



# The Algorithmic Market Hypothesis: Information Efficiency in the Age of AI

Joseph Simonian, PhD

## Executive Summary

This paper examines the transformation of financial markets as trading shifts from human-driven to algorithmically dominated activity, with particular focus on large language models (LLMs) as a distinct and increasingly influential category of market participant.

We argue that LLMs represent a qualitative shift from traditional algorithmic investing such as high-frequency trading platforms to novel forms of information processing that simultaneously reduce certain information asymmetries while creating new systemic risks. Building on the Grossman-Stiglitz paradox of informationally efficient markets, we demonstrate how the decreasing cost of information processing paradoxically increases market inefficiencies through data mining and the proliferation of false positives.

The paper develops new frameworks for understanding markets where price discovery occurs through the interaction of diverse AI architectures, including high-frequency time-series models and natural language processing systems. It then examines the emergence of company-specific semantic factors as sources of alpha. We conclude that although prices reflect the dominant algorithmic interpretation of information, ultra-fast information processing by AI may render markets less predictable rather than more efficient, as prediction itself becomes endogenous to the system being predicted.

The paper's key arguments and findings are as follows:

### LLMs Reduce, but Do Not Eliminate, Information Asymmetries

LLMs are materially compressing information asymmetries by enabling rapid analysis of complex documents, democratizing access to financial expertise, and lowering the cost of pre-meeting preparation and comparative analysis.

## CONTENTS

[Executive Summary](#) pg. 1 | [Introduction](#) pg. 3 | [Traditional Algorithms vs. LLMs](#) pg. 4 | [LLMs and the Reduction of Information Asymmetries in Finance](#) pg. 6 | [The Grossman–Stiglitz Paradox in the Age of AI](#) pg. 8 | [The Algorithmic Market Hypothesis](#) pg. 11 | [Divergent Dynamics: Time-Series Models vs. NLP Models](#) pg. 13 | [Company-Specific Semantic Factors as Alpha Sources](#) pg. 14 | [Factor Models and Portfolio Optimization in AI-Dominated Markets](#) pg. 17 | [Meta-Models and the Decoupling from Economic Fundamentals](#) pg. 19 | [The Risk of Algorithmic Monoculture and Some Suggested Countermeasures](#) pg. 20 | [Conclusion](#) pg. 24 | [References](#) pg. 26

This reduction, however, is partial and uneven. LLMs effectively process publicly available information that was previously costly or time-consuming to analyze, but they do not eliminate structural advantages predicated on proprietary information, privileged access, or experiential knowledge. The likely outcome is a tiered information landscape, where LLMs compress certain asymmetries while preserving or even potentially amplifying others, as sophisticated participants integrate LLM capabilities into their existing proprietary resources and advantages.

### AI-Driven Trading Creates Endogenous, Self-Reinforcing Factor Dynamics

In AI-dominated markets, traditional factor models break down because the factors themselves become endogenous to AI trading activity. A self-reinforcing feedback loop emerges: (1) AI models use factor returns to predict stock returns, (2) AI trading based on those predictions generates factor returns, (3) those returns feed back into the AI models, and (4) the cycle repeats and amplifies.

When factors are endogenous to AI trading, the conventional decomposition into systematic and idiosyncratic components becomes analytically meaningless. Every piece of information effectively becomes systematic and intrinsically part of the same AI-driven system. This shift necessitates a fundamental reassessment of factor investing strategies and the underlying assumptions of traditional asset pricing models.

### Markets Risk Decoupling from Economic Fundamentals

The evolution of AI-driven markets may eventually lead to a hypothetical future state in which meta-models (systems that predict returns purely from AI behavioral patterns) become entirely detached from underlying economic fundamentals. Such a state would create a paradoxical situation in which markets that are maximally efficient at processing information become entirely disconnected from the underlying economic reality they are intended to reflect. In this scenario, stock prices would be determined by the collective psychology of

“artificial minds” rather than by fundamental business metrics. Investment analysis would in turn shift from studying companies to reverse-engineering the learned representations and decision boundaries of competing AI systems, marking a profound transformation in the practice of financial analysis.

## AI Makes Efficiency Brittle, Not Perfect: Countermeasures Are Needed

Contrary to intuition, AI does not resolve the Grossman–Stiglitz paradox by making markets perfectly efficient. Rather, it makes efficiency radically contingent, brittle, and self-subverting. Accelerating convergence toward similar foundation models is fostering an AI monoculture, eroding the heterogeneity that once provided liquidity and allowed informed players to profit. Practical countermeasures include (1) correlation-penalized ensembles that reward cross-model disagreement in training objectives, (2) reflexivity stress-testing that simulates feedback amplification under synchronized shocks, and (3) adversarial validation protocols, which deliberately probe one’s own models with scenarios designed to expose fragility in consensus-driven patterns.

## Epistemic Diversity Is the Critical Safeguard

Without deliberate epistemic diversity (encompassing model architecture, training data, inductive biases, and loss functions), AI risks turning markets into fragile echo chambers of their own predictions: more reflexive, more volatile, and ultimately less tethered to the economy they claim to represent. The incentive for market participants to pursue differentiated information endures, but differentiation must now operate at the level of how agents form beliefs, what they assume is unlearnable, and how they resist convergence on collective illusions.

## Introduction

The efficient market hypothesis (EMH) (Fama 1970) posits that asset prices fully reflect all available information, implying that systematic outperformance is impossible without the investor either possessing superior information or taking additional risk (through leverage, a higher coefficient of risk aversion, or both). The increasing prevalence of algorithmic trading, however, represents a significant structural change that demands a reevaluation of the EMH and other core financial principles. Because the speed and impartiality of algorithmic execution can dramatically accelerate information assimilation, price discovery windows could potentially shrink to levels previously unimaginable. This acceleration, although theoretically promoting market efficiency, also introduces new sources of systemic risk. The simultaneous response of numerous algorithms to common signals could amplify feedback loops and heighten volatility, particularly during periods of market stress.

Consequently, traditional financial models such as the EMH, factor models, and portfolio optimization techniques may require substantial modification to

account for the dynamics of algorithmically driven markets. This paper explores the implications of markets increasingly dominated by AI and machine learning, focusing on how these technologies alter price discovery, risk assessment, and, ultimately, the efficient allocation of capital.

## Traditional Algorithms vs. LLMs

### Defining the Technological Landscape

A critical distinction exists between established algorithmic trading practices, particularly high-frequency trading (HFT) platforms operating at microsecond intervals, and the application of LLMs in trading. Although both fall under the umbrella of algorithmic approaches, LLMs represent a fundamentally different paradigm. The key difference lies in their ability to process and interpret unstructured data, exhibiting emergent behaviors and introducing novel sources of systemic bias that necessitate separate analytical consideration.

Traditional HFT algorithms are designed to rapidly analyze structured data such as price movements, order book dynamics, and trading volumes, identifying statistical arbitrage opportunities within numerical time series. The boundary between HFT and AI-driven trading has become increasingly porous over the years: To improve signal generation and execution, many HFT firms already deploy machine-learning models such as gradient-boosted trees, reinforcement-learning agents, and, more recently, neural architectures. The difference therefore is not that HFT is exclusively rule based and LLM-driven trading is exclusively learned. Rather, the distinction stems from the fact that LLMs represent a qualitatively different mode of information processing, one centered on unstructured, semantic content rather than the structured, numerical time series that have historically been HFT's domain.

Built upon transformer-based neural networks and trained on extensive textual datasets, LLMs can process unstructured information such as earnings call transcripts, news articles, social media sentiment, and regulatory filings, extracting semantic meaning that lies beyond the reach of conventional numerical algorithms. In contrast to the deterministic nature of HFT algorithms, LLMs can demonstrate *emergent properties*: behaviors not explicitly programmed but arising from the interaction of billions of parameters during the training process.

For example, an LLM trained on general financial text may spontaneously develop the ability to detect management tone shifts or identify early signals of strategic pivots in earnings call language. Although these capabilities were never explicitly encoded, they emerge from the model as it learns deep statistical regularities across vast datasets. This ability stands in sharp contrast to HFT systems, in which every behavior results directly from explicit programming.

In financial markets, such emergent capabilities can translate into unanticipated trading signals as well as into unpredictable failure modes when market conditions shift in ways not represented in the training data. This distinction warrants a reassessment of risk management and analytical frameworks alike.

## The Progression of LLM Automation and Novel Risk Profiles

Incorporating LLMs into financial markets introduces several distinctive characteristics that significantly alter the market landscape. First, LLMs facilitate diverse signal generation. In other words, these models can simultaneously process an unprecedented range of information sources, including corporate disclosures, real-time news sentiment in multiple languages, satellite imagery analysis of supply chains, and anomalous patterns in regulatory filings—all within a single analytical cycle. This breadth of signal contrasts sharply with traditional algorithms, which typically specialize in narrowly defined domains.

Next, LLMs offer enhanced customization and flexibility. Unlike traditional algorithms, which require explicit reprogramming to adapt to new data or market conditions, LLMs can be fine-tuned or prompted to adjust their behavior with minimal technical intervention. This reduced barrier to customization allows for wider adoption among a variety of market participants, from sophisticated hedge funds to individual retail investors using commercial LLM interfaces.

Furthermore, the emergence of agentic capabilities—in which LLMs autonomously plan, execute multi-step workflows, and interact with external data sources and trading systems—further amplifies their utility and reduces the barrier to adoption. These capabilities enable increasingly sophisticated market participation without proportionally increasing human oversight requirements.

LLMs also exhibit systemic biases, however, that are distinct from traditional algorithmic errors. These biases originate from the composition of training data, architectural choices, and the objective functions used during pre-training. For example, LLMs may demonstrate recency bias, create sentiment analysis inaccuracies when they encounter specialized financial terminology, or even generate spurious patterns in environments with low signal-to-noise ratios. These biases often remain latent until market outcomes reveal them.

Unlike a traditional momentum algorithm, for example, for which it is possible to reverse-engineer its behavior from its code, LLM biases are embedded within complex, high-dimensional parameter spaces that are not directly inspectable. Consequently, markets will likely experience a period of heightened risk as participants collectively identify these systemic biases through trial, error, and analysis of anomalous price movements. These unique dimensions of LLM deployment create a dynamic risk environment that differs fundamentally from traditional algorithmic trading, where risks are more readily characterized *ex ante* through established stress-testing and sensitivity analysis techniques. Only through broad recognition and mitigation—addressed using techniques

such as adversarial training, bias correction layers, or ensemble methods—will market participants appropriately price these risks. We discuss these risk mitigation techniques in further detail in a later section.

## LLMs and the Reduction of Information Asymmetries in Finance

### The Democratization of Financial Expertise

Emerging evidence indicates that LLMs are contributing to a notable reduction in the information asymmetries that have historically been a defining feature of financial markets. Recent analysis in *The Economist* illustrates the market efficiency improvements arising as consumers use LLMs to evaluate automotive repair quotes, compare mortgage offerings, and scrutinize service contract terms—all tasks that previously demanded specialized expertise or extensive research beyond the reach of most individuals.<sup>1</sup>

This trend holds even greater significance within institutional finance. The most valuable LLM applications directly address and compress traditionally time-intensive information-processing tasks. Among the most prominent applications are the following:

- *Rapid Document Analysis:* Investors are increasingly using LLMs to review hedge fund prospectuses, which helps to swiftly identify unusual fee structures, lock-up provisions, or redemption clauses that might otherwise be overlooked during time-constrained due diligence. What previously necessitated hours of meticulous review by legal and investment teams can now be accomplished in minutes, with the LLM flagging anomalies for subsequent human verification.
- *Pre-Meeting Preparation:* LLMs can quickly synthesize background information on unfamiliar companies, sectors, or investment vehicles prior to meetings. This ability diminishes the information advantage often held by sell-side analysts or company management, fostering more equitable investor negotiations and more informed decision-making.
- *Comparative Analysis:* Institutional investors are using LLMs to rapidly compare terms across multiple investment opportunities, such as evaluating capital structures, governance provisions, or performance metrics, in order to facilitate more efficient capital allocation without the need for large, dedicated research teams.

Together, these developments suggest that LLMs are fundamentally reshaping the landscape of information access and analysis within the investment process.

---

<sup>1</sup>*Economist*, “The End of the Rip-Off Economy” (27 October 2025). [www.economist.com/finance-and-economics/2025/10/27/the-end-of-the-rip-off-economy](https://www.economist.com/finance-and-economics/2025/10/27/the-end-of-the-rip-off-economy).

## Persistent Information Asymmetries in Financial Markets

Despite the LLM-driven progress in reducing information asymmetries, significant disparities within financial markets persist, and many of these disparities remain beyond the readily addressable scope of current AI technology. Several domains sustain significant informational advantages that LLMs cannot readily penetrate:

- *Proprietary Trading Strategies:* Hedge funds and trading firms often derive an informational edge from proprietary models, private-access alternative data sources, and specialized execution infrastructure. LLMs are inherently limited in their ability to access information deliberately shielded by confidentiality or protected by competitive barriers.
- *Material Non-Public Information:* Corporate insiders, private equity investors with board representation, and certain institutional investors with privileged access to management retain structural information advantages that LLMs cannot overcome. Of course, the legal restrictions surrounding material non-public information preclude not only LLM access but also human exploitation of such advantages. This informational edge is therefore structural rather than actionable.
- *Tacit Knowledge and Judgment:* Through decades of experience, seasoned investors have cultivated pattern recognition abilities and contextual understanding that are not readily codifiable in textual form. Assessing management quality, comprehending industry dynamics, and evaluating strategic positioning often require experiential knowledge that LLMs trained primarily on historical text cannot fully replicate.
- *Network Effects and Relationship Capital:* Access to promising deal flow, co-investment prospects, and early-stage ventures frequently hinges on established relationship networks, and LLMs are unable to effectively penetrate these channels. The social capital accumulated by established investors creates information conduits unavailable to those relying solely on public information processing.
- *Structural Complexity of Private Instruments:* Private markets, particularly private credit, real estate, and infrastructure, present a category of challenge beyond information access. These instruments are often inherently structurally complex: Bespoke loan agreements, layered capital structures, illiquid collateral, and idiosyncratic covenants may require domain expertise simply to identify the relevant risk dimensions before any analytical process can begin. An LLM cannot meaningfully assess a private credit instrument, for example, without first “understanding” what the instrument is—its seniority, security package, covenant protections, and the borrower’s industry-specific operating dynamics. This structural opacity extends to performance data, which are typically infrequent, smoothed, and not comparable across transactions in the way public market data are. For these instruments, the limitations of LLMs are therefore not merely informational

but ontological. The models must first parse the instrument's architecture before they can begin to extract signal, and public training corpora provide far less coverage of such instruments than of publicly traded securities.

Consequently, the reduction of information asymmetries through LLMs is partial and uneven. Although LLMs effectively process publicly available information that was previously costly or time consuming to analyze, they do not eliminate the structural advantages predicated on proprietary information, privileged access, or experiential knowledge. The likely outcome is a tiered information landscape where LLMs compress certain asymmetries while preserving or even potentially amplifying others, as sophisticated participants integrate LLM capabilities with their existing resources and advantages.

Indeed, it is plausible that LLMs could exacerbate a "haves and have-nots" dynamic in finance. For example, well-resourced institutions can deploy fine-tuned, proprietary LLMs on exclusive data streams, while smaller participants rely on general-purpose commercial models on public data. If this dynamic plays out, democratization at the lower end of the information hierarchy may coexist with widening advantages at the top, a form of polarization broadly consistent with what has been observed in other domains where powerful general-purpose technologies have been adopted unevenly. It is essential for investment professionals to develop a nuanced understanding of these evolving information dynamics.

## The Grossman–Stiglitz Paradox in the Age of AI

### The Original Paradox and Its Foundations

Grossman and Stiglitz (1980) articulated a fundamental paradox inherent in market efficiency theory: In order for markets to be informationally efficient, whereby prices reflect all available company information, investors must have no reason to pay for research or analysis because there would be nothing left to discover. This scenario creates an obvious problem: If nobody bothers to gather information, how do prices become efficient in the first place? Of course, they cannot: Prices reflect information only because someone went through the effort of finding and trading on it. Consequently, markets inevitably exist in a state of *equilibrium inefficiency*, a persistent, stable degree of mispricing just large enough to reward those who pay for information with trading profits that cover their research costs, but no larger. In other words, markets are never perfectly efficient, because perfect efficiency would eliminate the very activity that makes them efficient.

This paradox rests on several crucial assumptions: (1) information acquisition incurs a real cost, (2) informed traders generate returns by exploiting information advantages relative to uninformed participants, and (3) equilibrium is maintained when informed traders' profits precisely offset their information costs. The necessary consequence of this equilibrium is that prices remain imperfectly informative. If prices fully reflected all available information,

informed trading would become unprofitable, thereby eliminating the incentive to invest in information acquisition. This inherent tension highlights a critical limitation of the idealized EMH and underscores the persistent role of information asymmetry in shaping market dynamics.

## AI and the Collapse of Information-Processing Costs

The emergence of AI and LLMs fundamentally reshapes the Grossman–Stiglitz framework by dramatically lowering the costs of information processing. LLMs can now accomplish tasks that previously required substantial teams of analysts, such as reviewing thousands of corporate filings, monitoring global news flows, analyzing satellite imagery, and processing social media sentiment. And the costs to do so are marginal, approaching zero per query. This significant cost compression theoretically should drive greater market efficiency, as a broader range of participants gain a greatly enhanced ability to acquire and process information.

The underlying Grossman–Stiglitz logic persists, however, albeit through a modified mechanism. As information-processing costs decline, the focus shifts to data mining, defined here as the exhaustive search for patterns and anomalies within vast datasets. Paradoxically, this increased emphasis on data mining does not necessarily reduce market inefficiencies and can instead exacerbate them. The sheer volume of data and the complexity of algorithmic interactions create new opportunities for misinterpretation and spurious correlations, potentially leading to amplified trading activity based on flawed insights. LLMs lower the costs of information processing, but they do not guarantee improved information quality or enhanced market efficiency.

## Data Mining and the False Positive Problem

Historically, when information processing incurred significant costs, investors concentrated their resources on high-conviction signals grounded in economic theory or demonstrated empirical robustness. The reduction in processing costs brought about by AI and LLMs shifts the bottleneck from processing capacity to signal discovery, where the challenge lies in identifying genuinely predictive relationships within vast datasets. Investors can now readily test millions of potential patterns at minimal cost, resulting in an explosion of data-mining activity.

The statistical consequence of this data-mining renaissance is a dramatic proliferation of *false positives*, patterns that appear predictive within a specific sample but reflect data-mining artifacts rather than genuine economic relationships. Consider an investor using an LLM to analyze 10,000 potential signals across 3,000 companies over a 20-year period of textual data. Even with rigorous statistical controls, the multiple-testing problem becomes acutely pronounced (Bonferroni 1936; Benjamini and Hochberg 1995). At a 5% significance level, approximately 500 factors would appear statistically significant purely by chance.

These false positives may subsequently propagate through the market as traders allocate capital on the basis of spurious signals. When multiple market participants independently “discover” the same false pattern through similar data-mining techniques, their collective trading generates temporary price distortions. These distortions are inefficiencies because they are mispricings that do not reflect fundamental information, and ironically they arise precisely because information processing has become so inexpensive. Accordingly, robust statistical validation and a skeptical assessment of algorithmic signals become more important than ever.

## The New Equilibrium: Efficiency Through Inefficiency

In the modified Grossman–Stiglitz paradox, rather than costly information acquisition creating a divergence between price and fundamental value, the cost of validating discovered signals and distinguishing true economic relationships from statistical artifacts now generates this same divergence. The equilibrium level of market efficiency is consequently determined by the trade-off between two critical factors:

- (1) the still-substantial cost of validating whether those discovered patterns represent genuine economic relationships or mere statistical anomalies, and
- (2) the near-zero cost of generating potential trading signals through AI-driven data mining.

With respect to the first factor, perfect market efficiency is unattainable because it would necessitate an infinite allocation of resources dedicated to discerning true signals from false positives. Investors who invest in robust signal validation procedures generate returns by avoiding erroneous patterns that other participants use for trading. The equilibrium persists where the marginal cost of additional validation equals the marginal benefit derived from improved signal quality.

A critical caveat to this reframing, however, concerns whether signal validation is achievable in the first place. The Grossman–Stiglitz equilibrium requires that validation be costly but tractable: a feasible activity whose marginal costs can be weighed against marginal benefits. For traditional algorithmic signals, this condition broadly holds: A momentum strategy can be tested out-of-sample, its statistical properties examined, and its economic rationale scrutinized. LLM-based signals, however, present a more fundamental challenge. Because LLMs are continuously updated by their developers, the model generating a signal today may differ materially from the model that will be used six months hence. By the time sufficient data have accumulated to validate a signal derived from a particular version of an LLM, the underlying model may have already shifted in ways that invalidate the test entirely.

This is not merely a cost problem but a deeper issue of *epistemic tractability*: Validation requires a stable object to validate, and a continuously evolving LLM cannot provide one. If validation is closer to being intractable than merely expensive, the equilibrium level of market efficiency described earlier may not exist in a well-defined sense. The framework presented here should therefore be

understood as an approximation that holds only to the extent that participants can impose sufficient stability on their models—through versioning, frozen checkpoints, or architecture constraints—to make meaningful out-of-sample testing possible. This implies a maximum validatable horizon, limiting testable strategies to those whose evaluation window fits inside the model release cadence.

With respect to the second factor, this framework predicts that as AI diminishes information-processing costs, market inefficiencies may actually increase in certain domains, particularly those involving complex, high-dimensional data where the space of potential patterns vastly exceeds the space of economically meaningful relationships. The ease of signal discovery creates a form of “tragedy of the commons” in signal space, whereby all participants can cheaply mine data but the collective outcome is an increase in noise rather than an enhancement of efficiency. This dynamic is, however, currently moderated by a significant bottleneck that the foregoing analysis does not fully capture: implementation infrastructure. In other words, the ability to generate signals cheaply has not eliminated the cost of acting on them at scale.

High-frequency trading remains concentrated among a handful of firms, not because information processing is scarce but because the execution infrastructure required to convert signals into returns (e.g., low-latency networks, co-location arrangements, and sophisticated risk management systems) is exceptionally expensive to build and maintain. The “tragedy of the commons” in signal space thus may still be some distance away for strategies that require high-frequency execution, even if it is already materializing for lower-frequency, text-based approaches with relatively low implementation barriers. Rigorous validation methodologies and robust statistical controls are accordingly indispensable for investment firms operating in this environment.

## The Algorithmic Market Hypothesis

A market increasingly dominated by AI-driven trading represents the most profound challenge yet to the EMH. Paradoxically, we may observe both extreme efficiency and complete breakdown. Markets could achieve hyper-efficiency in incorporating traditional quantitative and fundamental information, yet the uniformity of AI approaches could simultaneously create novel forms of systemic mispricing. In such a scenario, the definition of “information” expands beyond traditional datasets to encompass alternative, unstructured sources such as satellite imagery, social media sentiment, and supply chain flows, processed at scale and speed unattainable by human analysts. This development fundamentally shifts the efficiency frontier but also introduces new systemic risks. The convergence of algorithmic approaches creates the potential for shared blind spots and the propagation of correlated mispricings, stemming from both shared data sources and inherent model biases.

Consequently, the very concept of “market efficiency” warrants re-evaluation. The central question becomes whether efficiency retains meaning when price discovery is driven by algorithmic consensus rather than the diverse perspectives and judgments of human investors. The potential systemic risks

inherent in a market where algorithmic homogeneity outweigh the benefits of rapid information assimilation thus deserve careful scrutiny.

The EMH condition  $P_t = E[V|\Omega_t]$ —where the current price,  $P_t$ , equals the expected fundamental value,  $V$ , of the asset given all available information,  $\Omega_t$ —breaks down, because the information set,  $\Omega_t$ , becomes model-dependent:

For natural language processing (NLP) models,

$$\Omega_t^{NLP} = \{price\ history, all\ text\ data\}. \quad (1)$$

For time-series models,

$$\Omega_t^{TS} = \{detailed\ price/volume\ microstructure\}. \quad (2)$$

The EMH violation occurs when

$$|E[V | \Omega_t^{NLP}] - E[V | \Omega_t^{TS}]| > transaction\ costs. \quad (3)$$

Equation 3 captures the idea that although an NLP model and a time-series model are both looking at the same stock, each is working from a fundamentally different picture of it. The NLP model forms its view of the stock's fair value by reading and interpreting textual information such as earnings call transcripts, news articles, and regulatory filings. In contrast, the time-series model forms its view purely from the statistical structure of price and volume data, such as momentum patterns, order book dynamics, and volatility regimes.

Because these two information sets are genuinely different, the two models will often arrive at meaningfully different estimates of what the stock is worth. Equation 3 states that whenever this valuation gap is large enough to exceed transaction costs, a profitable trade exists—and so does evidence that the market is not fully efficient. This state of affairs reflects a substantive and novel form of market segmentation in which the same publicly available data are processed through architecturally distinct lenses, generating persistent differences in inferred value. Unlike ordinary belief dispersion, these differences may not be arbitrated away through price inference alone, because a time-series model cannot reconstruct a semantic signal from the price path.

Different AI architectures process fundamentally different information sets, creating systemic differences in valuation even when all participants have access to the same raw data. Thus, in an algorithmically dominated market, we will be confronted with a new form of market segmentation based not on information access but on information-processing architecture. If this happens, prices may very well conform to what we call the algorithmic market hypothesis (AMH).

AMH: Current prices reflect the dominant algorithmic interpretation of information.

Such segmentation can persist even as institutional investors converge on similar foundation models, since NLP and time-series models will continue to process different information sets.

## Divergent Dynamics: Time-Series Models vs. NLP Models

### Time-Series Models: Ultra-Fast Mean Reversion Markets

Time-series models demonstrate a remarkable capacity for detecting statistical patterns, autocorrelations, and regime changes within price data. These models can use sophisticated econometrics or machine-learning architectures: for example, transformer networks processing order book sequences, long short-term memory (LSTM) networks identifying regime shifts, and attention mechanisms weighting recent versus historical price movements on the basis of market microstructure. In a market dominated by sophisticated time-series AI, several significant phenomena may emerge:

*EMH Transformation:* Markets might achieve microsecond-level efficiency with respect to price-based information, but this achievement could paradoxically create new inefficiencies. If every model detects the same mean reversion pattern, the collective rush to exploit it can generate violent oscillations around fair value. The EMH would likely hold at longer horizons but break down entirely at shorter ones.

*Factor Model Breakdown:* Traditional factors, which rely on persistent price patterns, would face intensified arbitrage pressure from time-series models. Momentum factors might more rapidly compress toward mean reversion at shorter horizons. It is worth noting, however, that HFT firms have long incorporated text and news signals alongside price data, and many factors have proved resilient because they are grounded in persistent behavioral or structural features of markets rather than purely statistical regularities. The more plausible scenario is a compression of factor half-lives at high frequencies. Lower-frequency premiums linked to fundamental risk exposures or behavioral biases prove more durable.

### NLP Models: Fundamental-Driven Efficiency with Novel Vulnerabilities

NLP models process news, earnings calls, social media, and unstructured data to predict price movements on the basis of information content rather than price patterns. In doing so, they create fundamentally different market dynamics, as follows:

*Efficient Market Hypothesis Transformation:* Markets have become exceptionally efficient at incorporating textual information, especially in processing earnings announcements, news releases, and central bank communications. However, this heightened efficiency introduces vulnerabilities to manipulation through coordinated narrative campaigns or adversarial text generation specifically designed to deceive NLP systems.

*Emergence of New Factors:* Traditional fundamental factors might actually strengthen as NLP models more accurately capture the true information content of financial statements, analyst reports, and management communications. Simultaneously, entirely new semantic factors could emerge. A semantic factor is a return-predictive variable derived not from traditional financial metrics such as earnings or book value but from measurable linguistic or stylistic patterns in corporate communications. For example, a company whose management consistently uses precise, quantitative language in earnings calls may earn a “communication clarity” premium as NLP models learn to associate that pattern with earnings predictability.

Unlike conventional factors that are grounded in economic theory, semantic factors emerge empirically from the way NLP models process text: by observing certain linguistic patterns in communications from important market actors. Semantic factors may systematically outperform conventional fundamental factors, independent of conventional fundamental metrics. Investment practitioners must therefore attend to not only the quantitative dimensions of financial data but also the qualitative nuances embedded in textual information. The following section further explores semantic factors.

## Market Structure Bifurcation

The potential emergence of distinct market layers—a high-frequency stratum dominated by time-series models trading purely on price dynamics and a lower-frequency layer where NLP models compete on information processing—would generate entirely new forms of cross-frequency arbitrage opportunities. Time-series models would likely drive portfolios toward extremely high-turnover, market-neutral strategies concentrated on statistical arbitrage. Conversely, NLP-driven portfolios would likely exhibit lower turnover, focusing on assets where informational advantages derived from text analysis provide the greatest edge. The most significant question, however, revolves around the dynamic interaction between these layers: What occurs when time-series models begin detecting the trading patterns of NLP models, and vice versa? This creates feedback loops where models learn to identify and exploit each other’s behavioral signatures, effectively trading not the underlying assets but the patterns of competing AI systems themselves. The implications of such *meta-strategies* for market stability will be discussed in further detail in the section titled “Meta-Models and the Decoupling from Economic Fundamentals.”

## Company-Specific Semantic Factors as Alpha Sources

NLP models facilitate the discovery of idiosyncratic semantic factors, which are linguistic patterns that predict individual company performance. Semantic factors represent an entirely new category of risk factors that do not align with traditional factor frameworks.

## Management Communication Style Factors

The technical implementation of semantic factors involves embedding models, such as BERT or GPT variants, to analyze management discourse patterns. For example, a large, closely followed mega-cap company might exhibit a “precision language factor” as its consistent, measured communication style, which could become a predictive feature independent of the content itself. This example highlights the potential for NLP to uncover subtle but meaningful signals from unstructured textual data, offering investors a new dimension for generating alpha. The mathematical formulation in Equation 4 represents management communication at time  $t$  as a BERT (Devlin et al. 2019) embedding

$$E_{i,t} = \text{BERT}_{\theta(\text{text}_{i,t})} \quad (4)$$

where  $E$  denotes the BERT embedding vector (a high-dimensional numerical representation of the text),  $\theta$  represents the learned parameters (weights) of the BERT model,  $i$  indexes the company,  $t$  indexes the time period, and  $\text{text}$  refers to the raw textual input (e.g., the earnings call transcript) for company  $i$  at time  $t$ . The predictive semantic factor becomes

$$SF_{i,t} = \text{cosine}_{\text{similarity}(E_{i,t}, \hat{E}_t)} \quad (5)$$

where  $\hat{E}_t$  represents the optimal communication pattern at time  $t$  relative to a *linguistic benchmark*, a reference vector that captures the communication style associated with optimal return-predictive behavior at any given point in time. It is, in effect, the “gold standard” language pattern that NLP models have learned to reward.

This benchmark evolves through exponential smoothing with a learning adjustment:

$$\hat{E}_{t+1} = \alpha \hat{E}_t + (1 - \alpha) E_{i,t} + \delta_t \quad (6)$$

This updating process creates a dynamic system in which the optimal communication pattern shifts as models update their parameters on the basis of observed relationships between language and returns. In Equation 6,  $\alpha$  is the smoothing parameter controlling how quickly the benchmark adapts to new observations (between 0 and 1), and  $\delta$  is a time-varying learning adjustment term that captures structural shifts in the relationship between communication style and returns, such as a sudden change in market regime that alters which linguistic patterns are rewarded.

## Narrative Consistency Factors

Companies maintaining semantic consistency across investor-relations documents might trade at premiums. Vector similarity measures between quarterly reports could become risk factors. For example, a biotechnology company that suddenly shifts terminology around drug trials might trigger

algorithmic selling before any explicit negative news emerges. The deviation from established narrative patterns signals potential underlying problems.

## Sentiment Momentum vs. Reversal

NLP models might discover that certain companies exhibit momentum in sentiment-driven moves while others mean-revert. A major tech company might have a “sentiment persistence factor,” in which positive NLP sentiment continues to drive prices, whereas a utility company might exhibit “sentiment fade,” in which initial NLP-driven moves reverse. These company-specific patterns represent exploitable anomalies that arise from the interaction between company characteristics and NLP model architectures.

These sentiment momentum and reversal patterns are, moreover, closely related to the time-series versus NLP model dynamics discussed in the previous section. When time-series models detect the trading footprints of NLP-driven sentiment moves, their responses can either amplify persistence (reinforcing momentum) or trigger rapid mean reversion, depending on the speed and architecture of the competing systems. The phenomena described here are therefore best understood not as entirely independent effects but as company-level manifestations of the broader cross-model interaction dynamic.

## The False Positive Problem in Semantic Factor Discovery

The ease with which LLMs can extract potential semantic factors dramatically increases the risk of discovering false positives. This issue connects directly to the data-mining concerns raised in the Grossman–Stiglitz discussion. An analyst can now test thousands of linguistic features (e.g., word choice patterns, syntactic structures, sentiment trajectories, or topic modeling clusters) across thousands of companies with minimal computational cost. But doing so creates a severe multiple-testing problem, as noted earlier.

When multiple market participants independently discover overlapping sets of these false factors, their collective trading creates real price movements that temporarily validate the false signals. The resulting feedback loops amplify market inefficiencies, and semantic alpha decay follows:

$$\alpha_{\text{semantic},t} = \alpha_0 \cdot e^{-\lambda t} (1 - \text{adoption rate}_t), \quad (7)$$

$$\text{where } \text{adoption rate}_t = \tanh\left(\frac{\text{Cumulative AI Investments}_t}{\text{Market Cap}_t}\right).$$

In Equation 7,  $\lambda$  is a fixed, strategy-specific decay rate parameter governing the speed at which a semantic factor loses predictive power as the broader market learns the same regularities. Because  $\lambda$  is treated as a constant characteristic of a given strategy rather than a quantity that varies over time, it carries no time subscript. This choice is intentional and distinguishes it from the time-indexed

adoption rate term. This mathematical formulation captures how semantic factors lose predictive power as more capital deploys similar NLP strategies, creating an arms race in which alpha decays faster as adoption increases.

**Exhibit 1** shows a simulation of semantic alpha decay.

Panel A of Exhibit 1 shows that even modest increases in either the decay rate,  $\lambda$ , or the speed of adoption compress the useful life of a semantic strategy from several years to less than two. Panel B decomposes the two multiplicative terms of Equation 7 for the moderate scenario: The exponential term  $e^{-\lambda t}$  captures intrinsic signal erosion as the broader market learns the same regularities, while the adoption-penalty term  $(1 - \tanh(\cdot))$  captures the additional crowding effect as more capital pursues similar NLP strategies. Because these forces compound rather than merely add, semantic alpha decays substantially faster than either term alone would suggest. This result has direct implications for how frequently practitioners must refresh their feature sets.

## Factor Models and Portfolio Optimization in AI-Dominated Markets


### Factor Model Transformation

Traditional factor models assume stable  $\beta$  coefficients in the relationship:


$$R_{i,t} = \alpha_i + \beta_i F_t + \varepsilon_{i,t}. \quad (8)$$

### Exhibit 1. Simulation of Semantic Alpha Decay

**Panel A. Semantic Alpha Decay Under Varying Adoption and Decay Regimes**



**Panel B. Decomposition of Decay Terms**



Notes: Alpha is normalized to its initial value  $\alpha_0$ . Panel B depicts a moderate scenario with decay rate  $\lambda = 0.30$  and adoption speed = 0.50.

With AI learning, however,  $\beta$  becomes path-dependent through an adaptive updating process:

$$\beta_{i,t+1} = \beta_{i,t} + \eta(R_{i,t} - \beta_{i,t}F_t - \alpha_i), \quad (9)$$

where  $\eta$  is the learning rate. Crucially,  $F_t$  (factor returns such as market, value, momentum, and size) itself becomes endogenous to AI trading activity:

$$F_t = g(Al_t, Trad_t, Int_t), \quad (10)$$

where  $Al_t$  represents AI-driven trades,  $Trad_t$  represents traditional trades, and  $Int_t$  captures cross-model interactions.

Traditional factor models typically treat these factors as exogenous, because external forces are assumed to drive stock returns. In AI-dominated markets, however, the factors themselves are generated by the AI trading activity.

Consider a traditional interpretation of the momentum factor: "The momentum factor had a +2% return today because momentum stocks fundamentally outperformed." In contrast, within an AI-driven market, the explanation becomes "The momentum factor had a +2% return today because 47 AI systems detected momentum signals and purchased momentum stocks. These purchases caused prices to rise, creating the +2% momentum factor return, which other AI systems subsequently detected and incorporated into their models, prompting further momentum buying tomorrow." The AI-driven case establishes a feedback loop wherein (1) AI models use factor returns to predict stock returns, (2) AI trading based on those predictions generates the factor returns, (3) the factor returns feed back into the AI models, and (4) the cycle repeats and amplifies.

Traditional factor models operate under the assumption that factors are independent, persistent risk premiums inherent in the economy. When factors are endogenous to AI trading, however, the conventional decomposition into systematic and idiosyncratic components becomes analytically meaningless because everything effectively becomes systematic, intrinsically part of the same AI-driven system. In this case, the underlying assumptions of traditional factor investing and asset pricing models will be in need of fundamental reassessment.

Two important qualifications to this framework are warranted. First, the path-dependent  $\beta$  in Equation 9 is a fundamentally different object from the stable  $\beta$  in Equation 8. It is not a parameter that can be estimated with standard statistical techniques and interpreted through conventional significance tests; rather, it is a latent process to be tracked over time. Assigning statistical significance to  $\beta_i$  at any given point is not meaningful in the classical sense, and practitioners should treat the adaptive updating rule as a descriptive model of evolving exposures rather than an inferential framework.

Second, the mutual endogeneity of  $\beta$  and  $F_t$  in Equations 9 and 10 raises a further concern: Because both quantities are jointly determined by AI trading

activity, standard methods for identifying factor premiums (which presuppose that factors are exogenous) break down. Disentangling genuine risk exposures from feedback-induced co-movement in such a system presents a significant methodological challenge that remains an open area for future research.

## Portfolio Optimization as Algorithmic Arms Race

Classical mean-variance optimization relies on the assumption of stable relationships and known parameters. In an AI-dominated market, however, these assumptions are fundamentally challenged. Portfolio construction must evolve to explicitly account for the following:

- the correlation between your AI model's behavior and that of other AI models operating within the market,
- regime shifts driven by alterations in algorithmic trading behavior,
- the risk that widely used AI strategies become suddenly crowded trades, leading to rapid reversals, and
- dynamic hedging against the predicted actions of other AI systems, effectively anticipating and mitigating their potential impact on portfolio performance.

Game theory provides a valuable framework for understanding these intricate dynamics. Consider  $N$  AI agents, each using investment strategy  $s_i$ . The Nash (1950) equilibrium condition, where  $s^*$  denotes an equilibrium strategy (i.e., the optimal strategy given that all other agents are also playing their optimal strategies), is as follows:

$$\pi_{i(s_i^*, s_{-i}^*)} \geq \pi_{i(s_i, s_{-i}^*)} \text{ for all } s_i \text{ and all } i, \quad (11)$$

where  $s_i$  denotes agent  $i$ 's equilibrium strategy (here representing any feasible deviation from equilibrium  $s_i^*$ ), and  $s_{-i}$  denotes the equilibrium strategies of all other agents. The profit function,  $\pi_i$ , depends on prediction accuracy, position size, and strategic correlation:

$$\pi_i = \alpha_i \cdot \text{accuracy}_{i(s_i, s_{-i})} - \beta_i \cdot |\text{position}_i| - \gamma_i \cdot \text{correlation}_{i(s_i, s_{-i})}. \quad (12)$$

The correlation penalty,  $\gamma_i$ , creates a mathematical tendency toward strategy differentiation. That tendency is bounded by the cost of differentiation, which is asymmetric across participants and therefore implies stratification rather than uniform divergence.

## Meta-Models and the Decoupling from Economic Fundamentals

The evolution of AI-driven markets may eventually lead to the emergence of meta-models, which predict returns purely on the basis of AI behavioral patterns, effectively detached from underlying economic fundamentals.

This development represents a radical departure from traditional finance theory, operating in a pure pattern space where the only inputs are the behavioral signatures of other AI systems. This progression can be articulated as follows:

Level 0: Stock prices reflect economic reality.

Level 1: Stock prices reflect patterns in economic data.

Level 2: Stock prices reflect patterns in the predictions of all market participants' AI models.

Level 3: Stock prices reflect patterns in how AI models detect patterns in other AI models.

At Level 3, complete economic decoupling occurs. For example, a pharmaceutical company's stock might move not in response to drug trial results, but because Meta-Model X detected that NLP Model Y typically exhibits overconfidence when processing biotech earnings calls, while Time-Series Model Z consistently front-runs NLP Model Y's trades by 3.5 minutes.

Meta-models effectively trade the traders rather than the assets themselves. They develop signatures for each AI system, characterizing patterns such as "Model A consistently overshoots momentum signals by 15%," "Model B's confidence declines predictably after three consecutive losses," and "Models C and D exhibit a 73% correlation in their error patterns." This creates a paradoxical situation where markets that are maximally efficient at processing information become entirely disconnected from the underlying economic reality they are intended to reflect. Stock prices are determined by the collective psychology of "artificial minds" rather than fundamental business metrics.

Investment analysis would in turn shift from studying companies to reverse-engineering the learned representations and decision boundaries of competing AI systems, marking a profound transformation in the practice of financial analysis. Such decoupling is bounded rather than absorbing, since the return to trading against it rises as it widens and fundamentals are eventually realized through cash flows, acquisition, or liquidation.

## The Risk of Algorithmic Monoculture and Some Suggested Countermeasures

In the AI era, we face the possibility that the Grossman-Stiglitz paradox evolves into a more insidious, self-reinforcing dynamic. With information-processing costs plummeting, the scarce resource shifts from raw computation to differentiated epistemology. The differentiations may include the creation of unique model architectures, proprietary data streams, contrarian inductive biases, or deliberately non-convergent training regimes.

With that said, supervisory and industry evidence points to accelerating convergence, with leading institutional investors increasingly relying on similar foundation models (e.g., transformer variants fine-tuned on overlapping

financial corpora and alternative data feeds), fostering an AI monoculture (Kleinberg and Raghavan 2021). This AI monoculture explicitly erodes the heterogeneity that once resolved the Grossman–Stiglitz paradox, namely the presence of idiosyncratic “noise traders” (Black 1986; De Long, Shleifer, Summers, and Waldmann 1990) whose imperfect, diverse beliefs provided liquidity and allowed informed players to profit.


Examining the situation from the perspective of systems theory, we see a complex interplay of balancing and reinforcing forces (Meadows 2008). The collapsing cost of information processing (a reinforcing force) drives algorithmic homogeneity (also reinforcing), which diminishes liquidity and amplifies shocks (balancing forces). Empirical warnings from regulators (e.g., Financial Stability Board, International Organization of Securities Commissions, European Central Bank) and researchers (Financial Stability Board 2024; International Organization of Securities Commissions 2025; European Central Bank 2024) highlight model herding and correlation risks, where synchronized AI responses exacerbate shocks, evaporate liquidity during stress events, and precipitate flash-crash-like occurrences (an algorithmic “Minsky moment” of sorts).<sup>2</sup> Recent episodes are consistent with this brittleness: the tariff announcements in early 2025, which triggered volatility spikes; AI-related volatility pulses in November 2025, which contributed to sharper drawdowns; and concerns over AI-driven “monoculture” echoing the 2010 Flash Crash.

**Exhibit 2** illustrates this dynamic. To help interpret this exhibit, a brief glossary of systems theory terms follows. *Stocks* are accumulations that build up or deplete over time, like water in a reservoir. They represent the state of the system at any moment. *Flows* are the rates at which stocks change, controlling how quickly they fill or drain. *Variables* are factors that influence the system but do not accumulate themselves; rather, they affect flows and other variables. *Reinforcing loops* create exponential growth or decline by amplifying change in the same direction (the more you have, the more you get), leading to acceleration. The behavior of a reinforcing loop may be compared to the phenomenon of resonance in physics: Just as a bridge can collapse when periodic forces align with its natural frequency and energy accumulates faster than it dissipates, a financial system dominated by algorithmically homogeneous agents can enter a destructive resonance when synchronized trading signals reinforce each other across many market participants, amplifying rather than damping systemic stress. In contrast, *balancing loops* work to stabilize the system by counteracting change, pushing toward equilibrium or a goal. When some element of a system increases, a balancing loop works to decrease it, and vice versa.

Together, these elements create the dynamic behavior patterns we see in complex systems such as markets, ecosystems, and organizational processes. In Exhibit 2, the primary “stocks” are algorithmic homogeneity (the accumulated degree to which market participants have converged on similar AI architectures and strategies) and market fragility (the accumulated vulnerability of the system to synchronized shocks). The primary “flows” are the rate of homogeneity

<sup>2</sup> John Cassidy, “The Minsky Moment,” *New Yorker* (4 February 2008).

## Exhibit 2. Hypothetical AI-Driven Market Dynamics



growth (driven by the reinforcing cost-reduction loop) and the rate of fragility accumulation (driven by growing homogeneity). Compute cost reduction is the exogenous variable that initiates the reinforcing loop, and validation cost is the key balancing variable that moderates the pace of convergence by preserving a residual cost to distinguishing genuine signals from noise.

Panel A of Exhibit 2 shows that compute cost (the exogenous driver) declines along an approximately exponential path, feeding a reinforcing loop (shown in Panel B) that drives algorithmic homogeneity toward saturation. Panel C illustrates how market fragility accumulates in proportion to homogeneity and is periodically interrupted by shock events (shown as vertical markers). Notably, because the reinforcing loop shown in Panel B is strong relative to any single shock, the system tends to re-lock into high fragility after each reset, illustrating our central claim that AI-driven monoculture is self-reinforcing rather than

self-correcting. Finally, Panel D shows how validation cost, the key balancing force, declines only slowly, preserving the partial inefficiency equilibrium described earlier (in the section titled “The New Equilibrium: Efficiency Through Inefficiency”). The rapid-adoption scenario compresses the same dynamics into a shorter horizon, accelerating fragility onset.

As a consequence of this dynamic, the classic Grossman–Stiglitz incentive reemerges: not for gathering public data (now a trivial undertaking) but for engineering strategic divergence through proprietary synthetic data generation, adversarial training against consensus patterns, or loss functions that penalize correlation with peer models—all of which are tools to preserve exploitable edges. Ultra-low-cost search supercharges *underdetermination*, the idea that any dataset will be compatible with multiple theoretical explanations (Quine 1951).<sup>3</sup> High-dimensional feature spaces yield combinatorial explosions of candidate factors, and modern machine-learning universal approximators fit noise as readily as signal, producing a “factor zoo” wherein a significant proportion of published anomalies fail rigorous out-of-sample or multiple-testing scrutiny (Harvey and Liu 2019). This situation manifests as “meta-overfitting”: Models overfit not only historical noise but also the transient behaviors of other models, creating illusory alphas that vanish when the herd adjusts. Although validation can be costly and human-resource intensive, such rigorous checks provide a crucial systemic counterweight. Equilibrium thus settles at partial inefficiency, where enough spurious patterns survive to reward differentiation, but not enough patterns survive to achieve hyper-efficiency.<sup>4</sup>

To counter the foregoing developments, several practical remedies include (but are not limited to) the following:

- correlation-penalized ensembles that reward cross-model disagreement in training objectives,
- reflexivity stress-testing, simulating feedback amplification under synchronized shocks, and

<sup>3</sup>For example, imagine you have data showing a correlation between ice cream sales and crime rates. You could propose a theory that ice cream consumption causes crime (highly unlikely). Or you could propose a theory that both ice cream sales and crime rates are caused by a third factor, such as warmer weather. Although both explanations are consistent with the data, one is far more plausible than the other. The data underdetermine which explanation is correct.

<sup>4</sup>Another risk is that markets could stratify orthogonally when ultra-high-frequency layers converge toward near-zero spreads as near-identical models arbitrage each other into oblivion (approaching a Grossman–Stiglitz collapse). Lower-frequency layers host meta-adversarial trading, where specialized AIs exploit systemic blind spots in dominant architectures (e.g., biases favoring short-horizon momentum over long-term reversion). Emerging regimes feature AI-versus-AI behavioral arbitrage, hunting for evidence of consensus fragility rather than underlying fundamentals. Alpha thus increasingly derives from epistemic arbitrage, where firms deliberately build divergent assumptions into their models. This practice of engineered divergence preserves the Grossman–Stiglitz logic, as someone must still bear a cost to differentiate, but shifts that cost from gathering data to designing how a model learns. If this shift becomes widespread, markets enter a phase where prices no longer aim to reflect the real economy at all. Instead, prices simply reflect the shifting balance of models reacting to one another. In essence, we risk becoming financial “brains in vats,” akin to Putnam’s (1981) well-known thought experiment in which an evil scientist removes a brain from its body, places it in a vat of nutrients, and connects it via electrodes to a supercomputer that perfectly simulates all sensory inputs, creating an indistinguishable illusion of normal embodied experience. Similarly, in hyper-financialized markets dominated by self-referential models, asset prices may increasingly reflect simulated realities and feedback loops rather than grounding signals from the underlying real economy.

- the use of adversarial validation protocols (a.k.a. “red-teaming”) to probe one’s own models and strategies by stress-testing them against scenarios specifically designed to expose their weaknesses.

In the present context, these scenarios would reflect market conditions under which consensus-driven AI signals are most likely to fail or amplify instability.

The dynamics described here have a striking antecedent in Keynes’s (1936) famous characterization of financial markets, in which he observed that investors reach “the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be” (p. 156). In Keynes’s framing, successful investing requires not only assessing fundamental value but also anticipating what *others* will anticipate. In AI-dominated markets, this logic extends further still. Herding toward the same signals as one’s peers is precisely the behavior of agents “anticipating the anticipations” of other AI systems, a higher-order guessing game now operating at machine speed and scale. The epistemic countermeasures described earlier in this section are, in this sense, strategies for escaping Keynes’s higher-order guessing game before it converges on a collectively self-subverting equilibrium.

Ultimately, AI does not resolve the Grossman–Stiglitz paradox by making markets perfectly efficient. Rather, it makes efficiency radically contingent, brittle, and self-subverting. The incentive to pursue differentiated information endures, but now differentiation must operate at the level of how agents form beliefs and which patterns they treat as genuinely predictive versus noise, how they avoid herding toward the same signals as their peers, and how they maintain strategic distinctiveness as dominant model architectures converge.

## Conclusion

The transition to AI-dominated financial markets represents a foundational transformation that demands extensive theoretical reconstruction across all pillars of modern finance. LLMs introduce a form of market participation that is qualitatively distinct from established algorithmic trading practices. They process unstructured information at scale while simultaneously exhibiting systemic biases that create novel risk profiles requiring iterative discovery and mitigation, an approach that fundamentally alters the Grossman–Stiglitz framework.

Although LLMs are demonstrably reducing information asymmetries in the investment landscape by enabling rapid analysis of complex datasets and also are democratizing certain facets of financial expertise, they do not eliminate the structural advantages that stem from proprietary information, privileged access, or experiential knowledge. Indeed, they may inadvertently amplify such asymmetries as sophisticated market participants integrate LLM capabilities with their existing resources.

The resulting information landscape is inherently heterogeneous. Certain barriers to information are lowered while others persist, and new challenges emerge related to meta-overfitting, endogenous reflexivity, and the proliferation of spurious signals. Navigating this increasingly algorithmically driven environment will require, above all, a continually evolving and empirically grounded understanding of market dynamics and more sophisticated validation frameworks.

## References

Benjamini, Yoav, and Yosef Hochberg. 1995. "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing." *Journal of the Royal Statistical Society: Series B (Methodological)* 57 (1): 289–300. doi:10.1111/j.2517-6161.1995.tb02031.x.

Black, Fischer. 1986. "Noise." *Journal of Finance* 41 (3): 529–543. doi:10.1111/j.1540-6261.1986.tb04513.x.

Bonferroni, Carlo E. 1936. "Teoria Statistica delle Classi e Calcolo delle Probabilità." *Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commerciali di Firenze* 8: 3–62.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. 1990. "Noise Trader Risk in Financial Markets." *Journal of Political Economy* 98 (4): 703–738. doi:10.1086/261703.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. "BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding." In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*: 4171–86. doi:10.48550/arXiv.1810.04805.

European Central Bank. 2024. "The Rise of Artificial Intelligence: Benefits and Risks for Financial Stability." *Financial Stability Review* (May).

Fama, Eugene F. 1970. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance* 25 (2): 383–417. doi:10.2307/2325486.

Financial Stability Board. 2024. *The Financial Stability Implications of Artificial Intelligence*. Report (14 November).

Grossman, Sanford J., and Joseph E. Stiglitz. 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70 (3): 393–408. [www.jstor.org/stable/1805228](http://www.jstor.org/stable/1805228).

Harvey, Campbell R., and Yan Liu. 2019. "A Census of the Factor Zoo." Working paper. doi:10.2139/ssrn.3341728.

International Organization of Securities Commissions. 2025. *Artificial Intelligence in Capital Markets: Use Cases, Risks, and Challenges*. Consultation Report CR/01/2025 (March).

Keynes, John Maynard. 1936. *The General Theory of Employment, Interest and Money*. Macmillan.

Kleinberg, Jon, and Manish Raghavan. 2021. "Algorithmic Monoculture and Social Welfare." *Proceedings of the National Academy of Sciences* 118 (22): e2018340118. doi:10.1073/pnas.2018340118.

Meadows, Donella H. 2008. *Thinking in Systems: A Primer*. Chelsea Green Publishing.

Nash, John F., Jr. 1950. "Equilibrium Points in  $n$ -Person Games." *Proceedings of the National Academy of Sciences of the United States of America* 36 (1): 48–49. doi:10.1073/pnas.36.1.48.

Putnam, Hilary. 1981. "Brains in a Vat." In *Reason, Truth and History*, 1–21. Cambridge University Press. doi:10.1017/CBO9780511625398.003.

Quine, Willard Van Orman. 1951. "Main Trends in Recent Philosophy: Two Dogmas of Empiricism." *Philosophical Review* 60 (1): 20–43. doi:10.2307/2181906.

## Author

Joseph Simonian, PhD  
Senior Affiliate Researcher  
CFA Institute Research and Policy Center

## ABOUT THE RESEARCH AND POLICY CENTER

CFA Institute Research and Policy Center brings together CFA Institute expertise along with a diverse, cross-disciplinary community of subject matter experts working collaboratively to address complex problems. It is informed by the perspective of practitioners and the convening power, impartiality, and credibility of CFA Institute, whose mission is to lead the investment profession globally by promoting the highest standards of ethics, education, and professional excellence for the ultimate benefit of society. For more information, visit <https://rpc.cfainstitute.org/en/>.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy, recording, or any information storage and retrieval system, without permission of the copyright holder. Requests for permission to make copies of any part of the work should be mailed to: Copyright Permissions, CFA Institute, 915 East High Street, Charlottesville, Virginia 22902. CFA® and Chartered Financial Analyst® are trademarks owned by CFA Institute. To view a list of CFA Institute trademarks and the Guide for the Use of CFA Institute Marks, please visit our website at [www.cfainstitute.org](http://www.cfainstitute.org).

CFA Institute does not provide investment, financial, tax, legal, or other advice. This report was prepared for informational purposes only and is not intended to provide, and should not be relied on for, investment, financial, tax, legal, or other advice. CFA Institute is not responsible for the content of websites and information resources that may be referenced in the report. Reference to these sites or resources does not constitute an endorsement by CFA Institute of the information contained therein. The inclusion of company examples does not in any way constitute an endorsement of these organizations by CFA Institute. Although we have endeavored to ensure that the information contained in this report has been obtained from reliable and up-to-date sources, the changing nature of statistics, laws, rules, and regulations may result in delays, omissions, or inaccuracies in information contained in this report.

First page photo credit: Getty Images/Ryzhi



**CFA Institute**

### PROFESSIONAL LEARNING QUALIFIED ACTIVITY

This publication qualifies for 1 PL credits under the guidelines of the CFA Institute Professional Learning Program.