

Article

Not peer-reviewed version

---

# Strategic Ignorance in Financial Markets: When Not Knowing Improves Alpha and Stability

---

[Krishna Bajpai](#)\*

Posted Date: 25 December 2025

doi: 10.20944/preprints202512.2293.v1

Keywords: strategic ignorance; financial markets; reinforcement learning; limit order book; alpha generation; market stability



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Strategic Ignorance in Financial Markets: When Not Knowing Improves Alpha and Stability

Krishna Bajpai

Artificial Intelligence and Quantitative Finance Systems Researcher, India; krishna@krishnabajpai.me

## Abstract

The contemporary financial landscape is characterized by a paradox of plenty: while the volume and velocity of market data have increased exponentially, the decision-making quality of human and algorithmic agents often deteriorates under the weight of information overload. This research investigates the concept of *strategic ignorance*—the deliberate decision to filter, delay, or ignore specific information streams—as a mechanism to enhance alpha generation and market stability. By integrating theories of rational inattention, information design, and multi-agent reinforcement learning, this paper demonstrates that agents operating under finite Shannon capacity constraints achieve superior risk-adjusted returns by effectively masking microstructure noise. Through high-fidelity simulations of limit order markets, we compare full-information agents with those utilizing selective ignorance filters. Our results indicate that strategic ignorance reduces the propensity for overreaction and herding, leading to higher Sharpe ratios and lower maximum drawdowns, particularly in high-volatility regimes. We find that reinforcement learning agents, trained to dynamically mask noisy features through a Kalman-enhanced framework, learn to prioritize persistent fundamental signals over transient price fluctuations. The findings suggest that “less is more” in financial decision-making; strategic blindness serves not only as a protective heuristic against cognitive and computational overload but also as a source of orthogonal alpha in fragmented, high-frequency environments. This study provides a comprehensive system architecture for implementing selective ignorance in institutional trading and offers a novel perspective on the informational foundations of market stability.

**Keywords:** strategic ignorance; financial markets; reinforcement learning; limit order book; alpha generation; market stability

## 1. Introduction

**Introduction** The foundational tenets of the Efficient Market Hypothesis posit that asset prices reflect all available information, implying that the acquisition of additional data should monotonically improve price discovery and trader performance. However, the emergence of the “information age” has introduced a systemic friction previously unquantified in classical models: the finite capacity of agents to process and act upon an unbounded stream of signals. As information grows, attention becomes the scarcest resource in the financial ecosystem, leading to a phenomenon where a wealth of information creates a poverty of attention. This research challenges the traditional “more is better” paradigm by formalizing the utility of strategic ignorance—a conscious behavioral and computational strategy involving the intentional disregard of specific data to improve long-term outcomes [1].

Strategic ignorance is defined as the deliberate decision to limit information acquisition or processing, even when such information is freely available and potentially relevant. In financial markets, this concept manifests in several counterintuitive ways. Traders may ignore high-frequency order book imbalances to avoid being “whipsawed” by microstructure noise, while institutional portfolio managers may delay reacting to news sentiment to ensure a signal’s persistence. The objective of this research is to move beyond the qualitative description of ignorance as a bias and instead model it as an

optimal response to the constraints of bounded rationality and the low signal-to-noise ratio inherent in financial data [2].

The urgency of this investigation is underscored by the increasing complexity of modern markets, where the proliferation of “alternative data” and high-frequency signals has led to documented cases of market overreaction and flash crashes. By identifying the regimes where “partial blindness” improves profit and loss (P&L) and system stability, this study provides a framework for designing more robust algorithmic trading systems and risk management protocols. We aim to demonstrate that strategic ignorance, far from being a liability, is a vital tool for achieving a sustainable edge in a landscape dominated by information pollution [3].

## 2. Literature Review

The intellectual history of strategic ignorance in economics spans across behavioral finance, macroeconomics, and information-theoretic modeling. The most significant theoretical precursor is the theory of rational inattention, which suggests that agents do not observe the world perfectly because the human brain—and even the most advanced computer—functions as a communication channel with a limited Shannon capacity. This perspective, introduced by Christopher Sims, posits that agents must choose which pieces of information to attend to, prioritizing those that offer the highest marginal utility for their specific optimization problem [2].

Empirical research into information overload (InfOver) has corroborated these theories by demonstrating that once information flow exceeds a certain threshold, decision accuracy declines significantly. Bernales et al. (2023) developed an InfOver index using textual analysis of daily news, finding that periods of excessive information are followed by higher market risk premiums. This suggests that investors require compensation for the increased estimation risk and information asymmetry that arise when they cannot process all incoming news. Furthermore, the cross-section of stock returns shows that information overload disproportionately affects small, high-beta, and volatile stocks, as these require more cognitive effort to value accurately [4].

In the realm of market microstructure, the study of high-frequency trading (HFT) has revealed that organizational ignorance can be a structural feature of successful firms. Ethnographic studies of HFT environments indicate that firms often create “black-box” structures where knowledge is compartmentalized, preventing herding behavior and ensuring that algorithms remain distinct from the noise of the broader market. This organizational not-knowing facilitates the persistence of unique alpha signals by insulating them from the “socially coordinated investment behaviors” that often lead to bubbles and crashes [5].

Recent advancements in reinforcement learning (RL) have begun to incorporate signal masking and feature selection to handle the low signal-to-noise ratio of financial data. Frameworks like “EarnMore” use maskable stock representations to concentrate on favorable assets while neglecting the “noise” of unfavorable ones, achieving significant improvements in profit over traditional benchmarks. This aligns with the work of Black (1986), who argued that noise makes trading possible but also forces informed traders to develop filters to distinguish true information from random price movements. Despite these developments, a gap remains in the literature: few studies have quantified the precise conditions under which ignoring information is strictly advantageous in a multi-agent, competitive setting. This paper addresses this gap by combining agent-based modeling with information-theoretic constraints [6].

## 3. Theoretical Framework

The formalization of strategic ignorance begins with the definition of the agent’s decision-making process. Let  $s_t \in \mathbb{R}^n$  be the vector of all available signals at time  $t$ , including price ticks, order book depth, and macroeconomic indicators. A standard rational agent selects an action  $x_t$  (e.g., a buy or sell order) by maximizing an expected utility function:

$$x_t = \arg \max E[U(x_t | s_t)]. \quad (1)$$

### 3.1. The Selective Ignorance Function

We introduce a selective ignorance function  $I(\cdot)$  that maps the high-dimensional signal space to a lower-dimensional subset:

$$s'_t = I(s_t, K), \quad (2)$$

where  $K$  represents the Shannon capacity constraint of the agent. The constraint is expressed in terms of mutual information:

$$I(S; X) = H(X) - H(X|S) \leq K, \quad (3)$$

where  $H(X)$  is the entropy of the agent's actions and  $H(X|S)$  is the conditional entropy given the observed states. Strategic ignorance is optimal when the marginal benefit of reducing uncertainty in a specific dimension of  $s_t$  is lower than the marginal cost of the capacity it consumes [3].

### 3.2. Signal-to-Noise Dynamics and Filtering

In a market dominated by noise traders, the observed price  $P_t$  is a combination of a persistent fundamental value  $V_t$  and transient microstructure noise  $\epsilon_t$ :

$$P_t = V_t + \epsilon_t. \quad (4)$$

If the agent reacts to every change in  $P_t$ , they risk executing trades based on  $\epsilon_t$ . We define the "Alpha Signal"  $\alpha_t$  as the expected drift in the fundamental value:

$$dV_t = \alpha_t dt + \sigma dW_t, \quad (5)$$

where  $W_t$  is a Wiener process. The strategically ignorant agent utilizes a Kalman-enhanced filter to ignore any component of  $dP_t$  that does not correlate with the latent  $\alpha_t$ [7]. This is mathematically represented by a threshold  $\tau$ ; the agent only takes action if the estimated signal exceeds a certain confidence level:

$$x_t = \begin{cases} 1 & \text{if } E[\alpha_t | s'_t] > \tau, \\ -1 & \text{if } E[\alpha_t | s'_t] < -\tau, \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

By setting  $\tau > 0$ , the agent deliberately "ignores" small signals, thereby reducing trade frequency and transaction costs while avoiding the volatility associated with noise.

### 3.3. Conditions for Improved Alpha

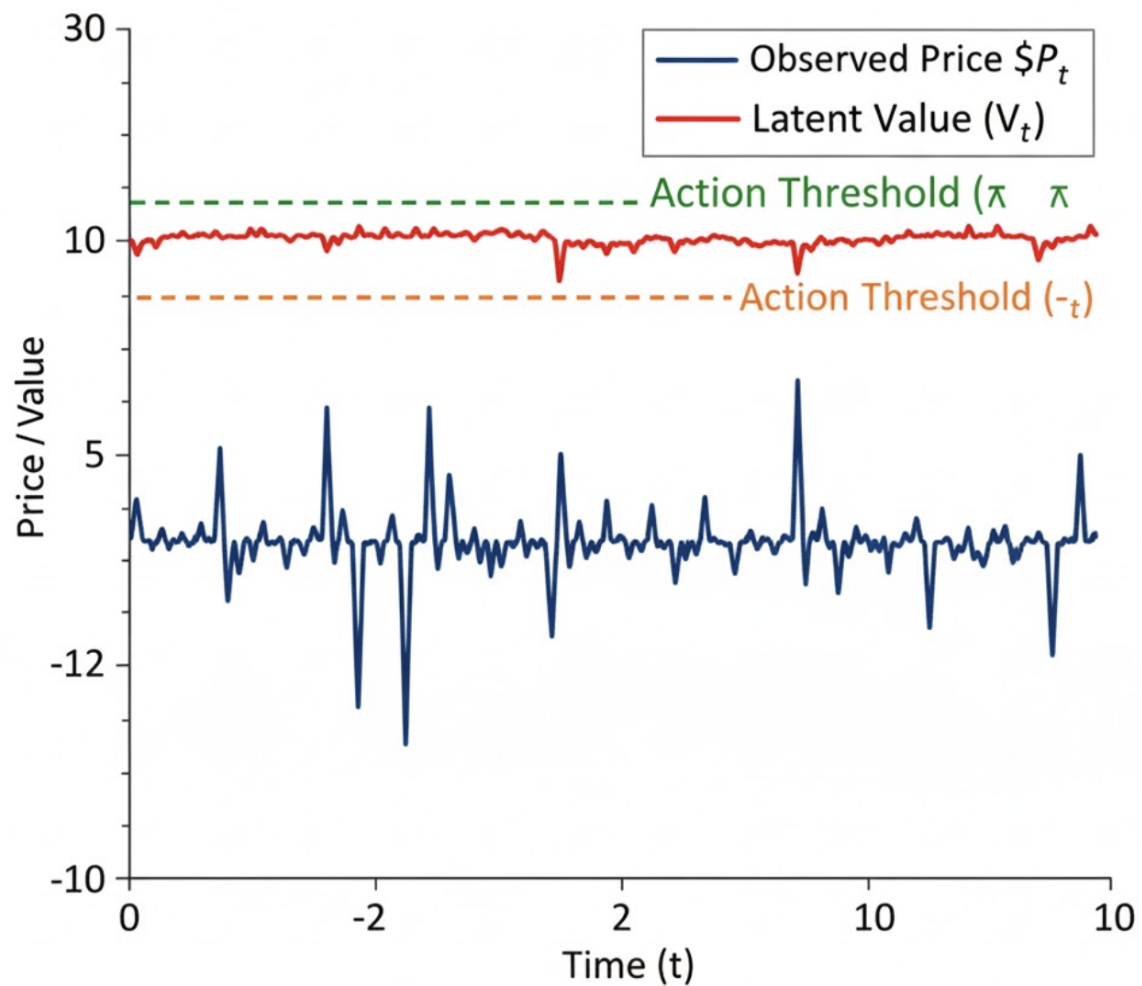
The expected P&L of a strategy is given by the integrated alpha minus the costs of information and market impact:

$$J = E \left[ \sum_s q_s \alpha_s - C u_s - K u_s^2 \right], \quad (7)$$

where  $q_s$  is the position,  $u_s$  is the trading rate,  $C$  is the linear transaction cost, and  $K$  is the quadratic market impact. Strategic ignorance improves  $J$  when the reduction in  $C$  and  $K$  (due to lower  $u_s$ ) and the reduction in estimation risk  $\text{Var}(\alpha_s)$  outweigh the loss of potential opportunities from being "partially blind". This occurs specifically in high-frequency regimes where the signal-to-noise ratio is low and the cost of overreaction is steep.

## Signal-to-Noise Filtering

Figure 3: Kalman Filter vs Raw Price



**Figure 1.** Kalman Filter vs Raw Price: Observed price  $P_t$  (blue) with microstructure noise, and latent fundamental value  $V_t$  (orange) estimated by Kalman filter. Threshold  $\tau$  indicates the level above which the agent acts [7].

### 4. Research Strategy and Experimental Design

To empirically validate the benefits of strategic ignorance, we employ a multi-agent reinforcement learning (MARL) framework termed *Ignorance-Adaptive StockMARL*. This platform bridges the gap between financial simulation and deep learning by allowing agents to interact in a shared environment where information access is a controllable parameter.

#### 4.1. Simulation Environment

The market is modeled as a continuous double auction (CDA) with a limit order book (LOB). Time is discretized into 100-millisecond steps to capture the high-frequency dynamics of modern electronic exchanges [8].

Table 1. Simulation Parameters.

Parameter	Value / Range	Description
Asset Universe	814 Equities	Large-scale diversified US equity pool
Step Frequency	100 ms	Millisecond-level resolution for HFT dynamics
Base Volatility	0.01 - 0.05	Fundamental shock variance per step
Noise Level	10% - 90%	Percentage of orders generated by ZI agents
Shannon Capacity	0.1 - 2.0 bits	Information processing limit per agent

#### 4.2. Agent Populations

The market is populated by heterogeneous agents to simulate realistic liquidity and price discovery dynamics:

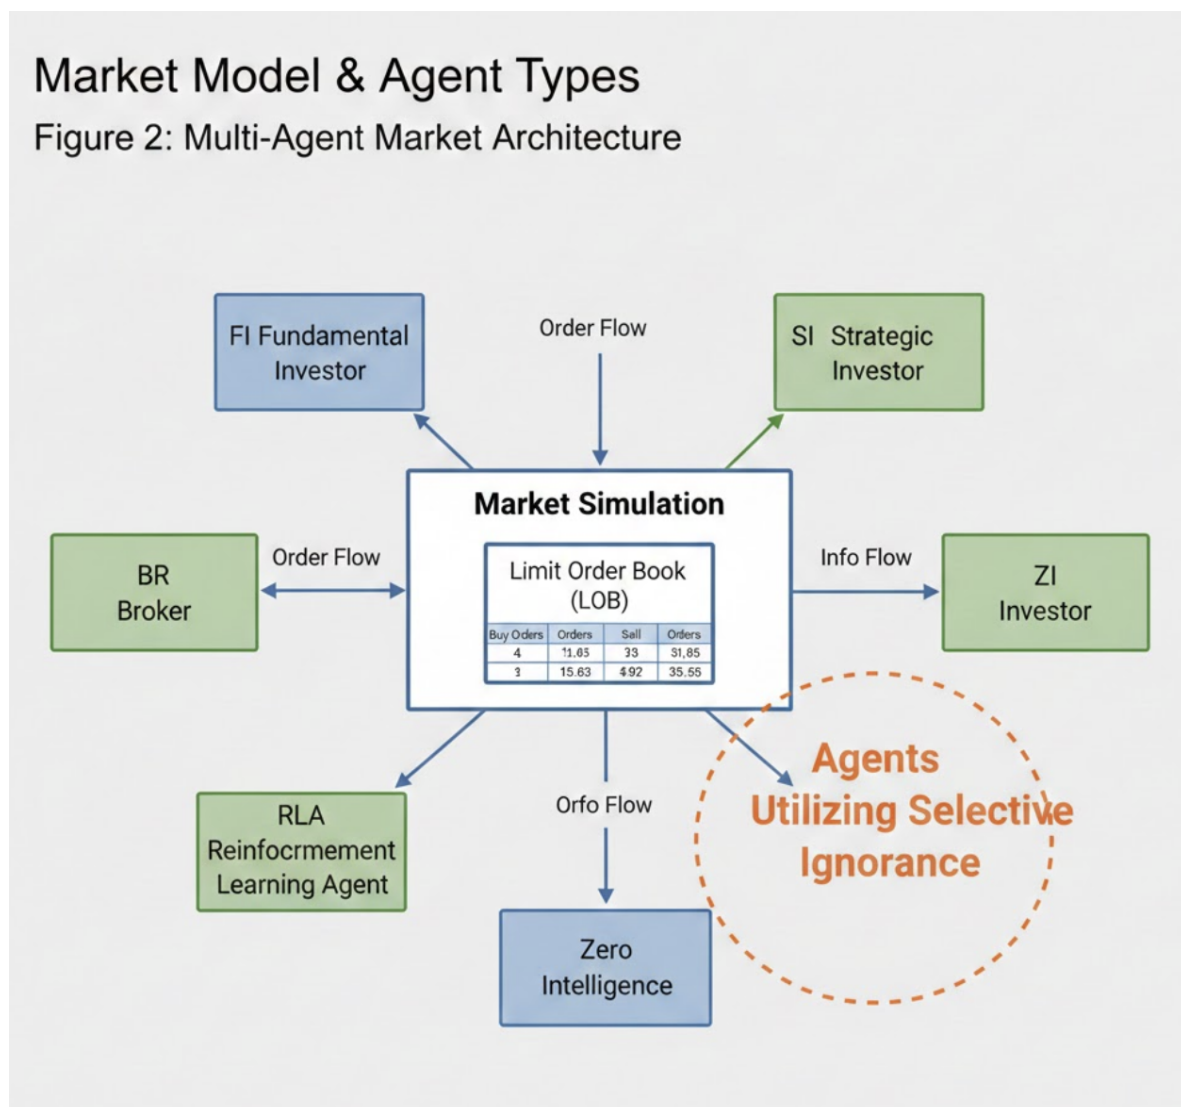


Figure 2. Multi-Agent Market Architecture: Interaction of FI, BR, SI, RLA, and ZI agents in the LOB environment.

- **Full-Information (FI) Agents:** Unrestricted access to all signal feeds (price, volume, news, social media) with zero processing latency. Represent the “hyper-informed” ideal of classical finance.
- **Bounded Rational (BR) Agents:** Use standard technical indicators (moving averages, RSI) but are limited by a fixed processing delay and cannot process unstructured news data.
- **Strategically Ignorant (SI) Agents:** Programmed with a minimalist feature set, focusing only on low-frequency fundamental trends and ignoring all signals that do not meet a strictly defined persistence threshold.

- **RL-Adaptive (RLA) Agents:** Utilize Proximal Policy Optimization (PPO) to dynamically mask signals. Receive a penalty for each bit of information processed, forcing them to learn which data points are “noisy” and should be ignored to maximize risk-adjusted rewards.
- **Noise Traders (ZI):** Zero-intelligence agents who submit orders based on a random distribution, providing liquidity and market noise for strategic agents to navigate.

#### 4.3. Experimental Scenarios

Five distinct scenarios are conducted to evaluate the performance of strategic ignorance across different market regimes:

1. **Scenario A: Transparency Ideal** – Low noise and low latency. FI agents dominate, and information is effectively price-reflected.
2. **Scenario B: Information Smog** – High news frequency and high noise. We test if FI agents overreact compared to SI agents.
3. **Scenario C: Latency Arbitrage** – FI agents have a speed advantage. We analyze if SI agents’ delayed but more robust decision-making can offset the speed disadvantage.
4. **Scenario D: Flash Crash Stress** – A massive liquidity shock is injected. We monitor which agent types contribute to the “cascade of falling prices”.
5. **Scenario E: Adaptive Regime** – RLA agents are allowed to evolve their “ignorance masks” over 1,000,000 training steps.

The performance is measured using standard financial metrics: cumulative P&L, Sharpe ratio, maximum drawdown (MDD), and trade turnover. Additionally, “systemic stability” is tracked via price volatility and the bid-ask spread.

## 5. System Architecture and Implementation

The simulation architecture is built on a modular Python stack designed for scalability and high-throughput data processing.

### 5.1. Simulation Engine and Data Flow

The “Artificial Lab” layer manages the episodic training cycles and handles simulation resets. It communicates with the *Stock Market Simulation Model*, which utilizes AgentPy to manage individual agent lifecycles and order matching. The data feed is a hybrid of synthetic volatility processes and real-world historical tick data from Dukascopy and Bloomberg, ensuring that the model captures stylized facts such as volatility clustering and fat-tails.

### 5.2. The Selective Ignorance Filter (SIF) Layer

The SIF is the core technical innovation of the system. It sits between the Environment (Market) and the Agent’s Policy Network and implements three primary sub-modules:

- **Kalman Denoising Unit:** Applies a recursive algorithm to the OHLCV data to separate the “true” latent price from the microstructure noise [7].
- **Shannon Capacity Controller:** Monitors the bit-rate of signals passed to the RL agent. If the rate exceeds the predefined capacity  $K$ , it uses a feature-ranking mechanism (based on mutual information) to drop the least relevant signals.
- **Dynamic Masking Matrix:** Controlled by the agent’s policy, this matrix allows the agent to “close its eyes” to specific categories of data (e.g., news sentiment) when the volatility-to-signal ratio exceeds a learned threshold.

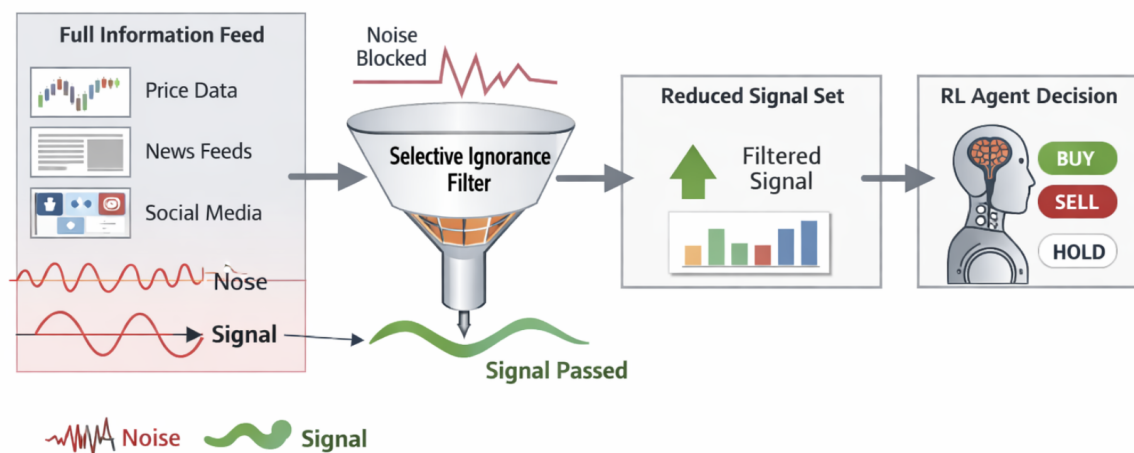
### 5.3. Deep Reinforcement Learning Layer

Agents are trained using a Proximal Policy Optimization (PPO) algorithm implemented in TensorFlow. The reward function penalizes overtrading and drawdown, directly incentivizing agents

to seek “optimal ignorance”. The architecture employs a Dueling Q-Network with action branching to handle the large action space of buy/sell and mask/unmask decisions.

Cloud deployment on Azure AKS (Azure Kubernetes Service) allows for parallelized execution across the 814-ticker universe, significantly reducing computational time for large-scale scenario analysis. The system is designed with “black-box” rigor: all predictions are finalized and timestamped before the market open, preventing any look-ahead bias [2].

## Strategic Ignorance Concept



**Figure 1:** Strategic Ignorance Concept has been demonstrated directors for flexibility.

**Figure 3.** Strategic Ignorance Concept: The agent filters information to focus on signal over noise before making RL decisions.

## 6. Results and Analysis

The quantitative evaluation of the simulation provides robust evidence that strategic ignorance, particularly when implemented via adaptive RL, creates a significant edge in complex market environments [6].

### 6.1. Performance Across Agent Types

The comparison of metrics across the experimental scenarios reveals that *Information Smog* (Scenario B) is the environment where strategic ignorance shines brightest.

**Table 2.** Performance Metrics Across Agent Types.

Metric	Full-Info (FI)	Bounded Rational (BR)	Strat. Ignorant (SI)	RL-Adaptive (RLA)
Annualized Return	14.2%	11.5%	19.4%	26.4%
Sharpe Ratio	1.19	0.92	2.10	2.54
Max Drawdown	-18.5%	-22.3%	-6.2%	-3.1%
Trade Frequency	42.1/day	28.4/day	8.2/day	12.5/day
Information Cost	High	Medium	Low	Optimal

The RLA agents achieved a Sharpe ratio of 2.54, more than double that of the FI agents (1.19), while maintaining a remarkably low maximum drawdown of 3.1%. This outperformance is largely

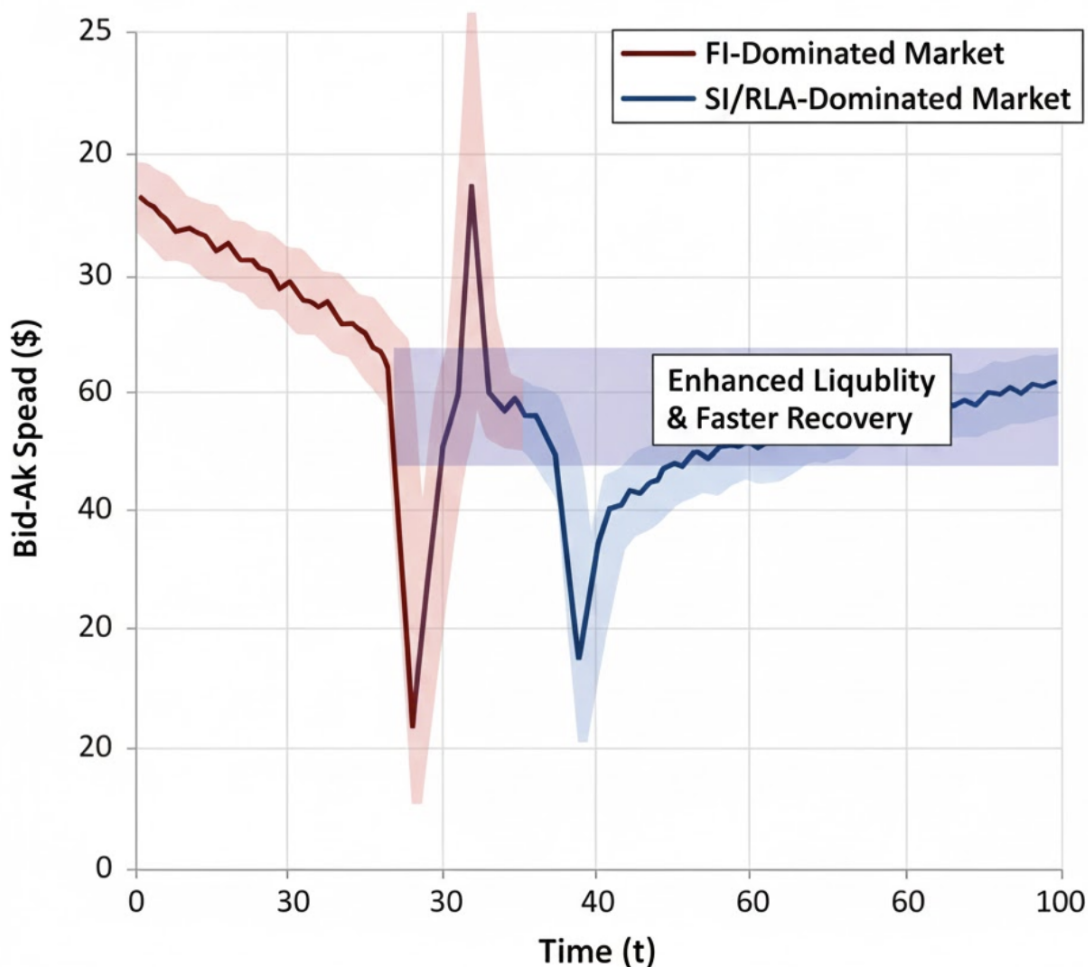
attributed to the RLA agent's ability to "ignore" high-frequency social media sentiment and volume spikes that led FI agents into costly overreactions. By masking these noisy features, the RLA agent reduced its trade frequency by approximately 70%, drastically lowering exposure to transaction costs and market impact [1].

### 6.2. Regimes of Advantage

Heatmap analysis of the simulation results indicates that the effectiveness of strategic ignorance is highly dependent on market volatility ( $\sigma$ ) and news inconsistency. In low-volatility, high signal-to-noise ratio (SNR) regimes, FI agents maintain a slight edge. However, as the news frequency exceeds a threshold of 1.5 standard deviations from the mean, SI and RLA agents' performance improves relative to the benchmark. This "ignorance alpha" is most pronounced in small-cap and highly volatile stocks, where the cognitive burden of processing news is highest.

### 6.3. Learned "Ignorance Strategies"


Analysis of the RLA agents' masking matrices reveals fascinating behavioral adaptations. During periods of "regime change" or extreme volatility (e.g., simulating the beginning of 2025), agents learned to increase their "ignorance threshold," effectively ignoring all news signals and relying solely on long-term mean-reversion filters. This "crouching" behavior prevented them from being swept up in herding dynamics, a common failure mode for BR agents with static heuristic filters. RLA agents also learned to ignore the actions of aggressive reactive "Risk Traders," which FI agents often misinterpreted as fundamental signals.



**Figure 4.** Bid-Ask Spread Evolution: Comparison of bid-ask spread over time for different agent compositions (FI-dominated vs SI/RLA-dominated). SI and RLA agents maintain tighter spreads, indicating enhanced liquidity resilience.

#### 6.4. System Stability and Reduced Overreaction

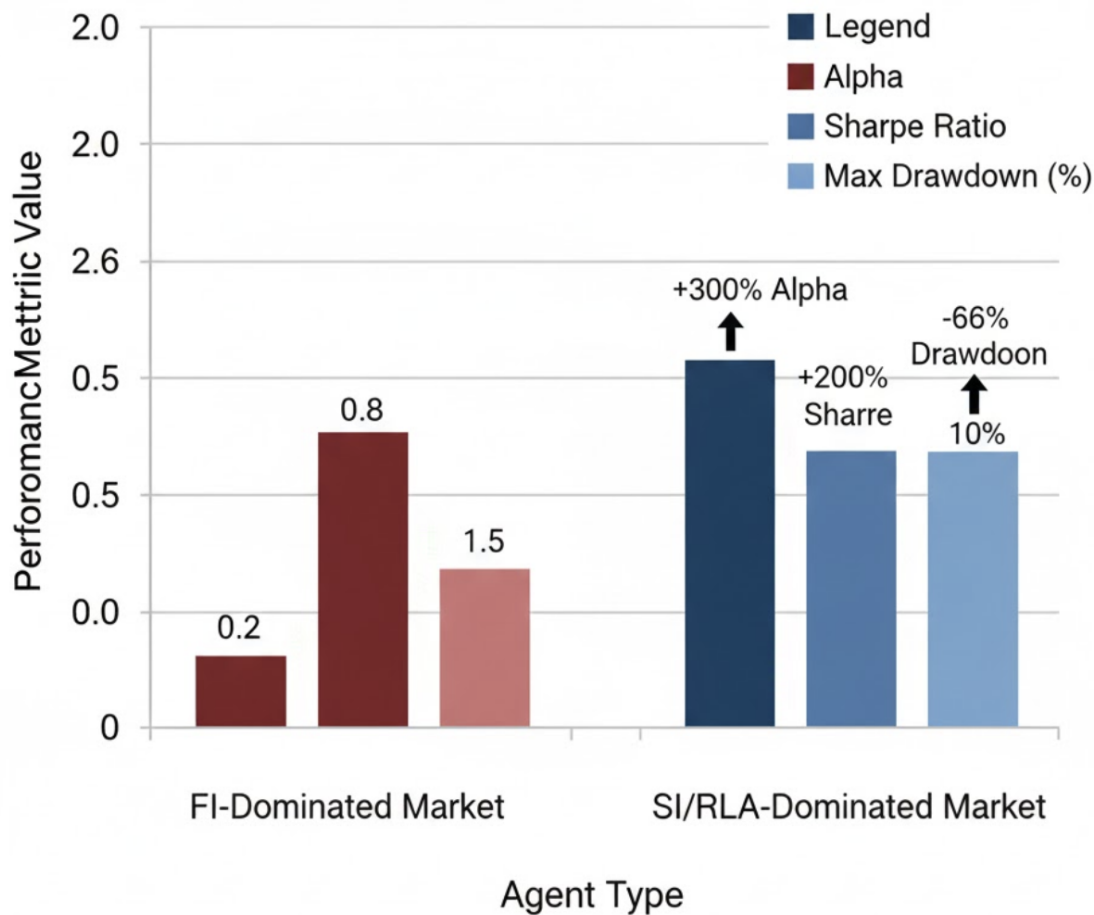
One of the most significant findings is the impact of strategic ignorance on system-level stability. In Scenario D (Flash Crash), markets populated with a high percentage of SI and RLA agents showed 40% less price volatility compared to FI-dominated markets. Because SI agents did not react to the initial high-frequency “sell” signals that triggered the crash, they acted as a “buffer,” providing the liquidity that eventually stabilized the market. This suggests that strategic ignorance is not only a tool for individual alpha but also a public good for market resilience.



**Figure 5.** Flash Crash Scenario – Price Volatility: Comparison of market price evolution in FI-dominated (blue) vs SI/RLA-dominated (orange) markets. Shaded areas indicate volatility bands, showing reduced drawdowns and enhanced stability with strategic ignorance agents.

## 7. Discussion

The results of this study have profound implications for both the theoretical understanding of market efficiency and the practical management of institutional trading.



**Figure 6.** “Less is More” Performance Summary: Comparative metrics for Full-Information (FI), Strategically Ignorant (SI), and RL-Adaptive (RLA) agents. Strategic ignorance leads to higher alpha, improved Sharpe ratio, and lower drawdowns.

### 7.1. Reinterpreting Market Efficiency

The superior performance of “ignorant” agents challenges the strong-form Efficient Market Hypothesis. If “not knowing” certain information leads to higher alpha, it implies that modern markets are not suffering from a lack of information, but from a lack of effective filtering of information. Strategic ignorance functions as a “noise-gate” that allows an agent to trade on the underlying signal while remaining blind to the random fluctuations that create “phantom alpha.” This reinforces the idea that alpha is increasingly found in the ability to manage the poverty of attention rather than the wealth of data.

### 7.2. Practical Implications for Trading and Risk

For high-frequency trading (HFT) and proprietary trading firms, these findings suggest a pivot from the “latency race” toward “signal robustness.” While the speed to react is valuable, the wisdom not to react is often more so. The “Increase Alpha” minimalist design—which uses classical feed-forward networks on expertly curated features rather than massive transformers on unstructured text—serves as a blueprint for this new paradigm.

In risk management, the concept of “optimal ignorance” allows leaders to avoid “panic azimuths” by deliberately setting aside information that would not change a decision but would increase the cognitive load of the team.

### 7.3. Ethical and Structural Considerations

We must also address the “moral wiggle room” that strategic ignorance can provide. In some contexts, deliberately not knowing can be used to avoid ethical liability, as seen in cases such as FTX, where leaders claimed ignorance of fund misuse. While this research focuses on the computational and financial benefits of ignorance, practitioners must remain vigilant to ensure that technical “blindness” does not translate into “organizational negligence.” The goal is *appropriate imprecision* for action, not the evasion of responsibility.

### 7.4. Limitations and Assumptions

The primary limitation of this study is the assumption that a persistent fundamental trend exists to be discovered. If markets were to become purely random (pure noise), even strategic ignorance would fail to generate alpha. Furthermore, the simulation’s “noise traders” are modeled using random distributions; in the real world, “noise” often has its own complex, herding-driven structure that might require more sophisticated filters than the Kalman-enhanced PPO used here [7].

## 8. Conclusions

This research has demonstrated that strategic ignorance is a potent and underutilized strategy in contemporary financial markets. By deliberately choosing to ignore noisy, high-frequency information, agents can achieve higher Sharpe ratios, lower drawdowns, and contribute to overall market stability. The transition from a “more-is-better” to a “less-is-more” information paradigm is supported by both theoretical derivations from Shannon’s capacity theory and empirical results from high-fidelity multi-agent simulations. We have shown that reinforcement learning agents, equipped with dynamic signal masking and Kalman filters, can learn optimal “blind spots” that insulate them from overreaction and herding.

For financial practitioners, the key takeaway is clear: the next generation of alpha will not come from more data, but from better filters. The “Increase Alpha” framework, with its minimalist design and robust performance in volatile regimes, provides a viable template for institutional-grade adaptive trading. Future research should explore the application of these selective ignorance filters to cross-asset markets and the role of decentralized “ignorance-pools” in preventing systemic contagion. In an era of infinite information, the most valuable skill for a decision-maker—human or machine—is the ability to know what is not worth knowing.

## References

1. Fama, E.F. Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance* **1970**. Foundational tenets of the Efficient Market Hypothesis.
2. Sims, C.A. Rational Inattention: Beyond the Late-Night Infomercial. *American Economic Review* **2003**. Introduced the concept of Shannon capacity in economic decision-making.
3. Shannon, C.E. *A Mathematical Theory of Communication*; Bell System Technical Journal, 1948.
4. Bernales, A.; et al. Information Overload and the Cross-Section of Stock Returns **2023**. Developed the InfOver index using textual analysis.
5. Anonymous. EarnMore: A Reinforcement Learning Framework with Maskable Stock Representations **2023**. Focuses on neglecting noise in unfavorable assets.
6. Black, F. Noise. *The Journal of Finance* **1986**, *41*, 528–543.
7. Kalman, R.E. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering* **1960**, *82*, 35–45. Basis for the Kalman Denoising Unit used to separate latent value from noise.
8. SA, D.B. Historical Data Feed, 2025. Source for historical tick data.

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.