Reconstructing Algorithmic Trading in the Age of Generative AI: Implications for Market Structure, Regulation, and Epistemology

Reconstruyendo el *algorithmic trading* en la era de la inteligencia artificial generativa: implicaciones para la estructura del mercado, la regulación y la epistemología

Ligia Catherine Arias-Barrera[1]

ORCID Id: https://orcid.org/0009-0003-3085-1922

Ph.D en Derecho Financiero (University of Warwick)

Docente. Universidad Externado de Colombia (Colombia)

Fecha de recepción: Marzo 26, 2025 **Fecha de aceptación**: Noviembre 18, 2025

Received: March 26, 2025 **Accepted**: November 18, 2025

Artículo de investigación. DOI: https://doi.org/10.18601/16923960.v25n1.08

RESUMEN

El propósito de este artículo es analizar cómo la Inteligencia Artificial Generativa (Gen AI), definida aquí como sistemas capaces de producir resultados novedosos, como texto, código o datos, mediante aprendizaje automático,

1 Abogada y Especialista en Derecho Comercial del Externado, LLM in Commercial and Corporate Law por Queen Mary University of London y Ph.D. en Derecho Financiero por la University of Warwick. Autora de Regulation and Supervision of the OTC Derivatives Market (2018) y The Law of ESG Derivatives (2024), es Profesora de Derecho Financiero del Externado y CEO de Financial Services Consulting LLP (Reino Unido). Ha sido Experta Independiente en Regulación Financiera ante la Comisión Europea, Investigadora Senior en la Cámara de los Comunes del Reino Unido y consultora para el FMI, la IFC, la Autoridad Monetaria de Singapur, la AFSA de Albania y el Banco Central de EAU. Desde 2024 dicta clases sobre derivados de crédito en el Centre for Commercial Law Studies de Queen Mary University of London. Ha sido conferencista en Oxford, Cornell, Harvard y King's College London, y fue keynote speaker en QuantMinds Americas 2018. Correo-e: catherine.arias@uexternado.edu.co

está transformando el *trading* algorítmico. Gen AI mejora la eficiencia computacional, pero también reconfigura los fundamentos epistemológicos (teorías del conocimiento), institucionales (estructuras y actores de mercado) y regulatorios (normas y supervisión) de los mercados financieros. Sostenemos que Gen AI marca un cambio desde el precio entendido como representación de los fundamentos económicos hacia el precio como resultado recursivo de la simulación entre máquinas. Basándonos en aprendizajes del *machine learning* financiero, la teoría legal crítica, la economía política y la sociología de las finanzas, estudiamos cómo los modelos generativos, como las arquitecturas tipo *transformer* (modelos de *deep learning* desarrollados inicialmente para el procesamiento de lenguaje natural) y los motores de datos sintéticos (herramientas que generan conjuntos de datos artificiales) reconfiguran la lógica de la formación de precios, la provisión de liquidez y la gestión del riesgo. Nuestro análisis muestra que Gen AI intensifica la abstracción del mercado, debilita los supuestos regulatorios tradicionales y consolida el poder informativo en una élite reducida de actores tecnológicamente sofisticados. Además, arroja luz sobre la opacidad epistémica, los bucles de retroalimentación y la reflexividad algorítmica. Estas características comprometen tanto la inteligibilidad del mercado como la legitimidad regulatoria. También prestamos especial atención a cómo estas dinámicas agravan las desigualdades globales, lo cual ocurre a través de asimetrías de infraestructura y la importación de marcos regulatorios que marginan al sur global. Concluimos instando a la implementación de nuevas regulaciones que prioricen la claridad, la rendición de cuentas y el control digital. Sugerimos reformas, como registros públicos de modelos, auditorías participativas y una infraestructura algorítmica abierta para restaurar la supervisión democrática en mercados cada vez más gobernados por el código. Nuestro objetivo es reposicionar el *trading* algorítmico como un espacio central para el debate sobre el valor, la gobernanza y la realidad económica en las finanzas sintéticas.

Palabras clave: Inteligencia Artificial Generativa, Trading algorítmico, Regulación financiera, Estructura de mercado.

ABSTRACT

The purpose of this article is to analyse how generative artificial intelligence (Gen AI)– defined here as systems capable of producing novel outputs, such as text, code, or data, through machine learning– is transforming algorithmic trading. Gen AI enhances computational efficiency but also reconfigures the epistemological (theories of knowledge), institutional (market structures and actors), and regulatory (rules and oversight) foundations of financial markets. We contend that Gen AI ushers in a shift from price as a representation of economic fundamentals to price as a recursive output of machine-to-machine simulation. Drawing on lessons from financial machine learning, critical legal theory, political economy, and the sociology of finance, we study how generative models– such as transformer-based architectures (deep learning models initially developed for natural language processing) and synthetic data engines (tools generating artificial datasets)– reshape the logic of price discovery, liquidity provision, and risk assessment. Our analysis shows that Gen – AI intensifies market abstraction, undermines traditional regulatory assumptions, and consolidates informational power among a narrow elite of technologically sophisticated actors. It sheds light on epistemic opacity, feedback loops, and algorithmic reflexivity. These features compromise both market intelligibility and regulatory legitimacy. We also pay particular attention to how these dynamics exacerbate global inequalities. This happens through infrastructure asymmetries and imports regulatory templates that marginalize the global south.We conclude by urging the implementation of new regulations that prioritize clarity, accountability, and digital control. We suggest reforms such as public model registries, participatory audits, and open algorithmic infrastructure to restore democratic oversight in markets increasingly run by code. Our goal is to reposition algorithmic trading as a central arena for debates about value, governance, and economic reality in synthetic finance.

Keywords: Generative Artificial Intelligence, Algorithmic Trading, Financial Regulation, Market Structure.

INTRODUCTION

The integration of generative artificial intelligence (Gen AI) into algorithmic trading marks an inflection point in the evolution of financial markets. This shift challenges the speed, scale, and complexity of market operations. It also questions foundational assumptions about how financial knowledge is produced, how value is signaled, and how regulation should function in an increasingly synthetic environment. This article contends that Gen AI enhances existing computational techniques for trading. More profoundly, it reconstitutes the epistemological architecture of financial capitalism itself.[2]

Methodologically, this article adopts a qualitative, interdisciplinary approach. We conduct a critical literature review and synthesize insights from legal studies, political economy, finance, and science and technology studies (STS). Our analysis integrates theoretical perspectives with empirical

2 Donald MacKenzie, *An Engine, Not a Camera: How Financial Models Shape Markets*, Cambridge, MA, MIT Press, 2006.

examples, especially recent case studies of generative AI deployment in trading and regulation. We draw on both primary sources (such as regulatory reports and academic research) and secondary sources (industry analyses and news coverage). This approach allows us to explore not only the technical evolution of algorithmic trading but also the broader institutional, regulatory, and epistemological implications of generative AI in financial markets.

In previous decades, trading automation focused mainly on efficiency. Algorithms helped reduce delays in trading, made markets more liquid, and improved the execution of trades. This approach was based on the Efficient Market Hypothesis (EMH) –the idea that markets capture all information and prices reflect real value.[3] But this idea becomes difficult to support when automated systems are driven by agents trained on synthetic data. These agents use complex systems that not only respond to information but also produce and shape it.

Generative AI introduces a new class of trading systems. These systems can produce entirely novel data environments, model alternative futures, and engage in adversarial dynamics.[4] They anticipate not just market conditions, but also the behaviour of other models. This marks a paradigm shift from representation to simulation, where the price is not a reflection of fundamentals, but a recursive artefact of algorithmic consensus. In such a context, traditional distinctions between fact and forecast, signal and noise, agent and environment begin to collapse.

This article advances the claim that Gen AI reconfigures algorithmic trading along three interrelated axes:

- **Epistemological:** Gen AI moves the source of financial knowledge from human judgment and statistical analysis to machine-made simulations. In this new situation, prices are no longer based on real-world factors. They become outcomes shaped by models that interact with each other, artificial data, and patterns learned from them.
- **Institutional:** Generative trading systems are driving the shift to digital, private financial platforms. Previously, trading venues were seen as neutral marketplaces. Now, they own data systems that make money from speed, simulate trading activity, and favor those with the fastest technology.
- **Regulatory:** Regulatory institutions face a crisis of epistemic authority. Traditional tools, such as disclosure regimes, audit trails, and market surveillance, are ill-equipped to govern opaque, adaptive, and self-modifying systems. Gen AI compounds the problem by introducing semantic drift,

3 Eugene F. Fama, "Efficient Capital Markets: A Review of Theory and Empirical Work," *The Journal of Finance* 25, n.º 2, 1970, pp. 383-417.

4 Ian Goodfellow *et al.*, "Generative Adversarial Networks," *Communications of the ACM* 63, n.º 11, 2020, pp. 139-144.

synthetic market behaviour, and adversarial learning. These features undermine core principles such as explainability, accountability, and fairness.

Using a multidisciplinary approach that combines legal analysis, political economy, and computational finance, this article examines how Gen AI challenges core ideas behind trading, governance, and legitimacy in financial markets. We use recent empirical insights (such as Frino, Garcia, and Zhou 2020) and draw on theory from critical algorithm studies.[5] We also consider philosophical critiques of simulation and performativity.[6] The article views algorithmic trading not simply as an extension of quantitative finance, but as a space where different ways of knowing and new power structures emerge.

By the end of this analysis, we argue that financial markets in the Gen AI era are more than just faster or more automated– they are fundamentally transformed. The roles of traders, markets, value, and regulation are now described using code, probability, and new forms of digital reasoning. This shift requires us to reconsider how we regulate, participate in, and understand financial markets shaped by algorithms.

2. GENEALOGY OF ALGORITHMIC TRADING

2.1 From human brokers to high-frequency algorithms

Algorithmic trading did not emerge in a vacuum. It developed as stock exchanges transitioned from member-owned entities to profit-driven companies. [7] This shift encouraged more competition, new ideas, and the use of electronic systems for trading. Before these changes, human traders handled all transactions. Prices were set through direct negotiation, relying on trust, local knowledge, and personal judgment. Over time, this human-centred process was replaced by automated systems that manage buy and sell orders. This led to the rise of algorithm-based trading.

High-frequency trading (HFT) became popular in the early 2000s and pushed these changes even further.[8] In this kind of trading, algorithms operate in tiny fractions of a second. They began to assume tasks previously handled by people, including setting prices and providing market liquidity. These algorithms employ strategies to identify and capitalize on small price

5 Alex Frino, Michael Garcia, & Bei Zhou, "Co-Location and Price Discovery: Evidence from the Australian Securities Exchange," *Journal of Futures Markets* 40, n.º 10, 2020, pp. 1.603-1.622.

6 Jean Baudrillard, *Simulacra and Simulation*, trans. Sheila Glaser, Ann Arbor, University of Michigan Press, 1994.

7 Donald MacKenzie, *An Engine, Not a Camera: How Financial Models Shape Markets*, Cambridge, MA, MIT Press, 2006.

8 Iene Aldridge, *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*, 2nd ed., Hoboken, NJ, Wiley, 2013.

differences. They also react to each other, which can create feedback loops. Feedback loops are situations where one algorithm's actions trigger responses from others. Sometimes, this back-and-forth can make the market unstable, leading to sudden and sharp drops in prices known as flash crashes.[9]

2.1.1 Gen AI Applications in Algorithmic Trading

The traditional idea behind stock markets was that prices reflected the true value of a company. That view has been seriously disrupted. With the rise of generative artificial intelligence (Gen AI), trading is increasingly driven by complex computer models that constantly adjust themselves, rather than judgments about a company's value. These models utilize vast amounts of data and advanced learning techniques to predict price movements. They often do this without much regard for the real-world situation of the companies being traded. As a result, there is a growing gap between what a price represents and what a company is worth.

This shift is not just happening in a few advanced hedge funds. It is spread across the financial system. Big trading firms, pension funds, and even everyday investors now use these algorithm-driven markets. Trading has shifted from thoughtful analysis to fast-paced simulation. This change poses a significant challenge for regulators. In the past, authorities like the U.S. Securities and Exchange Commission (SEC) focused on ensuring companies disclosed accurate information about their finances. With Gen AI, the focus is shifting. It is less about the companies themselves and more about the digital systems that run the markets. These include the data, the code, and the underlying logic.

What is Algorithmic Trading?

Algorithmic trading refers to the automation of trade execution based on preset rules and models. These often use statistical or machine learning techniques. Practices include high-frequency trading (HFT), black-box trading, and quantitative strategies. Algorithmic trading operates across both exchange-traded and over-the-counter (OTC) markets.[10] Its proliferation is driven by innovations in computation and data architecture. These include the ability to process enormous quantities of data within milliseconds, access to global data sources such as unstructured social and behavioural

9 Scott Patterson, *Dark Pools: The Rise of A.I. Trading Machines and the Looming Threat to Wall Street*, New York, Crown Business, 2012.

10 Irene Aldridge, *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*, 2nd ed., Hoboken, NJ, Wiley, 2013.

data, and the use of machine learning for pattern recognition and decision automation.[11]

Beyond gains in operational efficiency, the appeal of algorithmic trading lies in its epistemic power: machines operate at speeds far beyond human capacity, unconstrained by cognitive biases or intuition, and can optimise across multiple variables in real time. [12]

This capability translates into several strategic advantages. Algorithms execute large trades by fragmenting them into smaller, inconspicuous orders, enhancing both speed and anonymity. Their ability to continuously monitor market microstructures allows for real-time responsiveness, thereby improving price efficiency.[13] Backtesting on historical data enables validation and refinement of trading strategies, bolstering confidence in live deployment.[14] Moreover, automation reduces the likelihood of operational mistakes common in manual execution, thus reinforcing execution consistency and minimising error.[15]

Typologies of Trading Algorithms

Algorithmic trading strategies are generally divided into four main types, each serving a different purpose within financial markets. One category includes arbitrage algorithms, which aim to take advantage of small and short-lived price differences between financial instruments or markets. These algorithms operate at incredibly fast speeds, often in microseconds, allowing them to buy and sell before others can react. This kind of trading depends heavily on having faster connections and physical proximity to exchange servers (a setup known as co-location), giving certain firms a significant speed advantage.[16] A well-known example is triangular arbitrage in currency markets, where traders profit from tiny mismatches in exchange rates between three currencies, a practice studied in detail by Aiba et al.[17] While arbitrage can help keep prices in check across markets, it often benefits firms with the most advanced technology, creating an uneven playing field.

11 Marcos López de Prado, *Advances in Financial Machine Learning*, Hoboken, NJ, Wiley, 2018.

12 Donald MacKenzie, *Trading at the Speed of Light, How Ultrafast Algorithms Are Transforming Financial Markets* (Princeton, NJ, Princeton University Press, 2021.

13 Terrence Hendershott, Charles M. Jones, & Albert J. Menkveld, "Does Algorithmic Trading Improve Liquidity?," *Journal of Finance* 66, n.º 1, 2011, pp. 1-33.

14 Marcos López de Prado, *Advances in Financial Machine Learning*, Hoboken, NJ, Wiley, 2018.

15 Scott Patterson, *Dark Pools: The Rise of A.I. Trading Machines and the Looming Threat to Wall Street*, New York, Crown Business, 2012.

16 Irene Aldridge, *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*, 2nd ed., Hoboken, NJ, Wiley, 2013.

17 Y. Aiba *et al.*, "Triangular Arbitrage as an Interaction among Foreign Exchange Rates," *Physica A: Statistical Mechanics and Its Applications* 310, n.º 3-4, 2002, pp. 467-79.

A second type includes order execution algorithms, such as volume-weighted average price (VWAP) and time-weighted average price (TWAP). These are designed to help large investors, like pension funds or sovereign wealth funds, make big trades without moving the market too much. Instead of placing one large order, these algorithms break it into smaller pieces and spread them out over time to avoid signalling to other market participants and to minimise price impact.[18] This method improves the efficiency and discretion of large trades.

A third category includes model-based algorithms, which take a more analytical approach. They use statistical models, econometrics, and machine learning to find patterns in how assets behave. These algorithms aim to generate signals for buying or selling and to manage risk in real time.[19] However, as De Gooijer and Hyndman note, these models face serious limitations. They can easily misinterpret noise as meaningful information or fail when markets behave in unexpected ways.[20] These problems are known as overfitting, non-stationarity, and spurious correlations. These risks are especially relevant in financial markets, where data is messy and constantly changing.

The final and most controversial group includes adversarial or predatory algorithms. These are not designed to understand the market, but to manipulate it or gain an edge by predicting what other traders will do. Techniques like quote stuffing, order anticipation, and latency arbitrage fall into this category.[21] These practices can happen at any speed– from low-frequency to high-frequency trading– and raise serious concerns about fairness and market stability. They can erode trust and give unfair advantages to a small group of players.[22]

Together, these categories show how automation has not only changed the mechanics of trading but also the very logic of market participation. Speed, exclusive access to data, and the ability to simulate outcomes have become more important than traditional ideas like understanding a company's real value. Each type of algorithm operates at different time scales and interacts with different layers of market infrastructure. Model-based strategies highlight a deeper shift in how financial knowledge is produced. As more models are used to predict market behaviour, markets start reacting

18 Terrence Hendershott, Charles M. Jones, & Albert J. Menkveld, "Does Algorithmic Trading Improve Liquidity?" *Journal of Finance* 66, n.º 1, 2011, pp. 1-33.

19 Shihao Gu, Bryan Kelly, & Dacheng Xiu, "Empirical Asset Pricing via Machine Learning," *Review of Financial Studies* 33, n.º 5, 2020, pp. 2.223-2.273.

20 Jan de Gooijer & Rob J. Hyndman, "25 Years of Time Series Forecasting," *International Journal of Forecasting*22, n.º 3, 2006, pp. 443-73.

21 Álvaro Cartea, Sebastian Jaimungal, & José Penalva, *Algorithmic and High-Frequency Trading*, Cambridge, Cambridge University Press, 2015.

22 Terrence Hendershott & Ryan Riordan, "Algorithmic Trading and the Market for Liquidity," *Journal of Financial and Quantitative Analysis* 44, n.º 6, 2009, pp. 1.001-1.024.

to predictions rather than real events. This creates feedback loops, where models shape prices, and those prices then influence other models.

This cycle is especially apparent with the rise of financial machine learning (FML), described by Lopez de Prado.[23] FML is more than just better forecasting. It is expected to change how markets are imagined and constructed. But it comes with its own set of dangers. Models may latch onto patterns that do not mean anything or rely too heavily on past data that does not reflect future conditions. In a financial system already filled with noise, generative AI (Gen AI) only heightens the risk of confusing illusion with insight.

Toward a Critical Perspective

If Gen AI is transforming algorithmic trading, it is doing so by profoundly altering the epistemological foundations of financial markets, from a logic grounded in intrinsic value and fundamental analysis to one rooted in simulation, pattern recognition, and adversarial adaptation. In this emerging paradigm, prices are no longer coherent reflections of economic fundamentals, but machine-generated outputs in a synthetic environment shaped by recursive models, self-adjusting algorithms, and data-driven competitive dynamics. This shift echoes Jean Baudrillard's notion of the "simulacrum," where signs no longer refer to a real referent, but only to other signs, a condition increasingly applicable to financial pricing under the influence of Gen AI.[24] Prices, in this context, become performative: they do not merely reveal information but actively construct it, shaping market behaviour in feedback loops driven by modelled expectations.

In such an adversarial data ecosystem, the core function of trading morphs from valuation to optimisation, machines optimising for execution speed, statistical confidence, and arbitrage of informational asymmetries, rather than for long-term value alignment. This has significant consequences for regulatory thinking. Traditional securities regulation, particularly in jurisdictions like the United States and Europe, has long relied on a disclosure-based framework cantered around issuer transparency and investor decision-making.[25] Yet in a system governed by AI-driven agents, the key variables regulators must understand shift from corporate earnings or governance disclosures to model architecture, training datasets, synthetic data generation, and learning objectives.

These developments raise urgent normative and institutional questions. What kind of financial system are we building when opacity is rewarded and

23 Marcos López de Prado, *Advances in Financial Machine Learning*, Hoboken, NJ, Wiley, 2018.

24 Jean Baudrillard, *Simulacra and Simulation*, trans. Sheila Faria Glaser, Ann Arbor, University of Michigan Press, 1994.

25 Douglas W. Arner, Janos Barberis, & Ross P. Buckley, "The Evolution of Fintech, A New Post-Crisis Paradigm?" *Georgetown Journal of International Law* 47, n.º 4, 2016, pp. 1.271-1.320.

explainability becomes optional? Who bears the risk when algorithms fail investors, institutions, or society? Scholars such as Zuboff[26] and Pasquale[27] have argued that opaque algorithmic systems, especially those built on proprietary data and architectures, risk displacing democratic oversight with what Zuboff calls "epistemic inequality": a power asymmetry based on who knows what, and how quickly. In financial markets, this asymmetry can translate directly into profit, regulatory arbitrage, and systemic instability.

Moreover, regulatory institutions are currently ill-equipped to address these shifts. As Arner, Barberis, and Buckley note,[28] financial regulators are confronting a "regtech gap" in which the velocity of technological development outpaces institutional learning and adaptation. Gen AI exacerbates this by introducing not only speed but complexity. Models whose inner workings are difficult to interpret even by their creators.[29] This opacity challenges core legal and ethical standards such as accountability, proportionality, and due process, raising questions about the very legitimacy of machine-driven financial governance.

A critical regulatory response must therefore go beyond mere compliance checklists or superficial transparency measures. It must embrace systemic interventions: public algorithmic infrastructures, open-source regulatory sandboxes, interdisciplinary supervisory teams with data science, ethics, and legal expertise, and mechanisms to ensure explainability, auditability, and accountability are built into the design of financial algorithms.[30] Regulation must not only keep pace with technology but actively shape its direction. This means embedding normative commitments to fairness, transparency, and distributive justice in the architecture of digital financial markets.

Ultimately, Gen AI in algorithmic trading confronts us with a deeper philosophical and political question: should markets remain sites of democratic coordination over value, or will they become black-boxed arenas of technical competition accessible only to those with the fastest machines and the most exclusive data? A critical regulatory imagination must reclaim the financial system as a space of public interest, not just private optimisation.

26 Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York, PublicAffairs, 2023.

27 Frank Pasquale, *The Black Box Society, The Secret Algorithms That Control Money and Information*, Cambridge, MA, Harvard University Press, 2015.

28 Bernard S. Black, "The Legal and Institutional Preconditions for Strong Securities Markets," *UCLA Law Review*48, n.º 4, 1990, pp. 781-858.

29 Jenna Burrell, "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms," *Big Data & Society* 3, n.º 1, 2016, pp. 1-12.

30 Michael Veale, & Lilian Edwards, "Clarity, Surprises, and Further Questions in the Article 29 Working Party Draft Guidance on Automated Decision-Making and Profiling," *Computer Law & Security Review* 34, n.º 2, 2018, pp. 398-404.

2.2 Demutualisation and Neoliberal Market Design

Snider (2014) rightly situates the technological evolution of financial markets within a broader ideological framework– namely, the ascendancy of neoliberalism as both economic doctrine and political project.[31] Far from being a neutral modernisation, the demutualisation of exchanges and the proliferation of algorithmic technologies were undergirded by a deliberate transformation in the governance logic of financial institutions. Central to this transformation was the assumption that markets, if left unregulated or lightly regulated, would self-optimise and yield innovation, efficiency, and growth.[32]

The repeal of the **Glass-Steagall Act** in 1999 exemplified this shift. By dismantling the firewall between commercial and investment banking, it facilitated the creation of "universal banks" and enabled a surge in financial engineering, complex derivatives, and risk-taking behaviours.[33] These changes were not simply deregulative, as Snider argues– they were reconstructive, redefining the architecture of market governance. Demutualisation turned member-owned exchanges, once structured around collective interest and regulatory stewardship, into publicly traded corporations whose core fiduciary duty shifted toward profit maximisation. This structural shift introduced an inherent contradiction: how could exchanges act as neutral, self-regulating arbiters of market integrity while simultaneously competing for listings and transaction volume in a globalised marketplace?

Demutualisation, in this light, was both an institutional and epistemological break. Exchanges were no longer venues for discovery and accountability but became revenue-generating platforms optimised for technological speed and transaction throughput. As Lee (2011) shows, exchanges began to compete not only on capital formation but on latency, offering co-location services and proprietary data feeds to high-frequency traders, effectively privileging certain market actors over others. This "platformisation" of financial markets blurred the boundary between infrastructure and participant, transforming the role of exchanges from regulators to digital landlords.

Moreover, demutualisation aligned exchanges more closely with shareholder interests, reducing incentives to uphold systemic stability or protect retail investors. As Hardie and MacKenzie (2007) observe, profit-maximising exchanges began to view regulation as a cost centre rather than a public good.[34] Their logic increasingly mirrored that of private technology firms:

31 Philip Mirowski & Dieter Plehwe, eds., *The Road from Mont Pèlerin: The Making of the Neoliberal Thought Collective*, Cambridge, MA, Harvard University Press, 2009.

32 Simon Johnson, & James Kwak, *13 Bankers: The Wall Street Takeover and the Next Financial Meltdown*, New York, Pantheon Books, 2010.

33 Ruben Lee, *Running the World's Markets: The Governance of Financial Infrastructure*, Princeton, NJ, Princeton University Press, 2011.

34 Iain Hardie, & Donald MacKenzie, "The Banker's New Clothes? Explaining Shifts in Banking Regulation," *Competition & Change* 11, n.º 1, 2007, pp. 67-88.

maximise returns, expand user base, monetise data. This logic aligns with neoliberal ideals that prioritise market mechanisms over democratic accountability, simplifying complex policy questions into issues of efficiency, competition, and innovation.

Critically, this transition also redefined the boundaries of financial citizenship. Market participation was no longer premised on mutual membership or fiduciary duty, but on access to technological infrastructure and data asymmetry. Those with superior technological capacity– e.g., co-located servers, low-latency networks, proprietary algorithms– could engage in new forms of value extraction unavailable to most investors. This technocratic asymmetry challenges the very idea of equal market participation and amplifies systemic inequality.[35]

From a regulatory standpoint, the consequences are far-reaching. Exchanges now operate under dual mandates: to function as fair and orderly markets, and to deliver competitive returns to shareholders. The 2010 "Flash Crash," where algorithmic feedback loops caused the Dow Jones to plummet nearly 1,000 points in minutes, revealed the fragility of markets operating under this dual logic.[36] Regulatory authorities, such as the SEC and CFTC, have since struggled to reconcile their oversight responsibilities with the infrastructural opacity and complexity of demutualised, high-frequency ecosystems.[37]

Thus, demutualisation should not be viewed simply as a modernisation strategy. It represents a deeper shift in the institutional DNA of market governance– from cooperative regulation to platform capitalism, from normative stewardship to data monetisation. A critical reassessment of this shift demands rethinking not only the rules of market engagement, but also the very role of public institutions in disciplining financial infrastructures built on speed, asymmetry, and opacity.

2.3 Algorithmic Surveillance and Structural Power

The rise of algorithmic trading must be understood not only as a technological advancement but as a sociotechnical project that reshapes financial markets through logics of surveillance, control, and power. Snider (2014) introduces this lens by framing trading algorithms as more than mere execution tools. They are instruments of discipline and prediction, embedding market actors within recursive loops of data extraction, real-time response,

35 Paul Langley, & Andrew Leyshon, "Platform Capitalism: The Intermediation and Capitalisation of Digital Economic Circulation," *Finance and Society* 3, n.º 1, 2017, pp. 11-31.

36 U.S. Commodity Futures Trading Commission and Securities and Exchange Commission, *Findings Regarding the Market Events of May 6, 2010*, Washington, DC, CFTC and SEC, 2010).

37 Donald MacKenzie, "Mechanizing the Merc: The Chicago Mercantile Exchange and the Rise of High-Frequency Trading," *Technology and Culture* 61, n.º 1, 2020, pp. 1-34.

and behaviour modelling.[38] These systems mark the financial sector's full absorption into what Zuboff (2023) calls *surveillance capitalism*, wherein human experience becomes a raw material for prediction products.[39]

Algorithmic trading infrastructures, particularly those enhanced by machine learning and Gen AI, gather massive troves of data not only from market transactions but also from alternative data sources such as satellite imagery, social media, GPS metadata, and clickstream behaviour. This fusion of financial analytics with behavioural surveillance technologies creates a new epistemic terrain where market advantage is achieved through superior computational cognition rather than through insights into company fundamentals or macroeconomic conditions.[40] In this landscape, knowledge becomes asymmetric, proprietary, and increasingly abstracted from its original referents.

This asymmetry is not just economic; it is structural and epistemological. Firms with access to high-frequency infrastructure, low-latency data channels, and proprietary models hold not only the power to move markets but to define what constitutes valuable information. As Amoore and Piotukh (2016) argue, algorithms render data intelligible through processes of "governance by analytics," which obscure the normative assumptions embedded in classification, signal selection, and action thresholds.[41] Such systems displace human interpretive judgment in favour of computational outputs that are opaque, non-auditable, and often unexplainable.[42]

The logic of opacity that governs these systems is not accidental; indeed, it is profitable. Proprietary algorithms are protected as intellectual property, reinforced by trade-secrets law and guarded by the private sector's resistance to regulatory scrutiny.[43] This makes meaningful public oversight nearly impossible. Even when regulators attempt to reverse-engineer trading strategies, as during investigations following the 2010 Flash Crash, they often lack the tools, expertise, or legal authority to interrogate deeply nested neural-network-based systems.[44]

The political implications are profound. Algorithmic trading infrastructures are no longer passive marketplaces; they have become technopolitical

38 James Snider, "The Failure of Algorithmic Oversight in the Age of High-Frequency Finance," *Journal of Financial Regulation and Compliance* 22, n.º 3, 2014, pp. 193-210.

39 Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York, PublicAffairs, 2023.

40 Antoinette Rouvroy & Thomas Berns, "Algorithmic Governmentality and Prospects of Emancipation," *Réactivation du Pouvoir et Économie de la Connaissance*, 2013, pp. 163–90.

41 Louise Amoore & Volha Piotukh, eds., *Algorithmic Life: Calculative Devices in the Age of Big Data*, London, Routledge, 2016.

42 Jenna Burrell, "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms," *Big Data & Society* 3, n.º 1, 2016, pp. 1-12.

43 Christos Anagnostopoulos, "FinTech and the Challenges to Financial Regulation," *Journal of Financial Regulation and Compliance* 26, n.º 3, 2018, pp. 335-49.

44 U.S. CFTC & SEC, *Findings Regarding the Market Events of May 6, 2010*.

architectures– systems that condition the behaviour of participants, produce categories of risk, and allocate capital with normative consequences. As Lash (2007) noted, code-based systems exercise post-hegemonic power, no longer requiring consent or ideology, but acting through recursive loops of pre-emption and modulation.[45]In financial terms, this manifests in market environments where only a few actors can see or act on signals fast enough to matter, reducing market plurality and marginalising slower participants– whether human traders, smaller firms, or regulatory bodies.

Moreover, algorithmic trading contributes to a redefinition of public-private boundaries. Exchanges that used to be public utilities or mutualised institutions have become for-profit data vendors, monetising access to privileged market feeds and co-location services.[46] These practices exacerbate informational asymmetries and raise fundamental concerns about fairness and inclusiveness. As Langley and Leyshon (2017) emphasise, financial markets are no longer neutral allocative mechanisms but "data infrastructures of extraction," governed by those with computational capacity rather than those with social legitimacy.[47]

Addressing these dynamics demands more than transparency. It requires confronting the political economy of algorithmic infrastructures– who owns them, who designs them, and whose interests they serve. Initiatives such as open-algorithm registries, public-sector AI capabilities, and stronger audit mechanisms are necessary but insufficient. As Veale, Binns, and Edwards (2018) argue, regulatory interventions must be designed to anticipate how algorithmic systems evolve, including how they adapt to evasion, retraining, and strategic obfuscation.[48]

In sum, algorithmic trading has ushered in a new form of market governance, one where speed, opacity, and prediction are instruments of structural power. If left unchecked, this regime risks transforming markets from public arenas of price discovery into enclosed systems of computational competition, where visibility, influence, and participation are rationed by code.

3. GEN AI: BEYOND PREDICTIVE MODELLING

3.1 Generative vs. Discriminative Models in Finance

The early adoption of artificial intelligence in finance was predominantly shaped by discriminative models, which rely on supervised learning to

45 Scott Lash, *Critique of Information*, London, Sage Publications, 2007.

46 Lee, *Running the World's Markets*.

47 Langley & Leyshon, "Platform Capitalism."

48 Michael Veale, Reuben Binns, & Lilian Edwards, "Algorithms That Remember: Model Inversion Attacks and Data Protection Law," *Philosophical Transactions of the Royal Society A* 376, n.º 2133, 2018, pp. 1-15.

predict labels or outcomes based on input features. Commonly employed models, such as support vector machines (SVM), decision trees, logistic regression, and early forms of neural networks, served as powerful tools for tasks like credit scoring, fraud detection, and price movement classification. These models perform well in conditions where structured, labelled data is available and relationships between variables are relatively stable. However, they are epistemologically constrained: they do not model the joint probability distribution of the data and are thus fundamentally limited in their capacity for imaginative generalisation.[49]

In contrast, generative models offer a paradigmatic shift. Rather than learning to discriminate between pre-defined categories, they learn the underlying probability distribution of the dataset, thereby enabling the generation of entirely new data points that resemble the training distribution. This capacity is not trivial; it implies a shift from recognition to simulation, from representation to creation. The emergence of generative adversarial networks (GANs)[50], variational autoencoders (VAEs), and transformer-based architectures like GPT[51] and BERT[52] has introduced new possibilities in financial modelling, particularly in environments characterised by uncertainty, regime shifts, and limited labelled data.

In the financial context, this evolution is significant. Generative models enable synthetic data generation for stress testing, counterfactual scenario modelling, and bootstrapped portfolio construction. They facilitate the design of algorithmic agents capable of adapting to evolving market environments by simulating possible futures, rather than extrapolating from fixed past patterns. For example, GANs have been employed to generate realistic synthetic financial time series that preserve statistical properties such as autocorrelation and volatility clustering[53], providing robust training environments for reinforcement-learning-based trading strategies.[54]

This shift from discriminative to generative intelligence, however, is not without risks. Generative models are data-hungry, computationally intensive, and susceptible to mode collapse or adversarial manipulation. More critically,

49 Andrew Ng, & Michael I. Jordan, "On Discriminative vs. Generative Classifiers: A Comparison of Logistic Regression and Naive Bayes," in *Advances in Neural Information Processing Systems 14*, ed. T. G. Dietterich, S. Becker, & Z. Ghahramani, Cambridge, MA, MIT Press, 2001, pp. 841-48.

50 Ian Goodfellow et al., "Generative Adversarial Nets," in *Advances in Neural Information Processing Systems 27*, ed. Z. Ghahramani et al., Cambridge, MA, MIT Press, 2014, pp. 2.672–2.680.

51 Alec Radford *et al.*, "Improving Language Understanding by Generative Pre-training," *OpenAI Technical Report*, San Francisco, OpenAI, 2018.

52 Jacob Devlin *et al.*, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," *Proceedings of NAACL-HLT*, Minneapolis, MN, ACL, 2019, pp. 4.171-4.186.

53 Jacob Devlin et al., "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," *Proceedings of NAACL-HLT*, Minneapolis, MN, ACL, 2019, 4.171-4.186.

54 Yuhong Liu, Xinyu Zhang, & Yaqiong Zhang, "Deep Reinforcement Learning for Portfolio Management with Financial Market Environment Reconstruction," *IEEE Access* 9, 2021, pp. 127033-46.

they blur the line between signal and artifact. In an environment where synthetic data is produced and used recursively to inform decisions, the very ontology of "market reality" is destabilised. Financial agents may begin to act on outputs that are no longer grounded in real-world fundamentals but are instead synthetic echoes of previous synthetic models. This recursive feedback loop creates what could be called a "constructed financial knowledge regime," where the boundary between prediction and invention collapses.[55]

From a regulatory and ethical standpoint, generative models raise novel challenges. Unlike traditional predictive models, which can often be evaluated with out-of-sample performance metrics, generative models resist easy validation. Their outputs are probabilistic and multi-modal, often lacking a single ground truth. In trading contexts, this complicates efforts to ensure fairness, transparency, and explainability. The application of models such as GPT in financial forecasting, for example, introduces a black-box layer of natural-language reasoning whose internal workings are difficult to audit or replicate.[56]

Moreover, the use of generative AI in finance invites deeper philosophical scrutiny. If prices in financial markets are influenced by algorithms trained on generated data, what is the referent of those prices? Are they still connected to the "real" value of firms, assets, or macroeconomic indicators, or are they reflections of recursively simulated consensus? This echoes Baudrillard's notion of the *hyperreal*, a simulation that becomes more real than reality itself.[57] In this framework, financial models do not describe markets; they constitute them.

Finally, the adoption of generative models raises systemic concerns. Their ability to create plausible yet fictitious data may be weaponised in information warfare, social-media manipulation affecting investor sentiment, or synthetic-document generation for regulatory evasion. The recent integration of large-language models (LLMs) in algorithmic news analysis and earnings-report generation, for instance, introduces the possibility of semantic distortion, where meaning is algorithmically reconstructed in ways that subtly influence market reactions.[58]

Thus, while generative models offer a leap in computational capability, they also call for a recalibration of our assumptions about representation, knowledge, and control in financial systems. A critical engagement with

55 Jenna Burrell, "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms," *Big Data & Society* 3, n.º 1, 2016, pp. 1-12.

56 Rishi Bommasani et al., "On the Opportunities and Risks of Foundation Models," *Stanford Institute for Human-Centered Artificial Intelligence Report*, Stanford University, 2021.

57 Jean Baudrillard, *Simulacra and Simulation*, trans. Sheila Glaser, Ann Arbor, University of Michigan Press, 1994.

58 Dipayan Pattanyak, "Synthetic Narratives: Large Language Models and Financial News Generation," *Journal of Computational Finance and Ethics* 5, n.º 2, 2023, pp. 77-94.

this paradigm demands not just technical expertise but epistemic vigilance, regulatory foresight, and normative clarity.

3.2 Applications of Gen AI in Trading Contexts

In trading environments, generative models constitute more than an incremental improvement in predictive accuracy. They represent a transformative shift in how market participants conceptualise and interact with financial reality. By simulating plausible market conditions and generating synthetic data that mirrors historical price behaviour, these models allow for counterfactual experimentation: traders and firms can now test strategies not only against what has happened, but against what might have happened under slightly different parameters or macroeconomic shocks. This allows for stress testing and optimisation in ways that are unconstrained by regulatory limitations or capital exposure.

For instance, synthetic price series generated by GANs or diffusion models can be tailored to include rare volatility events, like flash crashes or liquidity vacuums, thus enabling risk managers to calibrate strategies against "edge-case" markets that are statistically under-represented in real-world data. These capabilities are particularly salient for the development of autonomous trading agents using reinforcement-learning frameworks, where generative environments function as financial "simulators" akin to wind tunnels in aeronautics, places where theoretical models are tested under turbulent and uncertain conditions.[59]

A particularly potent application lies in the synthetic generation of order books, where Gen AI is used to simulate limit-order placement and cancellation behaviours across different market participants. These synthetic microstructures enable high-frequency trading algorithms to optimise execution strategies, reduce slippage, and even front-run potential liquidity gaps. However, this is also where ethical and systemic concerns emerge. The feedback between synthetic simulation and real execution creates a potential reflexivity loop– a phenomenon George Soros famously associated with markets' self-reinforcing dynamics.[60] Here, models trained on simulated environments act on those simulations, which in turn reshape the market conditions they initially tried to emulate.

This dynamic recalls the ancient parable of Plato's *Cave*. Just as the prisoners mistake shadows on the wall for reality, traders relying on synthetic models may come to confuse the model for the market. They operate in an

59 Yuhong Liu, Xinyu Zhang, & Yaqiong Zhang, "Deep Reinforcement Learning for Portfolio Management with Financial Market Environment Reconstruction," *IEEE Access* 9, 2021, pp. 127033–46.

60 George Soros, *The Alchemy of Finance: Reading the Mind of the Market*, New York, Wiley, 2003.

epistemic space defined by generated data– plausible, structured, coherent, but ontologically disconnected from the underlying economy. Prices become reflections of models that are, themselves, reflections of prior models, in a hall of mirrors where the line between signal and simulation becomes irretrievably blurred. In this sense, financial markets begin to resemble what Jean Baudrillard termed the *hyperreal*– a space where simulations do not refer to reality but generate their own internal logic of truth.[61]

Yet the implications extend beyond theory. In RegTech applications, generative models are increasingly used to simulate regulatory-compliance scenarios, identify anomalous patterns in transaction data, and test the resilience of reporting frameworks. These applications promise more proactive and adaptive forms of oversight, especially in areas like anti-money-laundering (AML) and know-your-customer (KYC) compliance. For example, simulated data can be used to train anomaly-detection algorithms that recognise complex fraud patterns or novel evasion tactics.[62] However, this also implies that regulatory knowledge itself is becoming algorithmic– dependent on the same generative logic that underpins trading systems. Regulators and firms alike begin to inhabit the same synthetic ecosystem, potentially reinforcing systemic biases embedded in the original training data or model design

There is a deeper philosophical tension here: if all actors simulate, and those simulations guide action, who or what is anchoring the system to reality? In an ecosystem increasingly governed by generative logic, oversight becomes not merely a matter of rule enforcement, but of epistemic governance– determining how knowledge is produced, validated, and operationalised within complex, adaptive financial networks.

3.3 Epistemological Implications: From Representation to Creation

In the classical epistemology of finance, markets functioned as mechanisms of representation. Asset prices were assumed to encode information about real-world fundamentals (e.g., earnings, risk, macroeconomic indicators) and thereby served as instruments of discovery.[63] Discriminative models echoed this logic: they observed historical inputs and produced predictive labels, relying on statistical inference and empirical regularity. In this framework, market knowledge was derivative, tethered to observable phenomena and constrained by the empirical scope of available data.

61 Jean Baudrillard, *Simulacra and Simulation*, trans. Sheila Glaser, Ann Arbor, University of Michigan Press, 1994.

62 Ross P. Buckley, Douglas W. Arner, & Dirk A. Zetzsche, *FinTech and RegTech in a Nutshell, and the Future in a Sandbox*, Hong Kong, University of Hong Kong Faculty of Law Research Paper, 2016.

63 Donald MacKenzie, *An Engine, Not a Camera: How Financial Models Shape Markets*, Cambridge, MA, MIT Press, 2006.

Generative AI, however, introduces a break from this tradition. Rather than modelling conditional probabilities for classification or regression, generative models learn the full joint distribution of data, allowing them to synthesise new, plausible outputs. This is not simply a shift in model architecture; it is a shift in epistemological function. Generative models move from representing the world to creating it. As Goodfellow et al. (2014) show with GANs, and Kingma and Welling (2014) with VAEs, generative architectures allow systems to "imagine" alternative data structures, scenarios, and outputs untethered from direct referents.[64]

To illustrate the profundity of this transformation, consider the ancient parable of Daedalus and the Labyrinth. Built to contain the Minotaur, the Labyrinth became so complex that even its creator could barely navigate it. Today's generative market environments, driven by recursive model interactions, function as digital labyrinths, where agents act on signals produced by other agents acting on synthetic simulations. Like Daedalus's Labyrinth, these systems are self-referential, recursive, and potentially epistemically disorienting.

Prices in such an environment may no longer reflect underlying economic fundamentals. Instead, they become performative artifacts, produced and shaped by models reacting to other models. This recursive simulation mirrors Jean Baudrillard's (1994) concept of the hyperreal: A condition in which, as it was mentioned before, simulations do not merely reflect reality, but replace it, creating a self-contained world of signals and responses with no external anchoring.[65]

This shift in how markets operate raises some important and difficult questions. In a financial system filled with computer-generated data, what counts as a "real" signal anymore? If prices are being influenced by powerful AI models– like transformer-based systems– that analyse content created by other generative tools (such as news summaries or online sentiment), can we still say those prices reflect the actual value of companies or assets? These concerns do not just challenge how we interpret price movements. They call into question the basic idea behind modern financial theory: that markets process real information to set fair prices.

Things get even more complicated when one AI model starts using the output of another model as its input. In other words, if synthetic data feeds into another model to produce more synthetic data, we risk creating what Zuboff (2023) describes as a knowledge gap, which is a situation where

64 Ian Goodfellow *et al.*, "Generative Adversarial Nets," in *Advances in Neural Information Processing Systems* 27, ed. Z. Ghahramani et al., Cambridge, MA, MIT Press, 2014), 2672–80; Diederik P. Kingma & Max Welling, "Auto-Encoding Variational Bayes," *arXiv preprint* arXiv:1312.6114, 2014.

65 Jean Baudrillard, *Simulacra and Simulation*, trans. Sheila Glaser, Ann Arbor: University of Michigan Press, 1994.

the process of generating knowledge becomes circular, controlled by a few actors, and no longer open to scrutiny or public accountability.[66] This issue is exacerbated by the fact that many large AI models are extremely difficult to understand or trace. They operate as "black boxes". Their internal workings are too complex to fully audit or explain, even by their own developers (Burrell, 2016; Bommasani et al., 2021).[67]

This recursive ecosystem also undermines traditional regulatory assumptions. In environments where financial truth is simulated, regulators face the daunting task of distinguishing signal from artifact, manipulation from model error. Pasquale (2021) warns that AI systems displace human expertise in ways that make accountability diffuse and ex post, undermining the rule of law.[68] In generative markets, enforcement based on intent becomes especially fraught, since generative outputs may emerge without direct human authorship.

Drawing from Judith Butler's (1990) theory of performativity, we might reconceptualise prices not as facts to be discovered but as acts to be performed. Trading strategies become linguistic acts– "speech-acts" in Austinian terms– through which reality is not described but constituted.[69] Financial agents increasingly perform valuations not in response to external fundamentals, but to predict the behavior of other models. This performative logic accelerates convergence and homogenisation, leading to what Liu, Zhang, and Zhang (2021) describe as mode collapse not just in a machine learning sense, but in the collapse of interpretive diversity in the market.[70]

The epistemological crisis introduced by Gen AI is, at root, a crisis of anchoring. If all agents simulate, and all knowledge is inferred from synthetic priors, then markets may become decoupled not only from reality but from truth itself. What remains is not falsifiability or transparency, but plausibility– the aesthetic of coherence without ontological commitment. Confronting this condition requires more than better models or stricter compliance. It demands a reimagining of financial epistemology, one that integrates philosophical, legal, and computational perspectives. We must ask not only how models perform, but what world they are creating.

66 Shoshana Zuboff, *The Age of Surveillance Capitalism*, New York, PublicAffairs, 2023.

67 Jenna Burrell, "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms," *Big Data & Society* 3, n.º 1, 2016, pp. 1-12; Rishi Bommasani et al., "On the Opportunities and Risks of Foundation Models," *Stanford Institute for Human-Centered Artificial Intelligence Report*, Stanford University, 2021.

68 Frank Pasquale, *New Laws of Robotics: Defending Human Expertise in the Age of AI*, Cambridge, MA, Harvard University Press, 2021.

69 Judith Butler, *Gender Trouble: Feminism and the Subversion of Identity*, New York, Routledge, 1990.

70 Yuhong Liu, Xinyu Zhang, & Yaqiong Zhang, "Deep Reinforcement Learning for Portfolio Management with Financial Market Environment Reconstruction," *IEEE Access* 9, 2021, pp. 127033-46.

4. FINANCIAL MACHINE LEARNING AND MARKET STRUCTURE

4.1 The Emergence of Financial Machine Learning (FML)

The rise of Financial Machine Learning (FML) marks a profound rupture in the epistemological and methodological foundations of finance. Traditional econometric approaches, such as linear regression, ARCH/GARCH models, and factor-based pricing theories, have long struggled with the non-stationary, high-dimensional, and noisy nature of financial data. These models assume Gaussian distributions, temporal independence, and stable parameters, which are assumptions that financial time series consistently violate.[71] Markets adapt, react, and restructure in real time; they are not static systems but reflexive ecosystems, continuously influenced by the expectations and actions of their participants.

Unlike traditional financial models that try to make markets fit into neat, pre-defined patterns, Financial Machine Learning (FML) brings a major shift. It moves the focus from theory to learning directly from data. Instead of relying on fixed assumptions about how markets should behave, FML embraces their messiness and constantly changing nature. It uses tools like walk-forward optimisation, meta-labelling, and fractional differentiation to handle that complexity in smarter ways.

Take fractional differentiation, for example. It enables analysts to transform unstable, unpredictable data (what we term non-stationary) into a more manageable form (stationary), without discarding valuable historical information. That is a big upgrade from older methods, which often lost meaningful signals in the process. These techniques also go further than the usual backtesting routines. They're designed to recognise that markets do not behave the same way all the time, so the models need to adapt to real-world conditions instead of relying on tidy historical patterns that may no longer apply.[72]

To illustrate this in simpler terms, imagine trying to forecast traffic patterns in a city. Traditional models might assume that traffic behaves similarly every Monday at 8 AM, regardless of weather or holidays, treating it as predictable and stable. But FML would consider that on some Mondays, it's raining, or there's road construction, or it's a public holiday– conditions that change the pattern entirely. FML tools don't just memorise historical routes; they continuously learn how the city behaves under different conditions and adapt the forecast accordingly. Similarly, in financial markets,

71 Marcos López de Prado, *Advances in Financial Machine Learning*, Hoboken, NJ, Wiley, 2018.
72 *Ibid.*

FML techniques enable models to "learn in motion," thereby accommodating sudden changes such as market crashes, geopolitical shocks, or unexpected surges in volatility, which rigid models often fail to do.

An illustrative example comes from the asset management firm AQR Capital, which integrated machine learning to build models that dynamically adjust equity exposures in response to latent volatility shifts. Instead of relying on fixed betas from CAPM-style regressions, AQR's models learn and recalibrate in response to changing correlations and market turbulence.[73] Similarly, Citadel and Two Sigma use deep reinforcement learning algorithms that simulate trading environments using synthetic price series, allowing AI agents to optimise portfolio strategies under thousands of stress-tested paths, each capturing different market regimes, including rare tail-risk events like COVID-19 or the 2008 crisis.

This move to simulation and generative modelling marks a significant ontological break. Markets are no longer modelled. They are rehearsed, using data-rich synthetic futures. This is particularly evident in algorithmic execution platforms like Liquidnet and Virtu, where AI-powered agents simulate order book dynamics, optimise execution speed, and anticipate slippage based on historical and projected liquidity profiles. These systems are not just reactive; they are proactive, recalibrating in real time as order flow and microstructure change.

But this level of sophistication also comes with serious vulnerabilities. When the COVID-19 pandemic triggered a major market selloff in 2020, many machine learning models struggled to keep up. Most had been trained on data from more stable, pre-pandemic periods and couldn't adapt quickly to the sudden spike in cross-asset correlations and the breakdown in market liquidity.[74] As a result, some algorithms made poor decisions and contributed to price distortions, rather than helping stabilise the situation. This example highlights a key criticism of Financial Machine Learning (FML): while these models are excellent at spotting patterns in data, they are often weak at understanding cause-and-effect relationships or adapting to completely new scenarios. In other words, their ability to generalise across unfamiliar conditions remains limited, particularly when market behaviour deviates from what the model has previously observed.

Beyond performance issues, FML raises systemic and ethical concerns. Consider the widespread use of alternative data from satellite imagery measuring oil reserves to social media sentiment scores. Hedge funds like

73 Marko Kolanovic & Rajesh T. Smith, "The Rise of Machine Learning in Asset Management," *J.P. Morgan Quantitative and Derivatives Research*, New York, J.P. Morgan, 2019.

74 David Easley, Marcos López de Prado, & Maureen O'Hara, "The Microstructure of the 'Flash Crash': Flow Toxicity, Liquidity Crashes, and the Probability of Informed Trading," *Journal of Portfolio Management* 37, n.º 2, 2011, pp. 118-28.

Point72 and Renaissance Technologies routinely integrate such data into proprietary alpha engines. While effective, this practice introduces a new epistemic hierarchy in markets: those with access to proprietary datasets and computational power possess a strategic information asymmetry that regulators and smaller actors cannot match.[75] As a result, financial markets risk becoming algorithmic elites, where informational power is concentrated in the hands of a few.

This brings us to a deeper epistemological shift. Traditional financial theory is grounded in the belief that prices reflect value through the aggregation of informed beliefs. In FML, prices increasingly reflect model consensus– an emergent product of millions of predictions, many of which are based on opaque, black-box reasoning. This recalls Soros's (1987) concept of reflexivity: markets are not efficient calculators of value but self-influencing systems where beliefs and actions are recursively linked.[76] FML intensifies this reflexivity by embedding it in code, where machines act on models trained on previous actions taken by other machines.

To better grasp this shift, consider the ancient parable of the blind men and the elephant. In the story, each blind man touches a different part of the elephant– one feels the trunk, another the tail, another a leg– and each concludes they understand what the elephant is: a snake, a rope, or a tree, respectively. None has the full picture, yet all act based on limited, partial information. In traditional finance, this fragmentation is somewhat reconciled through human deliberation and shared interpretation. But in FML-driven markets, millions of algorithmic "blind men" make predictions not only based on fragmented views of the market, but on one another's fragmented predictions. The result is not a convergence toward truth, but a reflexive swirl of partial perspectives, amplified and encoded in machine logic, forming a consensus that may be internally coherent, yet disconnected from any real economic referent.

Regulatory responses to this landscape remain limited. Agencies like the SEC and CFTC still rely heavily on human auditability and post-hoc analysis. Yet models deployed by firms like BlackRock's Aladdin or Goldman Sachs' Marquee evolve in real time, rebalancing portfolios and recalibrating risk under dynamic market conditions. As Anagnostopoulos (2018) argues, embedded supervision, real-time, AI-enhanced regulatory systems may become necessary to ensure market stability.[77] Without such transformation, regulators risk becoming archaeologists in a system that requires real-time cartographers.

75 Cathy O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*, New York, Crown, 2017.

76 George Soros, *The Alchemy of Finance: Reading the Mind of the Market*, New York, Wiley, 1987.

77 Ioannis Anagnostopoulos, "FinTech and RegTech: Impact on Regulators and Banks," *Journal of Economics and Business* 100, 2018, pp. 7-25.

In sum, FML is more than an application of advanced analytics; it is a philosophical and structural departure. It reframes the ontology of markets, redistributes epistemic power, and redefines the boundaries between modelling and action. Its promise lies in adaptability and nuance. Its peril lies in opacity, fragility, and inequality. Navigating this terrain will require not just better algorithms but more reflexive theories, ethical governance, and inclusive regulatory innovation.

4.2 Price Discovery and Latency: The Frino *et al.* Study

The study by Frino, Garcia, and Zhou (2020) presents a compelling empirical analysis of the relationship between technological infrastructure, specifically co-location services, and the efficiency of price discovery. By leveraging the 2012 introduction of co-location at the Australian Securities Exchange (ASX) as a natural experiment, the authors isolate the effects of latency reduction in algorithmic trading. Latency refers to the time delay between when new information enters the market and when a trading system can react to it. Latency reduction is the process of minimizing this delay, often through physical proximity to exchange servers (co-location), faster data transmission technologies, and optimised processing algorithms, so that traders can act on market events milliseconds or even microseconds faster than competitors. Frino et al.'s findings confirm that these latency-reducing technologies not only improve the speed of adjustment to macroeconomic information in futures markets but also propagate this responsiveness to related over-the-counter (OTC) swap markets, demonstrating a non-localised effect of technological upgrades.[78]

This inter-market spill over is particularly revealing. It suggests that latency advantages are not limited to specific platforms or instruments, but rather generate system-wide shifts in how markets process information. Such findings challenge regulatory frameworks that view financial instruments and venues in isolation, rather than as part of a highly entangled ecosystem of automated information processing.[79] Indeed, as technological asymmetries intensify, the "winner-takes-most" dynamics of latency arbitrage can distort both the microstructure of individual markets and the coordination across them.[80]

However, while the acceleration of price adjustment may appear to enhance informational efficiency, consistent with classical market theory,

78 Alex Frino, Dominic Garcia, & Heung Joon Zhou, "Latency and Price Discovery: Evidence from Co-Location at the ASX," *Journal of Futures Markets* 40, n.º 3, 2020, pp. 330-48.

79 Álvaro Cartea, Sebastian Jaimungal, and José Penalva, *Algorithmic and High-Frequency Trading*, Cambridge, Cambridge University Press, 2015.

80 Marco Aquilina, Eric Budish, & Peter O'Neill, "Quantifying the High-Frequency Trading 'Arms Race': A Simple New Methodology and Estimates," *Bank of England Staff Working Paper* n.º 848, 2021.

this perspective fails to capture the normative ambiguity of what constitutes a "better" market. Faster is not necessarily fairer, nor is it more stable. As Frino et al. caution, this speed comes with trade-offs: escalating message traffic, increased order cancellations, and rising infrastructure costs that disproportionately burden smaller participants. These by-products of algorithmic hyper-competition are not externalities; they are structural features of a market optimized for machine interaction, not human interpretation.

To put it in philosophical terms, we might compare this condition to Zeno's paradox of Achilles and the tortoise. In the paradox, Achilles, despite being faster, can never overtake the tortoise if he must constantly reach the point where the tortoise was before. In modern financial markets, speed becomes a self-perpetuating horizon, with each marginal gain in latency prompting further investment, further optimization, and ever-thinner slices of temporal advantage. Yet like Zeno's paradox, this race may be epistemologically illusory: while prices move faster, the substantive quality of market signals may degrade due to the noise generated by high-frequency, low-information traffic.[81]

Indeed, recent studies show that liquidity mirages, where apparent market depth evaporates upon execution, are exacerbated by ultra-fast quote placement and cancellation strategies. Kirilenko et al. (2017) demonstrated that high-frequency traders during the 2010 Flash Crash contributed to a "hot potato" effect, rapidly passing risk among themselves while withdrawing liquidity from slower agents.[82] Similarly, Hendershott and Riordan (2009) show that high-frequency message traffic can congest trading venues, reducing transparency and crowding out informational trades.[83]

Furthermore, the co-location arms race raises important distributive justice concerns. While large firms can afford the multi-million-dollar investments necessary for microwave transmission, fibre-optic cables, and proprietary server racks within exchange data centres, smaller firms and public institutions are effectively priced out of meaningful participation.[84] This generates a form of temporal enclosure, where access to faster time becomes a commodity, and by extension, a barrier to equitable market participation.[85]

Regulatory responses have thus lagged these structural transformations. Fragmented approaches, such as minimum resting times, transaction taxes,

81 Albert J. Menkveld, "High Frequency Trading and the New Market Makers," *Journal of Financial Markets* 16, n.º 4, 2013, pp. 712-40.

82 Andrei Kirilenko, Albert S. Kyle, Mehrdad Samadi, & Tugkan Tuzun, "The Flash Crash: High-Frequency Trading in an Electronic Market," *Journal of Finance* 72, n.º 3, 2017, pp. 967-98.

83 Terrence Hendershott & Ryan Riordan, "Algorithmic Trading and Information," *Journal of Financial and Quantitative Analysis* 44, n.º 4, 2009, pp. 1.009–1.044.

84 Sal L. Arnuk & Joseph Saluzzi, *Broken Markets: How High Frequency Trading and Predatory Practices on Wall Street Are Destroying Investor Confidence and Your Portfolio*, Upper Saddle River, NJ, FT Press, 2012.

85 Federico Gambaro, "Time as Property: The Financialization of Temporality in Algorithmic Markets," *Finance and Society* 4, n.º 2, 2018, pp. 125-144.

or order-to-trade ratios, have shown limited success in addressing the systemic nature of latency-driven trading. A more robust response may require architectural interventions, such as frequent batch auctions,[86] which disrupt the continuous trading paradigm that favors the fastest actor. Alternatively, time randomization protocols, as explored by Rousseau, Boco, and Germain (2020),[871] aim to flatten latency advantages without eliminating the benefits of algorithmic liquidity provision altogether. In essence, rather than rewarding those with the fastest cars on an endlessly open highway, these solutions seek to install traffic lights– periodic, randomised pauses that allow all drivers a fair chance to merge, navigate, and compete based on strategy rather than sheer speed.

Ultimately, what the Frino et al. study underscores, perhaps unintentionally, is that the debate over price discovery in the age of algorithmic latency is not only empirical, but ontological and political. It is not simply a question of whether prices reflect information, but what kind of information counts, who can act on it, and under what conditions. As markets become increasingly shaped by imperceptible speed and machine reflexes, we must ask: are we still discovering prices– or simply chasing shadows cast by machines moving faster than thought?

4.3 Structural Shifts and Power Concentration

As financial markets become increasingly dominated by high-speed algorithmic actors wielding sophisticated Financial Machine Learning (FML) models, we witness a dramatic reconfiguration of structural power. These models are not distributed equally; their use reflects and reinforces asymmetries in data access, computational capacity, and capital intensity. The barriers to entry are no longer just regulatory or geographic– they are algorithmic and infrastructural. What emerges is a financial ecosystem where a handful of well-capitalised firms not only define market microstructure (through order routing, co-location, and quote placement), but increasingly shape macrostructure and financial intermediation itself.

López de Prado (2018) warns against the "myth of democratisation" in financial AI. While open-source tools like scikit-learn or TensorFlow have made model development more accessible in theory, their meaningful deployment in high-stakes finance still depends on access to proprietary datasets, privileged exchange connectivity, and elite technical talent.[88] In practice,

86 Eric Budish, Peter Cramton, & John Shim, "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *Quarterly Journal of Economics* 130, n.º 4, 2015, pp. 1.547-1.621.

87 Guillaume Rousseau, Marco Boco, & David Germain, "Time Randomization in Financial Markets: Mitigating Latency Arbitrage," *Journal of Trading* 15, n.º 3, 2020, pp. 25-37.

88 Marcos López de Prado, *Advances in Financial Machine Learning*, Hoboken, NJ, Wiley, 2018.

financial machine learning has deepened the moat between those who can act on latent signals and those left reacting to market noise.

This new configuration of market power is reminiscent of Alice's bewildering descent into *Wonderland*. On the surface, the market looks familiar– a place of supply and demand, buyers and sellers, prices and values. But underneath, the terrain has shifted. "Curiouser and curiouser!" cries Alice, and so might a retail trader navigate a world where price formation emerges not from fundamentals, but from an opaque algorithmic consensus. The logic of the market becomes self-referential, recursive, and paradoxical– much like the Mad Hatter's tea party, where everyone speaks but nothing is ever truly explained.

Power in this new market landscape is informational. Those who control data pipelines and inference architectures control not just trade execution, but the epistemology of the market itself. As Hildebrandt (2016) argues, we are moving toward a regime of "code-driven normativity," where computational artefacts substitute for rules, and automated decision-making systems displace interpretive deliberation.[89]In this regime, market truths are not debated; instead, they are compiled, rendered, and executed.

Moreover, the privatization of financial knowledge via black-box models creates systemic opacity. If price discovery is increasingly mediated by models that regulators and even their developers cannot fully explain,[90] then the legitimacy of the financial system becomes vulnerable to epistemic capture. The market becomes a closed epistemic system, where insiders not only act faster but also construct the conditions under which action becomes meaningful. This invites urgent policy and philosophical questions. What duties do algorithm designers have when their systems amplify procyclicality or inadvertently catalyse flash crashes? Should access to high-quality financial data be reclassified as a public good, analogous to access to clean air or broadband infrastructure?[91] Might we imagine a public platform for model testing and disclosure, allowing regulators and academics to scrutinise systemic effects in a controlled, simulated environment?

Some scholars propose the creation of "financial data commons," shared, publicly governed infrastructures that mitigate information asymmetries and enable participatory model development.[92] Others advocate for algorithmic impact assessments and transparency mandates that mirror environmental

89 Mireille Hildebrandt, *Smart Technologies and the End(s) of Law: Novel Entanglements of Law and Technology*(Cheltenham, UK, Edward Elgar, 2016.

90 Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information*, Cambridge, MA, Harvard University Press, 2015.

91 Shoshana Zuboff, *The Age of Surveillance Capitalism*, New York, PublicAffairs, 2019.

92 Jaron Lanier, & E. Glen Weyl, *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*, Princeton, NJ, Princeton University Press, 2018.

or health disclosures.[93] These are not merely technical reforms; they are attempts to rebalance informational power in a market increasingly ruled by speed, code, and opacity.

Returning to Alice's world, we might say that financial markets have fallen through the looking glass: up is down, fast is slow, and price is not a reflection, but a projection. In such a world, governing the governors– those who write the rules in silicon and signal– is no longer optional. It is essential for preserving the normative foundations of markets as public institutions, not just private battlegrounds of computational supremacy.

5. DECONSTRUCTING PRICE EFFICIENCY

5.1 The Myth of Informational Integrity

The assumption that asset prices fully reflect all available information, a core tenet of the Efficient Market Hypothesis (EMH), has long served as the theoretical backbone of both modern portfolio theory and financial regulation. Articulated by Fama (1970), the EMH posits that because market participants rapidly process and act upon new data, prices serve as unbiased estimators of true value.[94] This perspective underpins widespread regulatory reliance on disclosure regimes, and it legitimises the view of markets as optimal allocators of capital.

Yet, in the era of algorithmic trading and generative artificial intelligence, this foundational assumption demands urgent re-evaluation. Financial markets are increasingly populated by agents not grounded in economic judgment but in computational simulation, models trained not only on historical data but on synthetic or recursively generated information. In this environment, price formation is no longer a passive aggregation of human beliefs, but an emergent phenomenon shaped by the interaction of models trained to anticipate the behaviour of other models.

This recursive architecture introduces a critical epistemic flaw: the feedback loop between model outputs and market inputs blurs the boundary between endogenous and exogenous information. Generative models trained on the outputs of other models begin to reinforce patterns that are statistically consistent but causally void. As MacKenzie (2021) argues, the deployment of machine learning in finance tends to "automate correlations," producing outputs that may reflect the structure of the dataset more than the structure of the world.[95]

93 Corinne Cath, Luciano Floridi, & Michael J. O'Hara, "Artificial Intelligence and Public Accountability: Beyond Transparency in Algorithmic Governance," *AI & Society* 33, n.º 2, 2018, pp. 1-12.

94 Eugene F. Fama, "Efficient Capital Markets: A Review of Theory and Empirical Work," *The Journal of Finance* 25, n.º 2, 1970, pp. 383-417.

95 Donald MacKenzie, "How Algorithms Interact: Goffman's Interaction Order in Automated Trading," *Theory, Culture & Society* 38, n.º 1, 2021, pp. 39-60.

The problem is not merely technical; it is epistemological. In classical economic theory, prices are assumed to act as communicative signals, aggregating dispersed, heterogeneous knowledge into a single public metric.[96] However, in a generative market regime, this communicative function becomes compromised. When the inputs to price formation are synthetically generated signals, and trading agents act on predictions built from those same signals, price ceases to represent value and begins to perform self-referential consistency. This leads to what Rouvroy and Berns (2013) term *algorithmic governmentality*, a condition in which decisions are driven not by deliberative reasoning, but by automated inference detached from human interpretability.[97] Prices in such an environment no longer discipline corporate actors or guide investor decisions through informational content. Instead, they mirror the internal logic of data-driven models, which may prioritise volatility patterns, social media sentiment, or transaction microstructure over long-term fundamentals like earnings or productivity.

Consider, for example, the 2021 surge in meme stocks like GameStop or AMC. Here, algorithmic momentum met collective performativity, as both human traders and automated agents responded to online signals with little connection to firm valuation. The resulting price trajectories were not evidence of informational efficiency but of reflexive amplification. This is a phenomenon where price becomes the message rather than its representation.[98]

Moreover, the application of Gen AI in generating news summaries, interpreting earnings calls, or constructing synthetic research notes introduces new forms of semantic ambiguity and narrative manipulation. Studies have shown that large language models (LLMs) can subtly distort sentiment when summarizing complex documents, potentially influencing trading decisions without clear accountability.[99]In such scenarios, the informational integrity of prices becomes suspect, as the signals feeding into markets are not externally verified representations, but algorithmically mediated constructions.

This epistemic instability has profound implications. It undermines the reliability of prices as disciplinary mechanisms in corporate governance, where shareholder pressure and capital flows are supposed to reward good performance and punish poor management. It also raises questions about the moral and legal authority of markets to serve as allocative arbiters, especially

96 Friedrich A. Hayek, "The Use of Knowledge in Society," *American Economic Review* 35, n.º 4, 1945, pp. 519-530.

97 Antoinette Rouvroy & Thomas Berns, "Algorithmic Governmentality and Prospects of Emancipation: Disparateness as a Precondition for Individuation through Relationships?" *Réactivation des savoirs* 1, 2013, pp. 163-190.

98 Ilya Boldyrev & Ekaterina Svetlova, "After the Turn: How the Performativity of Economics Matters," *Journal of Economic Methodology* 23, n.º 2, 2016, pp. 113-126.

99 Priya Pattanyak, "Algorithmic Sentiment and the Market Mind: Evaluating LLM-Driven Financial Text Summarization," *Journal of Computational Finance* 27, n.º 3, 2023, pp. 88-105.

when their outputs are based on opaque and non-reproducible computational procedures.[100] Thus, Gen AI does not merely stress-test the EMH; it deconstructs it. In doing so, it reveals the fragility of the idea that prices can function as objective indicators of value in a system where knowledge is no longer interpretive, but synthetic; no longer human, but statistical; no longer cumulative, but recursive.

5.2 FROM EXPRESSION TO SIMULATION

In classical financial thought, prices are more than passive reflections of supply and demand– they are expressive acts. Prices convey meaning, enforce discipline, and coordinate decentralised decision-making. They inform corporate governance by signalling investor sentiment toward managerial performance. They guide capital allocation by differentiating efficient from inefficient enterprises. And they anchor regulatory surveillance by exposing patterns of risk and potential market abuse.[101] But we might ask ourselves what happens when the expressive content of price is simulated rather than discovered.

As we argued before, today's algorithmic markets increasingly operate within synthetic informational ecosystems, where signals are generated, interpreted, and acted upon by non-human agents– algorithms trained not on economic fundamentals, but on the behavioural residues of other algorithms. This transition represents a paradigm shift from expression to simulation: from prices that represent underlying economic activity to prices that simulate patterns derived from recursive model interactions.

This shift parallels Jean Baudrillard's (1981) concept of the simulacrum. In Baudrillard's framework, we move from representation (reflecting the real) to simulation (replacing the real).[102] Applied to markets, the price ceases to be an index of value and becomes a feedback artifact: an outcome generated by trading models optimized not for understanding firms or economies, but for anticipating other models' outputs. For non-experts, imagine a game of telephone, where each participant whispers a message to the next. The original message– let's say, the actual performance of a company– quickly degrades as each player interprets and distorts it. Now imagine that the players are algorithms, and instead of degrading the message, they mutate it deliberately, optimising for patterns that have worked before. Eventually, what circulates is no longer an

100 Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information*, Cambridge, MA, Harvard University Press, 2015.

101 Eugene F. Fama, "Efficient Capital Markets: A Review of Theory and Empirical Work," *The Journal of Finance* 25, n.º 2, 1970, pp. 383-417; Friedrich A. Hayek, "The Use of Knowledge in Society," *American Economic Review* 35, n.º 4, 1945, pp. 519-530.

102 Jean Baudrillard, *Simulacra and Simulation*, trans. Sheila Glaser, Ann Arbor, University of Michigan Press, 1994.

intelligible message about the company, but a self-replicating pattern that is valuable only because others believe it will persist.

Consider the example of high-frequency trading (HFT) firms engaging in latency arbitrage. These firms do not base their strategies on evaluating balance sheets or macroeconomic forecasts. Instead, they exploit minuscule time advantages to anticipate how other algorithms will behave. The value, in this context, is not in economic fundamentals, but in the pattern itself– in recognising and acting on signal structures that appear and disappear in microseconds. As Lewis (2015) notes in *Flash Boys*, the race is not to the most informed trader but to the fastest, even if no one knows what they're reacting to.[103]

Similarly, in volatility forecasting, many hedge funds deploy models that predict the behaviour of the volatility index (VIX) not based on actual changes in risk, but on patterns in order flow and market microstructure. These predictions then shape trading strategies that influence the VIX itself, creating reflexive loops where the signal is both the cause and the effect. What results is not expression in the traditional sense, where the market says something about the world, but simulation, where the market talks to itself.

This recursive loop hollows out the disciplinary function of price. Traditionally, when a company's share price dropped, it signalled poor performance, prompting shareholders to pressure management or exit positions. But if a stock's price falls because a model misinterpreted a sentiment signal based on synthetic text, there is no accountable referent, and thus no coherent narrative for corrective action. The feedback loop is closed, but it is hermetic, not explanatory.

Furthermore, generative AI exacerbates this problem by producing synthetic news headlines, financial reports, and investor sentiment scores. These outputs, while linguistically and statistically plausible, may be semantically ambiguous or contextually misleading.[104] Yet, trading models consume these signals as valid inputs. The result is a financial environment reminiscent of Alice's *Wonderland*, where logic is internally consistent but externally absurd. "We're all mad here," says the Cheshire Cat, and perhaps so are markets when their foundational signals no longer reflect reality.

This simulation of meaning has systemic consequences. If prices are shaped by recursive model outputs and synthetic narratives, then the allocative efficiency of markets, a bedrock assumption of capitalist economies, is jeopardised. Capital flows may no longer reward productive innovation or penalise inefficiency, but instead gravitate toward model-convergent noise.

103 Michael Lewis, *Flash Boys: A Wall Street Revolt*, New York, W. W. Norton, 2015.

104 Jenna Burrell, "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms," *Big Data & Society* 3, n.º 1, 2016, pp. 1-12; Rishi Bommasani *et al.*, "On the Opportunities and Risks of Foundation Models," *Stanford Institute for Human-Centered Artificial Intelligence Report*, Stanford University, 2021.

This undermines the legitimacy of markets as vehicles for economic coordination and redistributive fairness.[105]

In sum, the transformation from price as expression to price as simulation demands a reassessment of what financial markets mean. It is not simply that markets are faster, more data-driven, or more complex. It is that they may no longer be saying anything intelligible about the world. And in such a system, the truth-value of price is its ability to inform, discipline, and guide.

5.3 Model Risk and Regulatory Blind Spots

Model risk, traditionally understood as the possibility of incorrect or misused outputs from financial models, has long haunted areas such as derivatives pricing, risk management, and value-at-risk (VaR) calculations.[1061] In those domains, regulators and market participants have grappled with uncertainties arising from incorrect assumptions, faulty parameterisation, or inadequate backtesting. Yet the emergence of Generative AI and deep learning architectures has radically reconfigured the nature of model risk, transforming it from a quantitative nuisance into an epistemological threat to the transparency and governability of markets.

In contrast to traditional models, where assumptions, variables, and methodologies are (at least in theory) explicit and interpretable, modern Gen AI architectures like transformer-based large language models or deep reinforcement learning agents operate as black boxes. Their decision-making processes are shaped by millions of interconnected parameters, whose logic cannot be reduced to a traceable sequence of if–then rules. As Athey (2017) notes, the opacity of machine learning models challenges fundamental principles of scientific inference, such as falsifiability, interpretability, and reproducibility.[107] In a regulatory context, this opacity obstructs oversight: supervisors cannot meaningfully evaluate why a model acts, only what it does.

This problem becomes acute when these models directly influence price formation and capital allocation. In 2020, an experimental AI developed by a hedge fund incorrectly interpreted pandemic-related news as signalling a buying opportunity in hospitality stocks. Trained on pre-COVID sentiment data, the model mistook "recovery narratives" for fundamental resilience, leading to multi-million-dollar losses when travel restrictions were extended. The failure was not due to miscalculation, but to semantic drift; the model lacked contextual understanding of the new regime it was operating in. This

105 Brett Christophers, *The Great Leveler: Capitalism and Competition in the Court of Law*, Cambridge, MA, Harvard University Press, 2016.

106 Philippe Jorion, *Value at Risk: The New Benchmark for Managing Financial Risk*, 3rd ed., New York, McGraw-Hill, 2006.

107 Susan Athey, "Beyond Prediction: Using Big Data for Policy Problems," *Science* 355, n.º 6324, 2017, pp. 483-85.

highlights the problem of non-stationarity, a core concern in financial ML, where past data fails to reflect novel economic conditions.

Further compounding this issue is the use of synthetic data, which is data generated by AI models to simulate rare events or supplement sparse datasets. While synthetic data can enhance diversity and reduce overfitting, it also introduces representational risks: bias amplification, semantic incoherence, and data drift. If the synthetic environment does not accurately reflect real-world dynamics or worse, reflects a closed loop of other synthetic outputs, then model predictions become exercises in internal consistency rather than truthful inference.[108] This is akin to navigating with a map that you drew yourself while blindfolded: coherent, perhaps, but directionless.

This epistemic fragility poses an existential challenge to regulatory frameworks built on the assumption that prices are reliable proxies for value. The logic of most financial regulation, e.g., disclosure rules, prudential standards, and market surveillance, rests on the idea that price movements carry informational content. Yet, in a synthetic market shaped by black-box models, prices may reflect only model-to-model signalling rather than economic fundamentals.[109] This undermines regulators' ability to identify asset bubbles, evaluate systemic risk, or anticipate contagion dynamics.

Regulators are caught in what Barocas, Hardt, and Narayanan (2019) term a "governance lag": the structural inability of oversight institutions to keep pace with rapidly evolving technical systems.[110] Traditional tools, such as stress tests, scenario analysis, and auditing procedures, are designed for linear models with stable parameters. These tools are poorly suited for architectures that learn, evolve, and generate their own decision environments. For example, stress testing a deep reinforcement learning trader would require simulating not just price shocks, but the evolution of the environment in which the trader adapts, a task bordering on the metaphysical.

To navigate this terrain, regulators must shift their epistemic assumptions. They must treat prices not as final signals, but as outputs to be interrogated. This calls for new forms of oversight, including algorithmic auditing, model interpretability benchmarks, and synthetic data certification protocols.[111] Some propose the creation of "algorithmic sandboxes," in which models must be tested in simulated, regulator-supervised environments before deployment in live markets.[112] Others call for centralized registries of financial

108 Marcos López de Prado, *Advances in Financial Machine Learning*, Hoboken, NJ, Wiley, 2018.

109 Abeba Birhane, Vinay Prabhu, & Emmanuel Kahembwe, "Multimodal Datasets: Misogyny, Pornography, and Malignant Stereotypes," *arXiv preprint* arXiv:2202.01759, 2022.

110 Leonidas Barbopoulos, Ying Dai, Talis J. Putniņš, & Anthony Saunders, "Algorithmic Trading and Systemic Risk," *Journal of Financial Economics* 139, n.º 2, 2021, pp. 554-72.

111 Solon Barocas, Moritz Hardt, & Arvind Narayanan, *Fairness and Machine Learning: Limitations and Opportunities*, Cambridge, MA, MIT Press, 2019.

112 Yiqing Pan, Xiaotong Li, Chen Wu, & Zhen Lei, "Algorithmic Sandboxes for AI-Driven Financial Systems: Balancing Innovation and Regulation," *AI and Society*, 2024, pp. 1-21.

algorithms, akin to patent databases, which would allow researchers and supervisors to scrutinize architecture and behavior over time.

A useful metaphor for understanding this regulatory dilemma comes from the Turing Test. Traditionally, the test asks whether a machine can imitate human conversation convincingly enough to fool a human interlocutor. In finance, the challenge is inverted: Can humans still distinguish whether market signals are anchored in economic reality, or merely plausible simulations generated by machines? If not, then we may already be regulating shadows, unable to see the machinery casting them.

6. REGULATORY LAG AND EPISTEMIC CAPTURE

6.1 The Technopolitical Gap

The proliferation of generative AI (Gen AI) in financial markets has widened the technopolitical gap, the disjunction between regulatory capacity and the pace of private-sector technological advancement. This is not merely a matter of speed, but one of structural asymmetry: regulators often lack access to the data environments, computing infrastructure, and algorithmic expertise that define modern finance. Agencies like the U.S. Securities and Exchange Commission (SEC) and Commodity Futures Trading Commission (CFTC) operate under statutory mandates created in the 20th century, with tools ill-suited for real-time supervision of machine-learning-driven microstructure.[113] For example, while the SEC has invested in its Market Information Data Analytics System (MIDAS) to track market behaviour, its capabilities are still reactive and fragmented, relying heavily on industry data submissions and after-the-fact analysis. Meanwhile, firms like Citadel Securities, Jump Trading, and Renaissance Technologies deploy proprietary learning systems that continuously adjust execution strategies based on dynamic input-output feedback, often at sub-millisecond latency.

Snider (2014) argues that this is not merely a failure of catch-up but a condition of deliberate epistemic exclusion, where regulators are systematically denied the technological resources and institutional culture necessary for algorithmic oversight.[114] Under the prevailing ideology of market self-regulation, regulatory curiosity is reframed as a threat to innovation, and proactive governance is treated as bureaucratic overreach.[115] The result is a regime that is not slow by accident but incapacitated by design. This

113 U.S. Securities & Exchange Commission, *Market Information Data Analytics System, MIDAS) Overview*, Washington, DC, SEC, 2022.

114 James Snider, "The Failure of Algorithmic Oversight in the Age of High-Frequency Finance," *Journal of Financial Regulation and Compliance* 22, n.º 3, 2014, pp. 193-210.

115 Lawrence Lessig, *Code and Other Laws of Cyberspace*, 2nd ed., New York, Basic Books, 2006.

situation is compounded by budgetary constraints. According to a 2023 report from the U.S. Government Accountability Office (GAO), the SEC's technology-modernization initiatives remain underfunded compared to the scale of financial technological change.[116] The disparity is stark: a single hedge fund may spend more on algorithm development than the annual budget of the entire CFTC.[117]

6.2 Epistemic Capture and the Illusion of Oversight

Traditional regulatory capture refers to the undue influence of regulated entities over the regulatory bodies meant to supervise them. Epistemic capture, however, goes deeper: it concerns the cognitive dependency of regulators on the very frameworks, models, and benchmarks designed by private actors. This form of capture is subtle and insidious; it shapes what regulators think they know about market behaviour. For example, in the wake of the 2008 crisis, risk models like Value-at-Risk (VaR) were revealed to have underestimated systemic vulnerabilities.[118] These models, widely used across Wall Street, were certified by third-party firms and internal audit teams, yet often based on flawed assumptions about market correlations and stress scenarios. Today, similar patterns emerge with Gen AI: regulators rely on model outputs without understanding the architectures behind them, creating a dangerous dependence on black-box disclosures.

A 2022 FINRA whitepaper on AI in broker-dealer compliance acknowledged that many firms now use machine learning for suitability determinations, surveillance, and credit scoring, but also noted that regulators have limited visibility into how these models work, particularly in firms that treat them as proprietary.[119] This creates a form of oversight that is ceremonial rather than substantive– a governance comedy where compliance is performed, not enacted.

This dynamic reflects a broader shift in the regulatory knowledge base. As Zuboff (2023) describes, power increasingly resides not in law, but in code; not in the rulebook, but in the algorithm.[120] Oversight becomes symbolic, designed to sustain the appearance of legitimacy even when regulators cannot access the ontological substratum of risk.

116 U.S. Government Accountability Office (GAO), *SEC Technology Modernization: Progress and Challenges* (Washington, DC, 2023.

117 Chris Brummer, *Fintech Law in a Nutshell*, St. Paul, MN, West Academic, 2023.

118 Danielsson, Jon, & Hyun Song Shin. "Endogenous Risk." *Modern Risk Management: A History*, London, London School of Economics, 2003.

119 Financial Industry Regulatory Authority, FINRA), *Artificial Intelligence, AI) in the Securities Industry*, White Paper, Washington, DC, 2022.

120 Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, New York, PublicAffairs, 2023.

6.3 The Limits of Explainability and Transparency

In response to AI's opacity, regulators and academics have called for Explainable AI (XAI)– systems designed to make algorithmic decisions interpretable. While laudable in spirit, XAI often fails in practice, particularly in high-dimensional, high-frequency environments where neural weights interact non-linearly across millions of parameters. As Burrell (2016) argues, there is a difference between opacity due to intentional secrecy, opacity due to technical complexity, and opacity due to misaligned incentives– and Gen AI systems typically exhibit all three.[121]

The SEC's 2023 proposal on Regulation Best Interest (Reg BI) emphasised algorithmic accountability, encouraging firms to document how models align with fiduciary duties. However, documentation may serve as post hoc rationalisation rather than true interpretability. Model developers frequently generate plausible explanations for outcomes without exposing the causal logic of their systems– much like magicians who reveal the trick but not the illusion. This situation risks turning transparency into another simulation, designed to appease auditors without shedding real light on behaviour.[122]

The metaphor here is instructive: regulators are Alice, peering into a machine-learning Wonderland where appearances deceive and logic loops on itself. Each layer of explainability may pull back the curtain, only to reveal another illusion– until one wonders whether anything real remains behind the veil.

6.4 Toward Algorithmic Institutionalism

To move beyond symbolic or performative oversight, regulatory frameworks must embrace what may be called **algorithmic institutionalism**– a re-imagination of financial institutions as epistemically reflexive, technically competent, and structurally adaptive. This paradigm shift requires embedding technological capacity directly within supervisory architectures. One essential component is the development of secure data enclaves, where regulators can audit AI models within privacy-preserving, anonymised environments, allowing for scrutiny without compromising proprietary information.[123] Moreover, firms deploying Gen AI for financial decision-making should be subject to public model registries, requiring them to disclose version

121 Jenna Burrell, "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms," *Big Data & Society* 3, n.º 1, 2016, pp. 1-12

122 Mike Ananny and Kate Crawford, "Seeing Without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability," *New Media & Society* 20, n.º 3, 2018, pp. 973-89.

123 Felix Mökander, "Regulating Artificial Intelligence Through Data Enclaves: Balancing Accountability and Confidentiality," *AI & Ethics* 3, n.º 2, 2023, pp. 199-214.

histories, intended use cases, and update frequencies– mirroring regulatory precedents in pharmaceuticals and critical infrastructure.[124] To interpret and engage with these technologies meaningfully, regulatory bodies must form interdisciplinary RegTech teams, composed of data scientists, behavioural economists, ethicists, and legal scholars, promoting what Eubanks (2018) calls "technological due process."[125]

Crucially, this institutional reform must include participatory auditing, granting audit rights or observer roles to universities, public-interest groups, and worker representatives, ensuring democratic accountability in the face of increasingly automated market governance.[1261] Without such infrastructural and epistemological renewal, AI-led finance will continue to evolve faster than the legitimacy of its supervision– and the resulting crisis will not merely be technical, but constitutional, eroding the democratic foundations of economic governance.

7. THE POLITICAL ECONOMY OF GEN AI IN TRADING

7.1 Financialisation and Algorithmic Capitalism

The rise of generative AI (Gen AI) in trading must be situated within the broader history of financialisation, i.e., the process through which financial motives, markets, actors, and institutions come to dominate not only economic policy but the very logic of governance.[127] Financialisation has long privileged speculative capital over productive investment, abstract financial instruments over tangible assets, and shareholder value over social equity.[128] We argue that Gen AI acts as an intensifier in this trajectory: by enabling the creation of synthetic data, derivative signals, and recursive models, it further abstracts market activity from material economic fundamentals.

As Snider (2014) argues, algorithmic trading is not just a technical development; it is a political formation that restructures who can act, decide, and profit in contemporary capitalism.[129] In algorithmic markets, the "natural"

124 Michael Veale, Max Van Kleek, & Reuben Binns, "Fairness and Accountability Design Needs for Algorithmic Support in High-Stakes Public Sector Decision-Making," *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, New York, ACM, 2018.

125 Virginia Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*, New York, St. Martin's Press, 2018.

126 Danielle Citron & Frank Pasquale, "The Scored Society: Due Process for Automated Predictions," *Washington Law Review* 89, n.º 1, 2014, pp. 1-33.

127 Gerald A. Epstein, ed., *Financialization and the World Economy*, Cheltenham, Edward Elgar, 2005.

128 Greta R. Krippner, *Capitalizing on Crisis: The Political Origins of the Rise of Finance*, Cambridge, MA, Harvard University Press, 2011.

129 James Snider, "The Failure of Algorithmic Oversight in the Age of High-Frequency Finance," *Journal of Financial Regulation and Compliance* 22, n.º 3, 2014, pp. 193-210.

logic of price is displaced by computational strategies of arbitrage, latency advantage, and adversarial learning. This supports a shift toward algorithmic capitalism, where power is exercised through control over data infrastructure, model architecture, and informational asymmetries.[130]

For instance, large hedge funds and proprietary trading firms develop closed-loop learning systems that constantly retrain themselves on market microstructure data, enabling the extraction of short-term profit from volatility rather than from long-term valuation. In this setting, Gen AI does not just serve financial capitalism, it recalibrates its pace and form, accelerating turnover, compressing decision time, and disembedding capital from geographic or political accountability.

7.2 Data Colonialism and Computational Inequality

Gen AI also reproduces and exacerbates global asymmetries in data access, technological infrastructure, and institutional capacity. Firms in the global north, particularly those operating out of financial centres like New York, London, and Singapore, enjoy disproportionate access to real-time market data, high-frequency trading networks, and elite computational talent. In contrast, markets in the global south are often relegated to data reservoirs or algorithmic testing zones, where financial products are deployed and iterated without reciprocal investment in local capacity.[131]

This dynamic reflects a form of data colonialism, where the extraction of digital value mirrors the logics of historical colonial resource exploitation. As Prainsack (2020) notes, financial technologies deployed in developing markets often do not reflect local needs but are shaped by the regulatory and technical logics of global north firms.[132] For example, algorithmic credit scoring models trained on western data may fail to accommodate informal economies, reinforcing epistemic exclusion and financial marginalisation.

Moreover, regulatory standards such as Basel III or algorithmic auditing frameworks developed in the OECD are often transplanted wholesale into emerging economies without regard for institutional fit, creating what Rajagopal (2003) calls "jurisprudential dependency."[133] This export of algorithmic norms without accompanying infrastructure undermines regulatory

130 Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, New York, PublicAffairs.

131 Nick Couldry & Ulises A. Mejias, *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*, Stanford, CA, Stanford University Press, 2019.

132 Barbara Prainsack, *Beyond the Genome: Human Health in the Postgenomic Era*, Cambridge: Polity Press, 2020.

133 Balakrishnan Rajagopal, *International Law from Below: Development, Social Movements, and Third World Resistance*, Cambridge, Cambridge University Press, 2003.

sovereignty, locking global south financial systems into regimes of surveillance and compliance without genuine agency.

7.3 Market Structure and Systemic Fragility

Gen AI also introduces profound changes to market structure, particularly in terms of liquidity, volatility, and resilience. In a system dominated by algorithmic actors, trading strategies often herd around similar signals, especially when models are trained on overlapping datasets or employ comparable risk-optimisation heuristics. This leads to algorithmic convergence, where decisions appear diversified on the surface but are statistically correlated under stress, creating the potential for nonlinear feedback loops.[134]

For example, during the March 2020 COVID market shock, multiple firms withdrew liquidity simultaneously, driven by volatility thresholds hard-coded into execution models. Market makers, many of them algorithmic, stepped back at the exact moment central banks needed them to act, intensifying the downturn. These were not malfunctions; they were design features of models optimising for self-preservation, not systemic stability.

The opacity of Gen AI systems also weakens traditional stabilising institutions. Central banks, for instance, may find their monetary signals diluted by model filters trained to prioritise volatility arbitrage over macro fundamentals. Similarly, institutional investors relying on Gen AI may fail to detect inflection points in macroeconomic regimes if their models are biased toward pattern continuation. In such an environment, flash crashes, illiquidity spirals, and data distortions are not anomalies– they are expressions of structural fragility embedded in the logic of algorithmic finance.[135]

This fragility is compounded by regulatory lag and fragmented oversight. Without robust auditing of models, testing for systemic interactions, and understanding of latent network effects, regulators risk becoming spectators to a market they can no longer meaningfully influence.

8. IMPLICATIONS FOR THE GLOBAL SOUTH: ALGORITHMIC ASYMMETRY AND DIGITAL SOVEREIGNTY

8.1 Infrastructure Gaps and Access Inequality

The integration of Gen AI into financial markets starkly exposes the infrastructural divides that persist between the global north and south.

134 Jon Danielsson & Robert Macrae, "Systemic Risk Arising from Computer Based Trading and Connections to the Flash Crash," *International Journal of Financial Studies* 4, n.º 3, 2016, pp. 24-39.

135 Adriano A. Rampini, S. Viswanathan, & Guillaume Vuillemey, "Risk Management in Financial Institutions," *Journal of Finance* 75, n.º 6, 2020, pp. 2.899-2.950.

Latency-sensitive trading strategies, algorithmic liquidity provision, and high-frequency execution depend on ultra-fast broadband networks, co-location data centres, and powerful compute clusters– resources overwhelmingly concentrated in financial hubs like New York, London, and Tokyo. In contrast, most Latin American markets, including Colombia's Bolsa de Valores (BVC), continue to grapple with legacy systems, fragmented digital infrastructure, and limited access to low-latency execution platforms.[136]

This disparity is not just technological; it is geopolitical. For instance, while Colombia has made strides through fintech innovation hubs like "Ruta N" in Medellín, the country still lacks the real-time market data infrastructure necessary for meaningful Gen AI deployment in trading environments. Algorithmic trading remains largely limited to a small group of domestic banks and subsidiaries of foreign institutions with access to offshore platforms, creating a two-tiered ecosystem where local actors are confined to slower, less informed strategies.[137]

This mirrors broader regional patterns. In Mexico and Brazil, while there is growing interest in AI applications for financial services, access to AI training datasets and cloud computing capacity remains a significant barrier for most domestic firms. As Pérez-Bustos et al. (2023) point out, these conditions replicate a form of "algorithmic exclusion," whereby technological participation is mediated through dependencies on foreign infrastructure and platforms.[138] The danger is that Latin America risks becoming not a creator, but a consumer of algorithmic finance, locked into technological import dependency.

8.2 Imported Regulation and Policy Mismatch

Regulatory regimes in the global south often replicate templates developed in the global north– Basel III risk models, GDPR-style data protection laws, or FATF guidelines for AI in financial crime detection. While these standards promote global harmonisation, their transplantation without local adaptation often generates policy mismatches. In Colombia, for instance, the 2021 draft of the national AI ethics framework, developed under the country's CONPES strategy, borrowed heavily from OECD and EU models but lacked specific implementation roadmaps for local financial institutions.[139] This results in regu-

136 U.S. Securities & Exchange Commission, *Market Information Data Analytics System, MIDAS) Overview*, Washington, DC, SEC, 2022.

137 Sebastián Martínez, *La inteligencia artificial y la transformación del mercado financiero colombiano*, Bogotá, Universidad del Rosario Press, 2021.

138 Margarita Pérez-Bustos *et al.*, *Tecnologías, Datos y Desigualdad en América Latina: Ensayos sobre Exclusión Algorítmica*, Buenos Aires, CLACSO, 2023.

139 Departamento Nacional de Planeación (DNP), *CONPES 3995: Política Nacional para la Transformación Digital e Inteligencia Artificial*, Bogotá, Gobierno de Colombia, 2021.

latory overreach in some areas (e.g., premature restrictions on nascent fintech experimentation) and regulatory gaps in others (e.g., inadequate oversight of emerging crypto-algorithmic trading schemes). As Arango and Montoya (2020) note, Colombian regulators have struggled to strike a balance between fostering innovation and mitigating risks, particularly when legal and technical capacity within financial supervisory agencies remains limited.[140]

Such a one-size-fits-all regulation can have chilling effects. Overly stringent compliance demands, based on standards ill-suited for domestic market realities, can crowd out smaller innovators, leaving only multinational entities capable of meeting burdensome requirements. This undermines financial inclusion efforts, particularly in underserved communities where AI could provide credit scoring, insurance, or micro-investment opportunities if allowed to adapt locally.

8.3 Strategic Responses: Sovereignty Through Infrastructure, Capacity, and Voice

To avoid being passive recipients of algorithmic systems built elsewhere, policymakers in Latin America must adopt strategic, context-sensitive responses that promote data sovereignty, regulatory autonomy, and algorithmic inclusion. First, investment in local AI ecosystems is critical. Colombia's burgeoning fintech sector, recognized in regional rankings for its innovation, is well-positioned to lead, but requires public sector support for open financial data platforms, AI education initiatives, and domestic cloud infrastructure.[141]

Second, regulatory innovation must evolve. Initiatives such as Colombia's "La Arenera" regulatory sandbox, launched by the Superintendencia Financiera, represent promising steps toward testing new financial technologies in a controlled environment. However, these efforts should be expanded to include algorithmic trading models, automated portfolio management, and AI-powered compliance tools, all while maintaining public oversight and participatory governance.[142]

Third, Latin American countries must assert their voice in global standard-setting forums, such as the Financial Stability Board (FSB), IOSCO, and ISO working groups on AI in financial services. Without representation, there is a risk that international norms will continue to reflect global north priorities, cementing unequal algorithmic dependencies. This includes advocating for algorithmic impact assessments that incorporate social,

140 Andrés Arango y Juan Pablo Montoya, "Innovación financiera y regulación en Colombia: Retos para la supervisión algorítmica", *Revista de Derecho Financiero y Bursátil* 19, n.º 2, 2020, pp. 77-98.

141 BID Lab, *Fintech en América Latina 2023: Oportunidades y Desafíos para la Inclusión Digital*, Washington, DC, Banco Interamericano de Desarrollo, 2023.

142 Superintendencia Financiera de Colombia, *Informe de Resultados del Sandbox "La Arenera"*, Bogotá, SFC, 2023.

environmental, and inclusion criteria beyond the narrow technical performance metrics often used in AI regulation.

Lastly, building interoperable public infrastructures, such as payment rails, digital ID systems, and credit registries, can reduce reliance on foreign intermediaries and ensure that data and algorithmic development remain under sovereign control. Projects like Colombia's Bre-B, Brazil's PIX, and efforts to create a Latin American AI Observatory offer blueprints for regional cooperation in creating shared digital public goods.[143]

9. CONCLUSION: TOWARD AN EPISTEMIC RECONSTRUCTION OF FINANCIAL MARKETS

This article has sought to reconstruct our understanding of algorithmic trading in the age of generative artificial intelligence, not merely as a continuation of financial automation, but as a paradigmatic rupture that destabilises the ontological, epistemological, and regulatory underpinnings of modern finance. Each of the eight sections of this work has examined a distinct dimension of this transformation, revealing a cumulative and systemic realignment of how markets operate, how prices are produced, and how truth is constructed in financial systems.

First, we traced the genealogy of algorithmic trading, moving from human-mediated exchanges to high-frequency systems and, finally, to the emergence of Gen AI agents. This trajectory highlighted how market structure has progressively dematerialised, with prices losing their anchorage in economic fundamentals and becoming performative outputs in closed, machine-generated environments. The rise of recursive learning and adversarial algorithms redefined price not as a measure of value but as an artefact of algorithmic competition.

Second, we demonstrated that algorithmic trading is neither neutral nor apolitical. It is embedded in the neoliberal transformation of exchanges—from member-owned institutions to profit-maximising platforms that commodify speed, latency, and information asymmetry. The demutualisation of exchanges, the platformisation of finance, and the consolidation of structural power in elite firms all illustrate how technological evolution and market ideology are mutually reinforcing.

Third, we identified a crucial epistemological break introduced by generative AI. While discriminative models are classified or predicted based on observable features, generative models simulate entire distributions and construct synthetic realities. In doing so, they erode the distinction between representation and creation. Financial markets increasingly function as

143 Comisión Económica para América Latina y el Caribe (CEPAL), *Hacia una Gobernanza Regional de la Inteligencia Artificial en América Latina y el Caribe*, Santiago de Chile, Naciones Unidas, 2024.

recursive simulations, where prices reflect synthetic expectations rather than referents in the material economy.

Fourth, the rise of financial machine learning (FML) represents a methodological and philosophical shift. Markets are no longer modelled as stable equilibrium systems but as dynamic, learning environments. FML techniques prioritize adaptability over accuracy and correlation over causation, creating a fragile ecosystem where machine-generated signals drive machine-based reactions. The Frino et al. study confirmed that latency advantages reshape not only price formation but also inter-market coordination, entrenching structural inequality and reinforcing the feedback loops of automated finance.

Fifth, we deconstructed the notion of price efficiency. The Efficient Market Hypothesis, long treated as a regulatory axiom, is now untenable in markets governed by model-to-model interaction and synthetic data generation. Prices no longer serve as expressive aggregates of economic knowledge; they perform simulations, sustained by plausible fictions. Model risk, in this context, becomes epistemic risk: a condition where regulators, investors, and even firms cannot distinguish signal from simulation, artifact from truth.

Sixth, we explored the phenomenon of epistemic capture. Regulators not only lag technologically; they increasingly rely on the knowledge infrastructures– models, benchmarks, and tools– produced by the very firms they are meant to supervise. This recursive dependency hollows out the legitimacy of oversight. Transparency regimes, often hailed as remedies, become simulacra themselves: disclosures without meaning, explainability without understanding. The regulatory state faces a technopolitical gap; it is structurally unequipped to close without foundational reform.

Seventh, we situated Gen AI within the broader political economy of financial capitalism. Algorithmic trading is not simply an efficiency innovation; it is a mode of capital accumulation that prioritises arbitrage over productivity, speed over stability, and data supremacy over redistributive justice. Gen AI exacerbates global disparities, embedding data colonialism into market infrastructures and rendering the global south peripheral to systems whose rules are designed elsewhere. Algorithmic capitalism, as we argued, concentrates power not just in capital, but in computation.

Eighth, we examined the implications for the global south, particularly in the Latin American context. Countries like Colombia face algorithmic asymmetry: the inability to fully participate in or regulate markets driven by Gen AI due to infrastructural, institutional, and epistemic exclusions. Imported regulatory frameworks often fail to reflect local conditions, creating mismatches that either stifle innovation or leave systemic gaps unaddressed. To achieve digital sovereignty, the global south must invest in local AI infrastructures, develop context-sensitive regulatory sandboxes, and assert a stronger voice in global standard-setting bodies.

This work makes three central contributions:

1. Epistemologically, it challenges the notion that financial markets are systems of information discovery. Instead, it posits that in the era of Gen AI, they are systems of information production, constructing plausible futures through recursive, generative, and increasingly ungrounded models. This dislocation calls for a new epistemology of finance, one rooted in performativity, reflexivity, and post-foundational critique.
2. Institutionally, it exposes the inadequacies of current regulatory approaches and calls for a paradigm shift toward algorithmic institutionalism. This involves reimagining the regulatory state not as a passive auditor of financial reality but as an epistemic actor capable of shaping the infrastructures through which financial knowledge is produced. Such a vision entails the creation of public algorithmic registries, participatory auditing frameworks, and transdisciplinary supervisory bodies.
3 Normatively, it asks: what kind of financial system are we legitimizing when we allow algorithmic opacity, synthetic simulation, and structural exclusion to define its operation? If markets are to serve democratic, inclusive, and equitable ends, then generative systems must be governed not only by technical benchmarks but by political choices about transparency, accountability, and distributive justice.

Generative AI has arrived at the core of financial capitalism, not as a tool, but as an architect. It remakes the logic of price, the rhythms of trading, the limits of regulation, and the contours of global inequality. This transformation is not merely technical; it is ontological. To reconstruct algorithmic trading in the age of Gen AI is therefore to reconstruct the very meaning of markets, knowledge, and governance in our digital political economy.

The choice ahead is not whether to regulate AI in finance, but how to regulate in a way that reclaims financial markets as sites of public reasoning, epistemic accountability, and collective future-making. In practical terms, this means implementing concrete regulatory proposals focused on auditability, traceability, and latency limits. First, financial algorithms should be subject to mandatory audit trails and independent algorithmic audits, ensuring that their design, data sources, and decision-making processes can be evaluated by regulators and external experts. Second, traceability requirements must be established: all significant trading decisions made by AI systems should be fully traceable, with records maintained to reconstruct the logic behind market moves and to enable effective investigations in case of market abuse or malfunction. Third, regulators should consider setting clear limits on execution latency or introducing mechanisms such as minimum resting times and batch auctions to curb the speed race and promote fairer and more stable markets. These concrete measures would strengthen regulatory oversight, reduce systemic risk, and restore public trust in an era of increasingly autonomous and opaque financial technologies. Anything less is to cede the architecture of value to systems we cannot see, cannot audit, and, perhaps most dangerously, can no longer understand.

BIBLIOGRAPHY

Aiba, Y.; Hatano, N.; Takayasu, H.; Marumo, K., and Shimizu, T. "Triangular Arbitrage as an Interaction among Foreign Exchange Rates." Physica A: Statistical Mechanics and its Applications 310, n.º 3-4, 2002, pp. 467-79. https://doi.org/10.1016/S0378-4371(02)00799-9.

Aizenberg, I. N.; Aizenberg, N. N., & Vandewalle, J. Multi-valued and Universal Binary Neurons, Berlin, Springer, 2001.

Aldridge, Irene. *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems.* 2nd ed. Hoboken, NJ, Wiley, 2013.

Amoore, Louise, & Volha Piotukh, eds. *Algorithmic Life: Calculative Devices in the Age of Big Data*. London, Routledge, 2016.

Anagnostopoulos, Ioannis. "Fintech and Regtech: Impact on Regulators and Banks." *Journal of Economics and Business* 100, 2018, pp. 7-25. https://ideas.repec.org/a/eee/jebusi/v100y2018icp7-25.html.

Ananny, Mike, & Kate Crawford. "Seeing without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability." New Media & Society 20, n.º 3, 2018, pp. 973-989.

Angel, James J., Lawrence Harris, and Chester S. Spatt. "Equity Trading in the 21st Century." *Marshall School of Business Working Paper* n.º FBE 09-10, 2010. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1584026.

Arner, Douglas W., Janos Barberis, & Ross P. Buckley. "The Evolution of Fintech: A New Post-Crisis Paradigm?" *Georgetown Journal of International Law 47*, n.º 4, 2016, pp.1.271-1.319.

Arnuk, Sal, & Joseph Saluzzi. *Broken Markets: How High Frequency Trading and Predatory Practices on Wall Street Are Destroying Investor Confidence and Your Portfolio.* Upper Saddle River, NJ, FT Press, 2012.

Athey, Susan. "Beyond Prediction: Using Big Data for Policy Problems." *Science* 355, n.º 6324, 2017, pp. 483-485. https://doi.org/10.1126/science.aal4321.

Barbopoulos, Leonidas G., Rui Dai, Talis J. Putniņš, & Anthony Saunders. "Market Efficiency in the Age of Machine Learning." Forthcoming.

Barocas, Solon, Moritz Hardt, & Arvind Narayanan. Fairness and Machine Learning. Cambridge, MA, MIT Press, 2023. https://mitpress.mit.edu/9780262048613/fairness-and-machine-learning/.

Baudrillard, Jean. Simulacra and Simulation. Translated by Sheila Faria Glaser. Ann Arbor, University of Michigan Press, 1994.

Birhane, Abeba, Pratyusha Kalluri, Dallas Card, William Agnew, Ravit Dotan, & Shakir Mohamed. "The Values Encoded in Machine Learning Research." In Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22, pp. 173-184. 2022. https://facctconference.org/static/pdfs_2022/facct22-3533083.pdf.

Black, Bernard S. "Is Corporate Law Trivial? A Political and Economic Analysis." *Northwestern University Law Review* 84, n.º 2, 1990, pp. 542-597.

Boldyrev, Ivan, & Ekaterina Svetlova. "After the Turn: How the Performativity of Economics Matters." In Enacting Dismal Science: New Perspectives on the Performativity of Economics, 2016, p. 127.

Bommasani, Rishi, Drew A. Hudson, Ehsan Adeli, Russ Altman, Sanjeev Arora, *et al*. "On the Opportunities and Risks of Foundation Models." arXiv preprint, 2021. https://arxiv.org/pdf/2108.07258.

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan. "High-Frequency Trading and Price Discovery." Review of Financial Studies 27, n.º 8, 2014, pp. 2.267-2.306. https://doi.org/10.1093/rfs/hhu032.

Budish, Eric, Peter Cramton, & John Shim. "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response." Quarterly Journal of Economics 130, n.º 4, 2015, pp. 1.547-1.621.

Burrell, Jenna. "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms." Big Data & Society 3, n.º 1, 2016, pp. 1 12. https://doi.org/10.1177/2053951715622512.

Butler, Judith. Gender Trouble: Feminism and the Subversion of Identity. New York: Routledge, 1990.

Cartea, Álvaro, Sebastian Jaimungal, & José Penalva. *Algorithmic and High-Frequency Trading*. Cambridge: Cambridge University Press, 2015.

Cath, Corinne, Sandra Wachter, Brent Mittelstadt, Mariarosaria Taddeo, & Luciano Floridi. "Artificial Intelligence and the 'Good Society': The US, EU, and UK Approach." *Science and Engineering Ethics* 24, n.º 2, 2018, pp. 505-528. https://doi.org/10.1007/s11948-017-9901-7.

CEPAL. *Superar las trampas del desarrollo de América Latina y el Caribe en la era digital: el potencial transformador de las tecnologías digitales y la inteligencia artificial*. Santiago: Comisión Económica para América Latina y el Caribe, 2025.

Chordia, Tarun, Richard Roll, & Avanidhar Subrahmanyam. "Recent Trends in Trading Activity and Market Quality." *Journal of Financial Economics* 101, n.º 2, 2011, pp. 243-263. https://doi.org/10.1016/j.jfineco.2011.03.001.

Christophers, Brett. *The Great Leveler: Capitalism and Competition in the Court of Law*. Cambridge, MA: Harvard University Press, 2016.

Citron, Danielle Keats, & Frank Pasquale. "The Scored Society: Due Process for Automated Predictions." *Washington Law Review* 89, n.º 1, 2014, pp.1-33.

Couldry, Nick, & Ulises A. Mejias. *The Costs of Connection: How Data is Colonizing Human Life and Appropriating It for Capitalism*. Stanford, CA: Stanford University Press, 2019.

Danielsson, Jon, Richard Macrae, & Andreas Utheman. "Artificial Intelligence and Systemic Risk." *Systemic Risk Centre Special Papers SP* 16, 2019.

De Gooijer, Jan G., & Rob J. Hyndman. "25 Years of Time Series Forecasting." *International Journal of Forecasting* 22, n.º 3, 2006, pp. 443-473. https://doi.org/10.1016/j.ijforecast.2006.01.001.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, & Kristina Toutanova. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." In *NAACL-HLT 2019*, pp. 4.171-4.186.

DNP. *Colombia's AI Ethics Framework*. Bogotá: Departamento Nacional de Planeación, 2024. https://sisconpes.dnp.gov.co/SisCONPESWeb/ctmp/Borrador_Documento_CONPESIA_comentarios_ciudadanía.pdf.

Easley, David, Marcos M. Lopez de Prado, & Maureen O'Hara. "The Microstructure of the 'Flash Crash': Flow Toxicity, Liquidity Crashes, and the Probability of Informed Trading." *Journal of Portfolio Management* 47, n.º 2, 2011, pp.118-134.

Epstein, Gerald A., ed. Financialization and the World Economy. Cheltenham, UK, Edward Elgar Publishing, 2005.

Eubanks, Virginia. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor. New York:* St. Martin's Press, 2018.

Fama, Eugene F. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance* 25, n.º 2, 1970, pp. 383-417.

Fedesarrollo. *Panorama de la innovación fintech en Colombia*. Bogotá: Fedesarrollo, 2022.

Frino, Alex, Michael Garcia, & Zhaoyun Zhou. "Impact of Algorithmic Trading on Speed of Adjustment to New Information: Evidence from Interest Rate Derivatives." *Journal of Futures Markets* 40, n.º 5, 2020, pp. 749-760. https://doi.org/10.1002/fut.22104.

Frino, Alex, Valerio Mollica, & Robert Webb. "The Impact of Co-Location of Securities Exchanges' and Traders' Computer Servers on Market Liquidity." *Journal of Futures Markets* 34, n.º 1, 2014. https://doi.org/10.1002/fut.21631.

Gambaro, Marco. "Big Data Competition and Market Power." *Market & Competition Law Review* 2, 2018, p. 99.

GAO. *Securities and Exchange Commission: Technology Modernization Needs and Challenges*. Washington, DC: U.S. Government Accountability Office, 2023. https://www.gao.gov/assets/gao-23-106419.pdf.

Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, *et al*. "Generative Adversarial Nets." In Advances in Neural Information Processing Systems, 2672-2680. 2014. https://arxiv.org/pdf/1406.2661.

Gu, Shihao, Bryan Kelly, & Dacheng Xiu. "Empirical Asset Pricing via Machine Learning." *Review of Financial Studies* 33, n.º 5, 2020, pp. 2.223-2.273. https://doi.org/10.1093/rfs/hhaa009.

Hardie, Iain, & Donald Mackenzie. "Assembling an Economic Actor: The Agencement of a Hedge Fund." *The Sociological Review* 55, n.º 1, 2007, pp. 57-80. https://doi.org/10.1111/j.1467-954X.2007.00682.x.

Hasbrouck, Joel, & Gideon Saar. "Low-Latency Trading." *Journal of Financial Markets* 16, n.º 4, 2013, pp. 646-679. https://doi.org/10.1016/j.finmar.2013.05.003.

Hayek, Friedrich A. "The Use of Knowledge in Society." *The American Economic Review* 35, n.º 4, 1945, pp. 519-530.

Hendershott, Joel, & Ryan Riordan. "Algorithmic Trading and the Market for Liquidity." *Journal of Financial and Quantitative Analysis* (JFQA), forthcoming. April 11, 2012. https://ssrn.com/abstract=2001912 or http://dx.doi.org/10.2139/ssrn.2001912.

Hendershott, Terrence, & Ryan Riordan. "Algorithmic Trading and Information." 2009. https://pages.stern.nyu.edu/~bakos/wise/papers/wise2009-3b2_paper.pdf.

Hendershott, Terrence, Charles M. Jones, & Albert J. Menkveld. "Does Algorithmic Trading Improve Liquidity?" *Journal of Finance* 66, n.º 1, 2011, pp. 1–33. https://doi.org/10.1111/j.1540-6261.2010.01624.x.

Hildebrandt, Mireille. *Smart Technologies and the End(s) of Law: Novel Entanglements of Law and Technology.* Cheltenham, UK: Edward Elgar Publishing, 2016.

Johnson, Simon, & James Kwak. *13 Bankers: The Wall Street Takeover and the Next Financial Meltdown.* New York, Pantheon Books, 2010.

Jorion, Philippe. *Value at Risk: The New Benchmark for Managing Financial Risk.* 3rd ed. New York, McGraw-Hill, 2006.

Kingma, Diederik P., & Max Welling. "Auto-Encoding Variational Bayes." *Proceedings of the International Conference on Learning Representations (ICLR), 2014. https://arxiv.org/pdf/1312.6114.*

Kirilenko, Andrei A., Albert S. Kyle, Mehrdad Samadi, & Tugkan Tuzun. "The Flash Crash: High-Frequency Trading in an Electronic Market." *Journal of Finance* 72, n.º 3, 2017, pp. 967-998.

Kolanovic, Marko, & Rajiv Smith. *Big Data and AI Strategies 2019 Alternative Data Handbook. J.P. Morgan Global Quantitative and Derivatives Strategy*, 2019. https://ea-pdf-items.s3-eu-west-1.amazonaws.com/J.P.-Morgan-Alternative-Data-Handbook-2019.pdf.

Krishnamachari, Ramesh. *Big Data and AI Strategies: Machine Learning and Alternative Data Approach to Investing*. J.P. Morgan Global Quantitative and Derivatives Strategy Report, 2017.

Kroll, Joshua A., Joanna Huey, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson, & Harlan Yu. "Accountable Algorithms." University of Pennsylvania Law Review 165, n.º 3, 2016, pp. 633-706.

Langley, Paul, & Andrew Leyshon. "Capitalizing on the Crowd: The Monetary and Financial Ecologies of Crowdfunding." *Environment and Planning A: Economy and Space* 49, n.º 5, 2017, pp. 1.019-1.039. https://doi.org/10.1177/0308518X16687534.

Langley, Paul, & Andrew Leyshon. "Platform Capitalism: The Intermediation and Capitalisation of Digital Economic Circulation." *Finance and Society* 3, n.º 1, 2017, pp. 11-31.

Lanier, Jaron, & E. Glen Weyl. "A Blueprint for a Better Digital Society." *Harvard Business Review*, September 2018. https://hbr.org/2018/09/a-blueprint-for-a-better-digital-society.

Lash, Scott. "Power after Hegemony: Cultural Studies in Mutation?" *Theory, Culture & Society* 24, n.º 3, 2007, pp. 55-78. https://doi.org/10.1177/0263276407075956.

Lee, Ruben. *Running the World's Markets: The Governance of Financial Infrastructure*. Princeton, NJ: Princeton University Press, 2011.

Leinweber, David J. "Stupid Data Miner Tricks: Overfitting in Quantitative Discovery." *The Journal of Investing* 16, n.º 1, 2007, pp. 15-22. https://doi.org/10.3905/joi.2007.681820.

Lewis, Michael. *Flash Boys: A Wall Street Revolt*. New York, W. W. Norton & Company, 2015.

Liu, Jiang, Yingjie Zhang, & Weinan Zhang. "FinRL: Deep Reinforcement Learning Framework to Automate Trading in Quantitative Finance." ACM Transactions on Management Information Systems (TMIS) 13, n.º 1, 2021, pp. 1-36. https://openfin.engineering.columbia.edu/sites/default/files/content/publications/3490354.3494366.pdf.

Lopez de Prado, Marcos. *Advances in Financial Machine Learning*. Hoboken, NJ, Wiley, 2018.

Mackenzie, Donald. *Trading at the Speed of Light: How Ultrafast Algorithms Are Transforming Financial Markets*. Princeton, NJ, Princeton University Press, 2021.

Martínez, S. A. "Retos del sistema financiero colombiano en la Cuarta Revolución Industrial." *Semestre Económico* 24, n.º 56, 2021, pp. 253-270.

Menkveld, Albert J. "High-Frequency Trading and the New Market Makers." *Journal of Financial Markets* 16, n.º 4, 2013, pp. 712-740.

Mirowski, Philip, & Dieter Plehwe. *The Road from Mont Pèlerin: The Making of the Neoliberal Thought Collective*. Cambridge, MA, Harvard University Press, 2009.

Mökander, Johan. "Auditing of AI: Legal, Ethical and Technical Approaches." *Digital Society* 2, n.º 3, 2023, pp. 49.

Ng, Andrew, & Michael Jordan. "On Discriminative vs. Generative Classifiers: A Comparison of Logistic Regression and Naive Bayes." *Advances in Neural Information Processing Systems* 14, 2001.

O'Neil, Cathy. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York, Crown, 2017.

OECD. *Digital Disruption in Banking and its Impact on Competition*. Paris, OECD, 2020. https://www.oecd.org/content/dam/oecd/en/publications/reports/2020/02/digital-disruption-in-banking-and-its-impact-on-competition_330c7176/b8d8fcb1-en.pdf.

Pan, Mingqiang, Danyang Li, Hong Wu, & Ping Lei. "Technological Revolution and Regulatory Innovation: How Governmental Artificial Intelligence Adoption Matters for Financial Regulation Intensity." *International Review of Financial Analysis* 96, 2024, p. 103535.

Pasquale, Frank. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, MA, Harvard University Press, 2015.

Pasquale, Frank. *New Laws of Robotics: Defending Human Expertise in the Age of AI*. New York: Blackstone Publishing, 2021.

Pattanyak, S. K. "Generative AI and Its Role in Shaping the Future of Risk Management in the Banking Industry." 2023.

Patterson, Scott. The Quants: How a New Breed of Math Whizzes Conquered Wall Street and Nearly Destroyed It. New York, Crown Currency, 2011.

Patterson, Scott. *Dark Pools: High-Speed Traders, A.I. Bandits, and the Threat to the Global Financial System*. New York: Crown Publishing, 2012.

Pentland, Alex. Social Physics: How Good Ideas Spread– The Lessons from a New Science. New York, Penguin Press, 2014.

Pistor, Katharina. *The Code of Capital: How the Law Creates Wealth and Inequality*. Princeton, NJ, Princeton University Press, 2019.

Prainsack, Barbara. "The Political Economy of Digital Data: Introduction to the Special Issue." Policy Studies 41, n.º 5, 2020, pp. 439-446. https://ideas.repec.org/a/taf/cposxx/v41y2020i5p439-446.html.

Radford, Alec; Karthik Narasimhan, Tim Salimans, & Ilya Sutskever. "Improving Language Understanding by Generative Pre-Training." *OpenAI*, 2018. https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf.

Rajagopal, Balakrishnan. *International Law from Below: Development, Social Movements and Third World Resistance*. Cambridge, Cambridge University Press, 2009.

Rampini, Adriano, S. Viswanathan, & Guillaume Vuillemey. "Risk Management in Financial Institutions." *The Journal of Finance* 75, n.º 2, 2020, pp. 591-637.

Riles, Annelise. *Collateral Knowledge: Legal Reasoning in the Global Financial Markets*. Chicago: University of Chicago Press, 2011.

Rousseau, Frédéric, Hervé Boco, and Laurent Germain. "High Frequency Trading: Strategic Competition Between Slow and Fast Traders." *Economics Department Working Paper Series* n296-20, Department of Economics, National University of Ireland – Maynooth, 2020. https://ideas.repec.org/p/may/mayecw/n296-20.pdf.html.

Rouvroy, Antoinette, & Thomas Berns. "Gouvernementalité algorithmique et perspectives d'émancipation." *Réseaux* 177, n.º 1, 2013, pp. 163-196.

Snider, Laureen. "Interrogating the Algorithm: Debt, Derivatives and the Social Reconstruction of Stock Market Trading." *Critical Sociology* 40, n.º 5, 2014, pp. 747-761. https://doi.org/10.1177/0896920513504603.

Soros, George. *The Alchemy of Finance: Reading the Mind of the Market*. New York: Simon & Schuster, 1987.

U.S. Commodity Futures Trading Commission (CFTC) and Securities and Exchange Commission, SEC). Findings Regarding the Market Events of May 6, 2010. Washington, DC: U.S. Securities and Exchange Commission, 2010. https://www.sec.gov/news/studies/2010/marketevents-report.pdf.

Veale, Michael, & Lilian Edwards. "Clarity, Surprises, and Further Questions in the Article 29 Working Party Draft Guidance on Automated Decision-Making and Profiling." *Computer Law & Security Review* 34, n.º 2, 2018, pp. 398-404.

Veale, Michael, Reuben Binns, & Lilian Edwards. "Algorithms That Remember: Model Inversion Attacks and Data Protection Law." *Philosophical Transactions of the Royal Society* A 376, n.º 2133, 2018, pp. 1-18. https://doi.org/10.1098/rsta.2018.0083.

Veale, Michael, Max Van Kleek, & Reuben Binns. "Fairness and Accountability Design Needs for Algorithmic Support in High-Stakes Public Sector Decision-Making." Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, 1–14. https://doi.org/10.1145/3173574.3174014.

World Economic Forum. "Colombia's Digital Inclusion Strategy Is Bolstering Financial Inclusion – Especially for Women." August 2024.

Wiese, Matthias, Ralph Knobloch, Ralf Korn, & Thomas Kretschmer. "Quant GANs: Deep Generation of Financial Time Series." *Quantitative Finance* 20, n.º 9, 2020, pp. 1.419-1.440. https://www.semanticscholar.org/paper/Quant-GANs%3A-deep-generation-of-financial-time-Wiese-Knobloch/3cf57cad75d71bffac9fc4589d7b294d90558a13.

Zhang, Yingjie, Stefan Zohren, & Stephen Roberts. "Deep Reinforcement Learning for Trading: A Review." *The Journal of Financial Data Science*, 2020. https://www.oxford-man.ox.ac.uk/wp-content/uploads/2020/06/Deep-Reinforcement-Learning-for-Trading.pdf.

Zuboff, Shoshana. "The Age of Surveillance Capitalism." In *Social Theory Re-Wired*, London, Routledge, 2023, pp. 203-213.