, implying that capital flows move prices far more than traditional theory predicts.<sup>ci</sup>

Next, whether the growth of passive investing erodes price discovery has also been examined, with genuinely mixed results. Coles, Heath, and Ringgenberg find that more index investing lowers information production but leaves price informativeness unchanged, because the mix of investors adjusts endogenously until the returns to active investing are unchanged.<sup>cii</sup> Sammon, by contrast, finds that higher passive ownership reduces the amount of information impounded into prices ahead of earnings announcements, with the rise in passive ownership over the past three decades accounting for a decline of roughly a quarter of the full-sample mean.<sup>ciii</sup> The concern we raise in this paper is that AI may alter the speed and correlation of the endogenous adjustment described by Coles, Heath, and Ringgenberg, thereby tilting the balance between these two findings. That adjustment depends on how investors behave in aggregate, not on how any single one performs.

Recent work has begun to examine how AI traders behave when many of them operate together, though so far in simulated rather than live markets. In a working paper, Dou, Goldstein, and Ji show that AI traders trained via reinforcement learning can settle into collusive, supra-competitive behavior without any agreement, communication, or intent,

thereby undermining competition and price efficiency.<sup>civ</sup> In another study, Gufler, Sangiorgi, and Tarantino (2026) find that, in multi-agent environments, deep reinforcement learning traders exhibit a “co-learning bias,” whereby the co-evolution of their portfolio policies leads each algorithm to overestimate the price impact of its own trades. This curbs trading aggressiveness and raises the algorithms’ aggregate returns above the Nash-equilibrium benchmark while lowering market efficiency relative to it.<sup>cv</sup>

These studies demonstrate what can happen in modeled environments; they do not provide evidence that the same dynamics are already present in real markets and should thus be interpreted accordingly. Taken together, however, they suggest that the corrective mechanism described by Grossman and Stiglitz could come under greater strain as AI participation expands. While existing evidence indicates that this mechanism can accommodate some degree of free riding, it may struggle when market participation includes a large, fast-moving, and correlated population of learning systems. This is a margin that the earlier literature did not have in view, and one that this paper seeks to highlight rather than resolve.

Moreover, much of the available evidence is drawn from the US or European markets, where market structure and retail participation differ materially from those in India. One relevant comparison comes from Korea, another retail-heavy Asian market. Studying KOSPI order flow from January 2020 to February 2025, Kang shows that during panic-driven retail surges, institutional price impact deteriorates sharply in small-cap stocks but remains resilient in large-caps, making the effects on price discovery highly uneven across firms.<sup>cvi</sup> Although the study concerns human herding and is based on a specific market and period, it suggests that concentrated and correlated trading by AI systems could also affect price discovery unevenly rather than uniformly across the market.

Furthermore, some of the questions we raise are already being addressed in the literature on algorithmic monoculture, factor crowding, and third-party concentration. What remains substantially unexamined is how these strands connect, and what they jointly imply for price discovery in a retail-heavy, derivatives-dominated market such as India, as AI participation expands.

Two clarifications are worth making before closing this discussion. We are not suggesting that AI must be resisted, or that fundamental analysis is likely to disappear. Rather, we make a narrower point: the effects of AI cannot be judged only at the level of a single model or a single institution. What matters, and what is hardest to see from inside any

one firm, is what happens when many systems operate at once. *Whether AI improves individual investment decisions is not the deepest question. A deeper one is whether the growing use of AI will change how markets discover value and, eventually, what market prices represent.*

## **Acknowledgement**

*The authors acknowledge the valuable ground research support provided by the students of the PGDM and PGDM-BM Finance Specialization Batch of 2025–2027 at SPJIMR, Mumbai, for this white paper on AI in Investment. Their efforts in gathering insights, reviewing emerging practices, and contributing to the research foundation of this work are deeply appreciated.*

## Endnotes

---

<sup>i</sup> IMARC Group. n.d. "Algorithmic Trading Market Size, Share & Forecast to 2034." IMARC Group. <https://www.imarcgroup.com/algorithmic-trading-market>.

<sup>ii</sup> Khanna, Aayush. 2025. "Colocation Hits 35.7%, Algo Trading Surges to 53%: The Tech-Driven Shift in NSE." *Investing.com India*, February 10, 2025. <https://in.investing.com/news/stock-market-news/colocation-hits-357-algo-trading-surges-to-53-the-techdriven-shift-in-nse-4655643>.

<sup>iii</sup> Securities and Exchange Board of India. 2024. Updated SEBI Study Reveals 93% of Individual Traders Incurred Losses in Equity F&O between FY22 and FY24; Aggregate Losses Exceed ₹1.8 Lakh Crores Over Three Years. Press Release No. 37/2024. Mumbai: SEBI. September 23, 2024. [https://www.sebi.gov.in/media-and-notifications/press-releases/sep-2024/updated-sebi-study-reveals-93-of-individual-traders-incurred-losses-in-equity-fando-between-fy22-and-fy24-aggregate-losses-exceed-1-8-lakh-crores-over-three-years\\_86906.html](https://www.sebi.gov.in/media-and-notifications/press-releases/sep-2024/updated-sebi-study-reveals-93-of-individual-traders-incurred-losses-in-equity-fando-between-fy22-and-fy24-aggregate-losses-exceed-1-8-lakh-crores-over-three-years_86906.html).

<sup>iv</sup> Godsall, Jonathan, et al. 2025. "How AI Could Reshape the Economics of the Asset Management Industry." McKinsey & Company. July 16, 2025. <https://www.mckinsey.com/industries/financial-services/our-insights/how-ai-could-reshape-the-economics-of-the-asset-management-industry>.

<sup>v</sup> Thorbecke, Catherine. 2023. "Google Shares Lose \$100 Billion after Company's AI Chatbot Makes an Error during Demo." *CNN Business*, February 9, 2023. <https://www.cnn.com/2023/02/08/tech/google-ai-bard-demo-error>.

<sup>vi</sup> Chen, Heather, and Kathleen Magramo. 2024. "Finance Worker Pays out \$25 Million after Video Call with Deepfake 'Chief Financial Officer.'" *CNN*, February 4, 2024. <https://www.cnn.com/2024/02/04/asia/deepfake-cfo-scam-hong-kong-intl-hnk>

<sup>vii</sup> Financial Stability Board. 2024. *The Financial Stability Implications of Artificial Intelligence*. Basel: Financial Stability Board. November 14, 2024. <https://www.fsb.org/uploads/P14112024.pdf>.

<sup>viii</sup> JioBlackRock Investment Advisers. 2026. "JioBlackRock Investment Advisers Launches 'Personalised Investment Advice' to Democratise Wealth Management for the People of India." Press release, February 3, 2026. <https://www.jioblackrock.com/press-releases/jioblackrock-investment-advisers-launches-personalised-investment-advice-to-democratise-wealth-management-for-the-people-of-india>.

<sup>ix</sup> Press Trust of India. 2025. "Retail Participation in Markets Increases; Demat Accounts Surge to 194 mn: Sebi." *Business Standard*, July 15, 2025. [https://www.business-standard.com/markets/news/retail-participation-in-mkts-increase-demat-accounts-surge-to-194-mn-sebi-125071500934\\_1.html](https://www.business-standard.com/markets/news/retail-participation-in-mkts-increase-demat-accounts-surge-to-194-mn-sebi-125071500934_1.html).

<sup>x</sup> Hyde, Kirsten. 2024. "The Meteoric Rise in India's Equity Derivatives Volume (And What Regulators Plan to Do about It)." *FIA MarketVoice*, August 16, 2024. <https://www.fia.org/marketvoice/articles/explainer-meteoric-rise-indias-equity-derivatives-volume>.

<sup>xi</sup> Securities and Exchange Board of India. 2025. *Interim Order in the Matter of Index Manipulation by Jane Street Group*. Order No. WTM/AN/MRD/MRD-SEC-3/31516/2025-26. Mumbai: SEBI. July 3, 2025. [https://www.sebi.gov.in/enforcement/orders/jul-2025/interim-order-in-the-matter-of-index-manipulation-by-jane-street-group\\_95040.html](https://www.sebi.gov.in/enforcement/orders/jul-2025/interim-order-in-the-matter-of-index-manipulation-by-jane-street-group_95040.html).

- <sup>xii</sup> Das, Prabhanu Kumar. 2026. "IndiaAI Mission: Only 400 Crore of the over Rs 10,000 Crore 5-Year Outlay Has Been Released." Medianama, April 4, 2026. <https://www.medianama.com/2026/04/223-indiaai-mission-400-crore-over-rs-10000-crore-5-year-outlay-released/>.
- <sup>xiii</sup> Brush, Silla. 2025. "BlackRock Has Built an AI Analyst, 'Asimov,' to Scan Filings." Bloomberg News, June 12, 2025, via *Financial Advisor Magazine*. <https://www.fa-mag.com/news/blackrock-has-built-an-ai-analyst---asimov---to-scan-filings-82891.html>.
- <sup>xiv</sup> BlackRock, Inc. 2026. *2025 Annual Report*. New York: BlackRock. [https://s24.q4cdn.com/856567660/files/doc\\_financials/2026/ar/BLK\\_AR25.pdf](https://s24.q4cdn.com/856567660/files/doc_financials/2026/ar/BLK_AR25.pdf).
- <sup>xv</sup> Son, Hugh. 2017. "JPMorgan Software Does in Seconds What Took Lawyers 360,000 Hours." *Bloomberg*, February 28, 2017. <https://www.bloomberg.com/news/articles/2017-02-28/jpmorgan-marshals-an-army-of-developers-to-automate-high-finance>.
- <sup>xvi</sup> Rooney, Kate. 2025. "Goldman Sachs Autonomous Coder Pilot Marks Major AI Milestone." *CNBC*, July 11, 2025. <https://www.cnbc.com/2025/07/11/goldman-sachs-autonomous-coder-pilot-marks-major-ai-milestone.html>.
- <sup>xvii</sup> Patri, Anirudh. 2025. "How Goldman Sachs Deployed Its AI Platform." Nanonets Blog, August 13, 2025. <https://nanonets.com/blog/goldman-sachs-ai-platform/>.
- <sup>xviii</sup> Vanguard. 2024. "Vanguard Expands Access to Advice by Reducing Asset Minimum for Robo Service, Digital Advisor." Vanguard Press Release. September 4, 2024. <https://corporate.vanguard.com/content/corporatesite/us/en/corp/who-we-are/pressroom/press-release-vanguard-expands-access-to-advice-by-reducing-asset-minimum-for-robo-service-digital-advisor-090424.html>.
- <sup>xix</sup> Calvery, Jennifer. 2024. "Harnessing the Power of AI to Fight Financial Crime." HSBC News and Views, June 10, 2024. <https://www.hsbc.com/news-and-views/views/hsbc-views/harnessing-the-power-of-ai-to-fight-financial-crime>.
- <sup>xx</sup> Karlsson, Kasper. 2024. "Renaissance Technologies and the Medallion Fund." *Quartr*, March 20, 2024, updated November 21, 2025. <https://quartr.com/insights/edge/renaissance-technologies-and-the-medallion-fund>.
- <sup>xxi</sup> HedgeCo Insights. "Two Sigma's 'AI-First' Internal Mandate — The Race for 'Operational Alpha' in the Age of Frontier Models." April 3, 2026. <https://hedgeco.net/news/04/2026/two-sigmas-ai-first-internal-mandate-the-race-for-operational-alpha-in-the-age-of-frontier-models.html>.
- <sup>xxii</sup> Acadian Asset Management. n.d. "Equity Strategies." Acadian Asset Management. Accessed 2026. <https://www.acadian-asset.com/investment-strategies/equity>.
- <sup>xxiii</sup> *Business Today*, "JioBlackRock Debuts First Active Equity Fund with AI-Driven Flexi Cap Strategy; Key Takeaways for Investors," September 25, 2025, <https://www.businesstoday.in/mutual-funds/story/jioblackrock-debuts-first-active-equity-fund-with-ai-driven-flexi-cap-strategy-key-takeaways-for-investors-495681-2025-09-25>.
- <sup>xxiv</sup> JioBlackRock Mutual Fund. n.d. "JioBlackRock Flexi Cap Fund." JioBlackRock Asset Management. Accessed August 2026. <https://www.jioblackrockamc.com/funds/jioblackrock-flexi-cap-fund>.
- <sup>xxv</sup> Trendlyne. n.d. "AI Summary: Stock Insights from Earnings Calls." Trendlyne. <https://trendlyne.com/ai-conference-calls/>.
- <sup>xxvi</sup> Tijori Finance. n.d. "Features." Tijori Finance. <https://www.tijorifinance.com/features/>.
- <sup>xxvii</sup> Zerodha. n.d. "What Is Nudge?" Zerodha Support. Accessed 2026. <https://support.zerodha.com/category/trading-and-markets/alerts-and-nudges/nudges/articles/what-is-nudge>.
- <sup>xxviii</sup> Zerodha. n.d. "What Is Kill Switch?" Zerodha Support. <https://support.zerodha.com/category/console/segments/killswitch/articles/what-is-the-kill-switch>.
- <sup>xxix</sup> Abrar, Peerzada. 2026. "Groww Builds AI-Powered Platform across Trading, Wealth and Fixed Income." *Business Standard*, February 28, 2026. [https://www.business-standard.com/companies/news/groww-builds-ai-powered-platform-across-trading-wealth-and-fixed-income-126022800536\\_1.html](https://www.business-standard.com/companies/news/groww-builds-ai-powered-platform-across-trading-wealth-and-fixed-income-126022800536_1.html).
- <sup>xxx</sup> Fama, Eugene F., and Kenneth R. French. 1993. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics* 33, no. 1: 3–56. [https://doi.org/10.1016/0304-405X\(93\)90023-5](https://doi.org/10.1016/0304-405X(93)90023-5).
- <sup>xxxi</sup> CFA Institute Research and Policy Center. 2025. "Agentic AI for Finance: Workflows, Tips, and Case Studies." CFA Institute. 2025. <https://rpc.cfainstitute.org/research/the-automation-ahead-content-series/agentic-ai-for-finance>.
- <sup>xxxii</sup> World Economic Forum and Accenture. 2025. "Artificial Intelligence in Financial Services." World Economic Forum White Paper. January 2025. [https://reports.weforum.org/docs/WEF\\_Artificial\\_Intelligence\\_in\\_Financial\\_Services\\_2025.pdf](https://reports.weforum.org/docs/WEF_Artificial_Intelligence_in_Financial_Services_2025.pdf).
- <sup>xxxiii</sup> PwC. 2017. *Sizing the Prize: PwC's Global Artificial Intelligence Study — Exploiting the AI Revolution*. PricewaterhouseCoopers. <https://www.pwc.com.au/government/pwc-ai-analysis-sizing-the-prize-report.pdf>.
- <sup>xxxiv</sup> Zoting, Shivani. 2025. "AI in Asset Management Market Size to Attain USD 38.94 Bn by 2034." Precedence Research, July 14, 2025. <https://www.precedenceresearch.com/ai-in-asset-management-market>.
- <sup>xxxv</sup> Sozzi, Brian. 2026. "Meta, Microsoft, Amazon, and Alphabet Are about to Spend a Shocking Amount of Money to Dominate the AI Era." *Yahoo Finance*, June 3, 2026.

<https://finance.yahoo.com/sectors/technology/article/meta-microsoft-amazon-and-alphabet-are-about-to-spend-a-shocking-amount-of-money-to-dominate-the-ai-era-115359575.html>.

<sup>xxxvi</sup> Khanna, Aayush. 2025. "Colocation Hits 35.7%, Algo Trading Surges to 53%: The Tech-Driven Shift in NSE." *Investing.com India*, February 10, 2025. <https://in.investing.com/news/stock-market-news/colocation-hits-357-algo-trading-surges-to-53-the-techdriven-shift-in-nse-4655643>.

<sup>xxxvii</sup> Association of Mutual Funds in India. 2025. *AMFI Annual Report Fiscal 2025*. Mumbai: AMFI. [https://www.amfiindia.com/Themes/Theme1/downloads/AMFI\\_AnnualMFRReport2025.pdf](https://www.amfiindia.com/Themes/Theme1/downloads/AMFI_AnnualMFRReport2025.pdf).

<sup>xxxviii</sup> Sen, Joydeep. 2025. "Investments in Mutual Funds: The Other Perspective." *Outlook Money*, December 11, 2025, updated December 13, 2025. <https://www.outlookmoney.com/invest/investments-in-mutual-funds-the-other-perspective>.

<sup>xxxix</sup> Bagattini, Giulio, Nassim Khelifi, and Valentina Mejdahl. 2026. *AI Adoption and Trends in Securities Markets: EU Evidence*. ESMA TRV Risk Analysis, ESMA50-481369926-30599. Paris: European Securities and Markets Authority. February 20, 2026. [https://www.esma.europa.eu/sites/default/files/2026-02/ESMA50-481369926-30599\\_TRV\\_Risk\\_Analysis\\_AI\\_adoption\\_and\\_trends\\_in\\_securities\\_markets.pdf](https://www.esma.europa.eu/sites/default/files/2026-02/ESMA50-481369926-30599_TRV_Risk_Analysis_AI_adoption_and_trends_in_securities_markets.pdf).

<sup>xl</sup> CFA Institute. 2024. "New CFA Institute Survey Reveals Industry-Wide Need for Standards and Education on AI/GenAI in Investing." Press release, August 27, 2024. <https://www.cfainstitute.org/about/press-room/2024/ai-in-investment-sector-survey>.

<sup>xli</sup> Leinweber, David J. 2007. "Stupid Data Miner Tricks: Overfitting the S&P 500." *The Journal of Investing* 16 (1): 15–22. <https://doi.org/10.3905/joi.2007.681820>.

<sup>xlii</sup> Hoque, Mohammad Rezoanul, Md Meftahul Ferdous, and M. Kabir Hassan. 2025. "Reinforcement Learning in Financial Decision Making: A Systematic Review of Performance, Challenges, and Implementation Strategies." arXiv preprint arXiv:2512.10913. December 2025. <https://arxiv.org/abs/2512.10913>.

<sup>xliii</sup> Adrian, Tobias, and Michael Fleming. 2022. "The Bond Market Selloff in Historical Perspective." *Liberty Street Economics* (Federal Reserve Bank of New York). July 5, 2022. <https://libertystreeteconomics.newyorkfed.org/2022/07/the-bond-market-selloff-in-historical-perspective/>.

<sup>xliv</sup> Kleinberg, Jon, and Manish Raghavan. 2021. "Algorithmic Monoculture and Social Welfare." *Proceedings of the National Academy of Sciences* 118, no. 22: e2018340118. <https://doi.org/10.1073/pnas.2018340118>.

<sup>xlv</sup> Meng, Shuchen, and Xupeng Chen. 2026. "Artificial Intelligence and Systemic Risk." arXiv preprint arXiv:2604.03272. March 23, 2026. <https://arxiv.org/abs/2604.03272>.

<sup>xlvi</sup> Khandani, Amir E., and Andrew W. Lo. 2011. "What Happened to the Quants in August 2007? Evidence from Factors and Transactions Data." *Journal of Financial Markets* 14, no. 1: 1–46. <https://doi.org/10.1016/j.finmar.2010.07.005>.

<sup>xlvii</sup> BlackRock, Inc. 2026. 2025 Annual Report. New York: BlackRock. [https://s24.q4cdn.com/856567660/files/doc\\_financials/2026/ar/BLK\\_AR25.pdf](https://s24.q4cdn.com/856567660/files/doc_financials/2026/ar/BLK_AR25.pdf).

<sup>xlviii</sup> Tiwari, Khushboo. 2024. "Data Breach: No Impact on Biz Ops, Says Motilal Oswal Financial Services." *Business Standard*, February 19, 2024. [https://www.business-standard.com/companies/news/no-impact-on-business-operations-motilal-oswal-financial-on-data-breach-124021901020\\_1.html](https://www.business-standard.com/companies/news/no-impact-on-business-operations-motilal-oswal-financial-on-data-breach-124021901020_1.html).

<sup>xlix</sup> Singh, Jagmeet. 2024. "LockBit Claims Cyberattack on Indian Broker Motilal Oswal." *TechCrunch*, February 15, 2024. <https://techcrunch.com/2024/02/15/lockbit-ransomware-cyberattack-india-brokerage-firm-motilal-oswal/>.

<sup>l</sup> Securities and Exchange Board of India. 2025a. *Safer Participation of Retail Investors in Algorithmic Trading*. Circular No. SEBI/HO/MIRSD/MIRSD-PoD/P/CIR/2025/0000013. Mumbai: SEBI. February 4, 2025. [https://www.sebi.gov.in/legal/circulars/feb-2025/safer-participation-of-retail-investors-in-algorithmic-trading\\_91614.html](https://www.sebi.gov.in/legal/circulars/feb-2025/safer-participation-of-retail-investors-in-algorithmic-trading_91614.html).

<sup>li</sup> Press Information Bureau. 2024. "Cabinet Approves Ambitious IndiaAI Mission to Strengthen the AI Innovation Ecosystem." Ministry of Electronics and Information Technology, Government of India. March 2024. <https://www.pib.gov.in/PressReleaseIframePage.aspx?PRID=2012355>.

<sup>lii</sup> Shumailov, Ilia, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal. 2024. "AI Models Collapse When Trained on Recursively Generated Data." *Nature* 631 (8022): 755–759. <https://doi.org/10.1038/s41586-024-07566-y>.

<sup>liii</sup> Wang, Lu, Whitney Kislring, and Eric Lam. 2013. "Fake Post Erasing \$136 Billion Shows Markets Need Humans." *Bloomberg*, April 23, 2013. <https://www.bloomberg.com/news/2013-04-23/fake-report-erasing-136-billion-shows-market-s-fragility.html>.

<sup>liv</sup> Gupta, Shivangini. 2025. "Trust NR Narayana Murthy? First Check if It's an AI Video." *Outlook Business*, October 1, 2025, updated October 7, 2025. <https://www.outlookbusiness.com/magazine/trust-nr-narayana-murthy-first-check-if-its-an-ai-video>.

<sup>lv</sup> McAfee. 2023. "Beware the Artificial Impostor: A McAfee Cybersecurity Artificial Intelligence Report." May 2023. <https://www.mcafee.com/content/dam/consumer/en-us/resources/cybersecurity/artificial-intelligence/rp-beware-the-artificial-impostor-report.pdf>.

- <sup>lvi</sup> FINOS (Fintech Open Source Foundation). n.d. "Hallucination and Inaccurate Outputs." AI Governance Framework for Financial Services. [https://air-governance-framework.finos.org/risks/ri-4\\_hallucination-and-inaccurate-outputs.html](https://air-governance-framework.finos.org/risks/ri-4_hallucination-and-inaccurate-outputs.html).
- <sup>lvii</sup> Naddaf, Miryam. 2023. "ChatGPT Generates Fake Data Set to Support Scientific Hypothesis." *Nature* 623: 895–896. November 22, 2023. <https://doi.org/10.1038/d41586-023-03635-w>.
- <sup>lviii</sup> Taloni, Andrea, Giulia Coco, Marco Pellegrini, et al. 2025. "Exploring Detection Methods for Synthetic Medical Datasets Created With a Large Language Model." *JAMA Ophthalmology* 143 (6): 517–522. <https://doi.org/10.1001/jamaophthalmol.2025.0834>
- <sup>lix</sup> Wang, Xinlin, and Mats Brorsson. 2025. "Can Large language model analyze financial statements well?" In *Proceedings of the Joint Workshop of the 9th Financial Technology and Natural Language Processing (FinNLP), the 6th Financial Narrative Processing (FNP), and the 1st Workshop on Large Language Models for Finance and Legal (LLMFinLegal)*, 196–206. Abu Dhabi, UAE: Association for Computational Linguistics. January 2025. <https://aclanthology.org/2025.finnlp-1.19/>.
- <sup>lx</sup> Kamble, Kiran, Melisa Russak, Dmytro Mozolevskyi, Muayad Ali, Mateusz Russak, and Waseem AlShikh. 2025. "Expect the Unexpected: FailSafe Long Context QA for Finance." arXiv preprint arXiv:2502.06329. February 10, 2025. <https://arxiv.org/abs/2502.06329>.
- <sup>lxi</sup> Luo, Sichun, Yi Huang, Shichang Meng, Fengyuan Liu, Mukai Li, Qinghua Yao, Zefa Hu, Junlan Feng, and Qi Liu. 2026. "AlphaQT-Bench: Diagnosing the Gap between Financial Code Generation and Quantitative Reasoning in LLMs." Findings of the Association for Computational Linguistics: ACL 2026, 2871–2886. July 2026. <https://doi.org/10.18653/v1/2026.findings-acl.138>.
- <sup>lxii</sup> Huang, Lei, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2025. "A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions." *ACM Transactions on Information Systems* 43 (2): 1–55. <https://doi.org/10.1145/3703155>.
- <sup>lxiii</sup> Huang, Lei, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2025. "A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions." *ACM Transactions on Information Systems* 43 (2): 1–55. <https://doi.org/10.1145/3703155>.
- <sup>lxiv</sup> Singha, Mainak. 2025. "Detecting AI Hallucinations in Finance: An Information-Theoretic Method Cuts Hallucination Rate by 92%." arXiv preprint arXiv:2512.03107. December 2, 2025. <https://arxiv.org/abs/2512.03107>.
- <sup>lxv</sup> Gao, Yunfan, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2023. "Retrieval-Augmented Generation for Large Language Models: A Survey." arXiv preprint arXiv:2312.10997, revised March 2024. <https://arxiv.org/abs/2312.10997>.
- <sup>lxvi</sup> Autio, Chloe, Reva Schwartz, Jesse Dunietz, Shomik Jain, Martin Stanley, Elham Tabassi, Patrick Hall, and Kamie Roberts. 2024. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. NIST AI 600-1. Gaithersburg, MD: National Institute of Standards and Technology. July 26, 2024. <https://doi.org/10.6028/NIST.AI.600-1>.
- <sup>lxvii</sup> Manakul, Potsawee, Adian Liusie, and Mark Gales. 2023. "SelfCheckGPT: Zero-Resource Black-Box Hallucination Detection for Generative Large Language Models." In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 9004–17. Singapore: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.emnlp-main.557>.
- <sup>lxviii</sup> Bushard, Brian. 2023. "Workers' ChatGPT Use Restricted at More Banks—Including Goldman, Citigroup." *Forbes*, February 24, 2023. <https://www.forbes.com/sites/brianbushard/2023/02/24/workers-chatgpt-use-restricted-at-more-banks-including-goldman-citigroup/>.
- <sup>lxix</sup> JPMorgan Chase. 2025. "LLM Suite Named 2025 'Innovation of the Year' by American Banker." June 3, 2025. <https://www.jpmorganchase.com/about/technology/blog/llmsuite-ab-award>.
- <sup>lxx</sup> Azhar, Saeed, and Pritam Biswas. 2025. "Goldman Sachs Launches AI Assistant Firmwide, Memo Shows." *Reuters*, June 23, 2025. <https://www.reuters.com/business/goldman-sachs-launches-ai-assistant-firmwide-memo-shows-2025-06-23/>.
- <sup>lxxi</sup> Wells Fargo. 2025. "Wells Fargo Names Saul Van Beurden to Lead Artificial Intelligence; Kleber Santos to Expand Responsibilities and Serve as Co-CEO of Consumer Banking and Lending with Van Beurden." Press release, November 20, 2025. <https://newsroom.wf.com/news-releases/news-details/2025/Wells-Fargo-Names-Saul-Van-Beurden-to-Lead-Artificial-Intelligence-Kleber-Santos-to-Expand-Responsibilities-and-Serve-as-Co-CEO-of-Consumer-Banking-and-Lending-with-Van-Beurden/default.aspx>.
- <sup>lxxii</sup> Securities and Exchange Board of India. 2026. "Master Circular for Research Analysts." February 6, 2026. [https://www.sebi.gov.in/legal/master-circulars/feb-2026/master-circular-for-research-analysts\\_99571.html](https://www.sebi.gov.in/legal/master-circulars/feb-2026/master-circular-for-research-analysts_99571.html).
- <sup>lxxiii</sup> Federal Reserve History. n.d. "Near Failure of Long-Term Capital Management." Federal Reserve History. <https://www.federalreservehistory.org/essays/lbcm-near-failure>.

- lxxiv Goldman Sachs Asset Management. "The 'Quant Quake' of August 2007: Ten Years Later." 2017. <https://www.gsam.com/content/dam/gsam/pdfs/common/en/public/articles/2017/quant-quake-reprint-news.pdf>
- lxxv U.S. Commodity Futures Trading Commission and U.S. Securities and Exchange Commission. 2010. "Findings Regarding the Market Events of May 6, 2010." September 30, 2010. <https://www.sec.gov/files/marketevents-report.pdf>.
- lxxvi Federal Reserve Bank of New York. 2012. "The Flash Crash: Two Years On." Liberty Street Economics, May 7, 2012. <https://libertystreeteconomics.newyorkfed.org/2012/05/the-flash-crash-two-years-on/>.
- lxxvii U.S. Securities and Exchange Commission. 2013. "In the Matter of Knight Capital Americas LLC." SEC Release No. 34-70694. October 16, 2013. <https://www.sec.gov/litigation/admin/2013/34-70694.pdf>.
- lxxviii Knight Capital Group. 2012. "Knight Capital Group Completes \$400 Million Equity Financing Agreement." Press release, August 6, 2012. <https://www.sec.gov/Archives/edgar/data/1060749/000119312512338098/d392396dex991.htm>.
- lxxix Danielsson, Jón, and Andreas Uthemann. 2025. "Artificial Intelligence and Financial Crises." *Journal of Financial Stability* 80: 101453. <https://doi.org/10.1016/j.jfs.2025.101453>.
- lxxx CrowdStrike. 2024. "Preliminary Post Incident Review: Content Configuration Update Impacting the Falcon Sensor and the Windows Operating System." July 24, 2024. <https://www.crowdstrike.com/en-us/blog/falcon-content-update-preliminary-post-incident-report/>.
- lxxxi Weston, David. 2024. "Helping Our Customers through the CrowdStrike Outage." The Official Microsoft Blog, July 20, 2024. <https://blogs.microsoft.com/blog/2024/07/20/helping-our-customers-through-the-crowdstrike-outage/>.
- lxxxii Securities and Exchange Board of India. 2025. "Update on the Jane Street Interim Order." Press Release No. 45/2025. July 21, 2025. [https://www.sebi.gov.in/media-and-notifications/press-releases/jul-2025/update-on-the-jane-street-interim-order\\_95491.html](https://www.sebi.gov.in/media-and-notifications/press-releases/jul-2025/update-on-the-jane-street-interim-order_95491.html).
- lxxxiii Kumar, Brajesh. 2026. "'Not Sure What to Trade and What Not, So Stopped Trading,' Jane Street Tells SAT; Next Hearing on June 23." Moneycontrol, April 22, 2026. <https://www.moneycontrol.com/news/business/markets/not-sure-what-to-trade-and-what-not-so-stopped-trading-jane-street-tells-sat-final-hearing-on-june-23-13896446.html>.
- lxxxiv Securities and Exchange Board of India. 2025b. Extension of Timeline for Implementation of SEBI Circular dated February 04, 2025 on 'Safer Participation of Retail Investors in Algorithmic Trading'. Circular No. SEBI/HO/MIRSD/MIRSD-PoD/P/CIR/2025/132. Mumbai: SEBI. September 30, 2025. <https://www.sebi.gov.in/legal/circulars/sep-2025/extension-of-timeline-for-implementation-of-sebi-circular-dated-february-04-2025-on-safer-participation-of-retail-investors-in-algorithmic-trading-96979.html>.
- lxxxv Ministry of Electronics and Information Technology. 2025. "MeitY Unveils India AI Governance Guidelines to Promote Safe and Responsible AI Adoption." Government of India. November 5, 2025. <https://www.newsonair.gov.in/meity-unveils-india-ai-governance-guidelines-to-promote-safe-and-responsible-ai-adoption>.
- lxxxvi Reserve Bank of India. 2025. Report of the Committee to Develop a Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI) in the Financial Sector. Mumbai: RBI. August 13, 2025. <https://rbidocs.rbi.org.in/rdocs/PublicationReport/Pdfs/FREEAIR130820250A24FF2D4578453F824C72ED9F5D5851.PDF>.
- lxxxvii Reddy, Bharath. 2026. "Decoding the Budget for India's AI Priorities." Takshashila Institution, February 3, 2026. [https://takshashila.org.in/content/blogs/20260203\\_ai-mission-budget-analysis.html](https://takshashila.org.in/content/blogs/20260203_ai-mission-budget-analysis.html).
- lxxxviii Ministry of Electronics and Information Technology. 2026. "Progress of IndiaAI Mission." Rajya Sabha Unstarred Question No. 1668. February 13, 2026. [https://sansad.in/getFile/annex/270/AU1668\\_ajohnN.pdf?source=pqars](https://sansad.in/getFile/annex/270/AU1668_ajohnN.pdf?source=pqars).
- lxxxix Prime Minister's Office. 2026. "PM's Address at India AI Impact Summit 2026." Government of India. February 19, 2026. [https://www.pmindia.gov.in/en/news\\_updates/pms-address-at-india-ai-impact-summit-2026/](https://www.pmindia.gov.in/en/news_updates/pms-address-at-india-ai-impact-summit-2026/).
- xc Fama, Eugene F. 1970. "Efficient Capital Markets: A Review of Theory and Empirical Work." *The Journal of Finance* 25 (2): 383–417. <https://doi.org/10.1111/j.1540-6261.1970.tb00518.x>.
- xci Sharpe, William F. 1991. "The Arithmetic of Active Management." *Financial Analysts Journal* 47 (1): 7–9. <https://doi.org/10.2469/faj.v47.n1.7>.
- xcii Pedersen, Lasse Heje. 2018. "Sharpening the Arithmetic of Active Management." *Financial Analysts Journal* 74 (1): 21–36. <https://doi.org/10.2469/faj.v74.n1.4>.
- xciii Armour, Bryan, Eugene Gorbatikov, and Maciej Kowara. 2026. "US Active/Passive Barometer Report: Year-End 2025." Morningstar. February 18, 2026. <https://www.morningstar.com/business/insights/research/active-passive-barometer>.

- <sup>xciv</sup> Carter, Drew, and Joe Bullard. 2026. "8 Charts on US Fund Flows in 2025: Highest Net Inflows Since 2021." Morningstar. January 16, 2026. <https://www.morningstar.com/funds/8-charts-us-fund-flows-2025-highest-net-inflows-since-2021>.
- <sup>xcv</sup> Garcia-Zarate, Jose. 2026. "Passively Managed Funds Outperform Active Peers Across Most Categories." Morningstar, February 25, 2026. <https://global.morningstar.com/en-gb/funds/passively-managed-funds-outperform-active-peers-across-most-categories>.
- <sup>xcvi</sup> Grossman, Sanford J., and Joseph E. Stiglitz. 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70 (3): 393–408. <https://www.aeaweb.org/aer/top20/70.3.393-408.pdf>.
- <sup>xcvii</sup> McLean, R. David, and Jeffrey Pontiff. 2016. "Does Academic Research Destroy Stock Return Predictability?" *The Journal of Finance* 71 (1): 5–32. <https://doi.org/10.1111/jofi.12365>.
- <sup>xcviii</sup> Gupta, Abhishek, and Stuart Doole. 2025. Factor Indexing Through the Decades: Lessons from 50 Years, Today's Challenges and Future Directions. MSCI Research Insights, July 2025. <https://www.msci.com/downloads/web/msci-com/research-and-insights/paper/factor-indexing-through-the-decades/factor-indexing-through-the-decades.pdf>.
- <sup>xcix</sup> Anand, Kartik, Sophia Kazinnik, Agnese Leonello, and Ettore Panetti. 2026. "Financial Stability in the Age of Artificial Intelligence: The Role of Algorithmic Architecture." ECB Research Bulletin, no. 143, May 21, 2026. <https://www.ecb.europa.eu/press/research-publications/resbull/2026/html/ecb.rb260521~4d8b12940b.en.html>.
- <sup>c</sup> International Monetary Fund. 2024. "Advances in Artificial Intelligence: Implications for Capital Market Activities." In *Global Financial Stability Report, October 2024: Steadying the Course: Uncertainty, Artificial Intelligence, and Financial Stability*. Washington, DC: International Monetary Fund. <https://doi.org/10.5089/9798400277573.082>.
- <sup>ci</sup> Gabaix, Xavier, and Ralph S. J. Koijen. 2021. "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis." NBER Working Paper 28967. <https://doi.org/10.3386/w28967>.
- <sup>cii</sup> Coles, Jeffrey L., Davidson Heath, and Matthew C. Ringgenberg. 2022. "On Index Investing." *Journal of Financial Economics* 145 (3): 665–683. <https://doi.org/10.1016/j.jfineco.2022.05.007>.
- <sup>ciii</sup> Sammon, Marco. 2025. "Passive Ownership and Price Informativeness." *Management Science* 71 (6): 4582–4598. <https://doi.org/10.1287/mnsc.2023.00836>.
- <sup>civ</sup> Dou, Winston Wei, Itay Goldstein, and Yan Ji. 2025. "AI-Powered Trading, Algorithmic Collusion, and Price Efficiency." NBER Working Paper 34054. <https://doi.org/10.3386/w34054>.
- <sup>cv</sup> Gufler, I., Sangiorgi, F., & Tarantino, E. (2026). *(Deep) learning to trade: An analysis of AI trading and market outcomes* [Working paper]. SSRN. <https://doi.org/10.2139/ssrn.5375160>
- <sup>cvi</sup> Kang, Sungwoo. 2026. "The Physics of Price Discovery: Deconvolving Information, Volatility, and the Critical Breakdown of Signal during Retail Herding." arXiv preprint arXiv:2601.11602. February 25, 2026. <https://arxiv.org/abs/2601.11602>.