Student Journal of Business and Economics
Vol. 1, No. 4 (2026), pp. 103–116 • ISSN 3144-0363
DOI: 10.67521/sjbe.2026.038

Original Research

Evaluating AI Disclosure Practices and Regulatory Gaps in Quantitative Finance

Ryan Hong

Regis Jesuit High School, 6400 S. Lewiston Way, Aurora, Colorado 80016, United States
Corresponding author email: ryan.yc.hong@gmail.com

Accepted: October 31, 2026

Abstract

As artificial intelligence becomes increasingly prominent within financial markets, regulators lack the frameworks necessary to assess whether AI-driven investment strategies comply with fiduciary standards, manage systemic risk responsibly, or produce returns through legitimate means. This gap is particularly prevalent among leading quantitative trading and investment firms, which deploy sophisticated machine learning systems while simultaneously disclosing little about how those systems are built or monitored. This paper analyzes the publicly available AI disclosure practices and performance of five major quantitative trading and investment firms (Renaissance Technologies, Two Sigma, D. E. Shaw, Citadel, and Jane Street) from Q1 2023 to Q1 2026, comparing returns from their publicly disclosed top-20 equity holdings against the S&P 500 using data from the website HedgeFollow. The findings reveal that all five firms score below 40 out of 100 on a disclosure rubric created for this study, with the largest gaps appearing in model governance, retraining procedures, risk limits, stress-response behavior, and incident-reporting protocols. At the same time, three firms—Renaissance Technologies, D. E. Shaw, and Jane Street—display peak cumulative return gaps ranging from approximately 40 to over 100 percentage points above the S&P 500 during the study period, demonstrating that strong performance can coexist with limited public visibility into the systems and controls behind that performance. The paper concludes with four targeted regulatory proposals: mandatory AI disclosure standards, explainability requirements, systemic risk monitoring, and human oversight attribution rules, all of which must be implemented while preserving market innovation.

Keywords: Artificial Intelligence; Quantitative Finance; Algorithmic Trading; Financial Regulation; AI Disclosure; Systemic Risk; Innovation

1. Introduction

1.1 The Problem: Why Regulating AI in Portfolio Management Is Critical

This paper argues that effective regulation of AI in portfolio construction must be rooted in empirical evidence which can be done by analyzing how leading quantitative hedge funds deploy these technologies and the results of this usage.

While AI and machine learning have the undeniable potential to develop superior trading strategies, there is a need for greater regulation and transparency about their use to protect investors and financial

institutions. For example, while AI applications may have supported firms in noticing market trends quicker and analyzing large amounts of data, they have also introduced systemic risks including flash crashes, herd behavior amplification, and opaque decision-making processes. A prominent example of these risks appeared during the 2010 Flash Crash, when quantitative trading algorithms contributed to a rapid market collapse that erased nearly $1 trillion in market value within minutes. Although the event was not caused by modern AI systems, it demonstrates how automated trading can amplify market instability and make the causes of financial disturbances difficult to identify (CFA Institute 2019).

Consequently, in this new era, effective regulation is necessary to balance innovation with accountability. In terms of methodology, this paper analyzes the AI disclosure practices and performance of five leading quantitative trading and investment firms from 2023 to 2026, benchmarking their returns against the S&P 500 to identify disclosure gaps and develop targeted regulatory proposals.

1.2 What is AI Deployment in Portfolio Construction and Asset Allocation?

In institutional investment management, AI is utilized by applying machine learning algorithms, deep neural networks, and reinforcement learning to a range of portfolio management functions. These functions include predicting market movements, generating trading signals, adjusting portfolio allocations, and evaluating portfolio performance. AI has empowered finance traders, enabling small trading teams to solve complex portfolio optimization problems efficiently and at scale (Ghatak et al. 2025; Sutiene et al. 2024). Complementing these functions is AI's ability to process significant amounts of alternative data in real time to execute strategies that no human analyst could replicate. Ultimately, AI-driven portfolio management not only increases efficiency but also reshapes the financial markets industry.

1.3 What are the issues with unregulated AI in finance?

With the widespread integration of artificial intelligence in financial markets, several significant risks have emerged. As mentioned earlier, one of the primary concerns is the potential for flash crashes, where AI-driven trading systems amplify market movements, leading to sharp and sudden declines in asset prices. This risk is demonstrated by the UK Government Office for Science, which argues that the Flash Crash exposed the possibility of positive feedback loops induced by automated technology, as execution algorithms interacted with high-frequency trading strategies in ways that amplified market stress, leaving financial markets vulnerable (Government Office for Science 2012).

Specific historical events such as the 2010 Flash Crash and the U.S. Treasury market disruption in 2014 demonstrate how these automated and high-frequency trading systems can magnify instability even in highly liquid markets (Svetlova 2022, 715). Furthermore, as AI adoption continues to expand, these feedback-driven dynamics underscore the need for stronger regulations designed to not only monitor individual models, but also their impact on market stability.

In addition to systemic risks, unregulated AI raises serious transparency and accountability challenges. Advanced models, particularly deep neural networks, are often highly complex and difficult to interpret, resulting in the "black box" problem (Financial Stability Board 2017, 13). This opacity makes it difficult for regulators and investors to understand how decisions are made and whether advanced models comply with fiduciary and risk management standards. For regulators such as the SEC, this lack of interpretability limits their ability to determine whether an AI model is executing a legitimate allocation strategy or slowly contributing to widespread systemic exposure. For investment advisers and institutional investors, the same opacity creates a fiduciary and compliance gap because they may be unable to fully verify whether AI-driven portfolio decisions align with client interests, disclosure obligations, and internal risk controls.

2. Methods

2.1 Proposal: Targeted Regulation

This paper proposes that an effective path to AI regulation in portfolio management is grounded in empirical analysis of existing disclosure practices at the industry's most technologically advanced firms. By

examining the publicly available financial reports of five leading US quantitative trading and investment firms—Renaissance Technologies, Two Sigma, D. E. Shaw, Citadel, and Jane Street—for the period 2023 to 2026, this paper identifies observable disclosure patterns, firm-specific practices, and critical gaps in transparency. In addition, this analysis draws on secondary industry sources to further observe how each firm deploys AI and manages associated risks. This empirical foundation is then used to develop a concrete regulatory proposal as a base for the future.

2.2 Methodology

This study employs a mixed approach combining quantitative performance analysis with qualitative assessment of disclosure practices to fully examine the role of AI in portfolio management. First, the performance of selected quantitative hedge funds is evaluated using publicly disclosed equity holdings data to compare cumulative returns against the S&P 500 and identify patterns of relative outperformance or underperformance. More specifically, this study uses a manager-weighted top-20 holdings metric which evaluates the performance of each fund's twenty largest publicly disclosed equity positions, weighting each holding according to its relative size in the manager's portfolio. Consequently, larger disclosed positions mean a proportionally greater impact on the calculated portfolio return than smaller positions.

To further analyze these results, fund performance is compared against established market benchmarks: specifically, the S&P 500. This benchmark was selected because it is widely considered a standard benchmark of equity market performance, offering this paper a consistent baseline for evaluating returns.

While quantitative metrics, such as the returns derived from the manager-weighted top 20 holdings, reveal whether AI-driven funds outperform traditional benchmarks, qualitative analysis of firm disclosures is also necessary to understand how that performance is generated. Drawing from these disclosures, the study argues that regulation should mandate AI governance frameworks, explainability standards, and performance benchmarking to ensure the fulfillment of fiduciary duties without constricting innovation.

2.3 Regulatory Baseline: Current Form ADV Disclosure Requirements

Furthermore, Form ADV, the SEC's primary investment adviser disclosure document, does not contain a dedicated section requiring firms to describe AI model architecture, training data or model-specific parameters. However, investment advisers remain subject to existing disclosure, compliance, and fiduciary obligations. Chris Stanley from Kitces recommends that advisers disclose material uses of AI in their Form ADV Part 2A brochures, including relevant risks, limitations, and oversight mechanisms (Stanley 2025). These recommendations illustrate how existing disclosure requirements could be supplemented with more standardized AI-specific information.

3. Results: Analysis of Leading Quant Firms

The following section provides a fund-by-fund overview of AI disclosure practices and relative performance for the five target firms. Performance figures represent statistics drawn from HedgeFollow.

3.1 Renaissance Technologies

3.1.1 AI Disclosure and Strategies

Renaissance Technologies, originally founded as Monemetrics by James Simons in 1978 and renamed Renaissance Technologies in 1982, is widely regarded as a pioneering quantitative hedge fund. The firm specializes in systematic trading, leveraging quantitative models derived from mathematical and statistical analysis, and incorporates statistical arbitrage, which uses machine learning to identify and exploit statistical relationships between assets (TrendSpider n.d.-b). However, Renaissance enforces an extreme level of internal and external secrecy as employees are required to sign non-disclosure agreements, and internal access to trading strategies is scattered across departments so individual researchers are only provided what is necessary for their respective roles (Quantified Strategies 2026). Furthermore, Simons himself wrote to his investors, claiming that "visibility invites competition" and that the firm's "only

defense is to keep a low profile" (quoted in Scrivner 2025). Consequently, Renaissance's limited public disclosure makes it difficult for outside investors and researchers to assess its model governance and risk controls.

3.1.2 Cumulative Returns vs. the S&P 500

Based on HedgeFollow's manager-weighted top-20 holdings metric, Renaissance Technologies' publicly disclosed equity holdings outperformed the S&P 500 during the study period, with the largest observed cumulative return gap exceeding 100 percentage points in Q3 2025. While the observed benchmark-relative outperformance is compatible with the possibility of effective quantitative investment methods, the publicly available data does not establish whether AI-driven systems contributed to these returns or whether those systems are explainable or compliant with current regulatory standards. Accordingly, the regulatory proposal in Section 4 argues that AI-driven portfolio construction should be evaluated through both performance outcomes and the transparency of the mechanisms behind them.

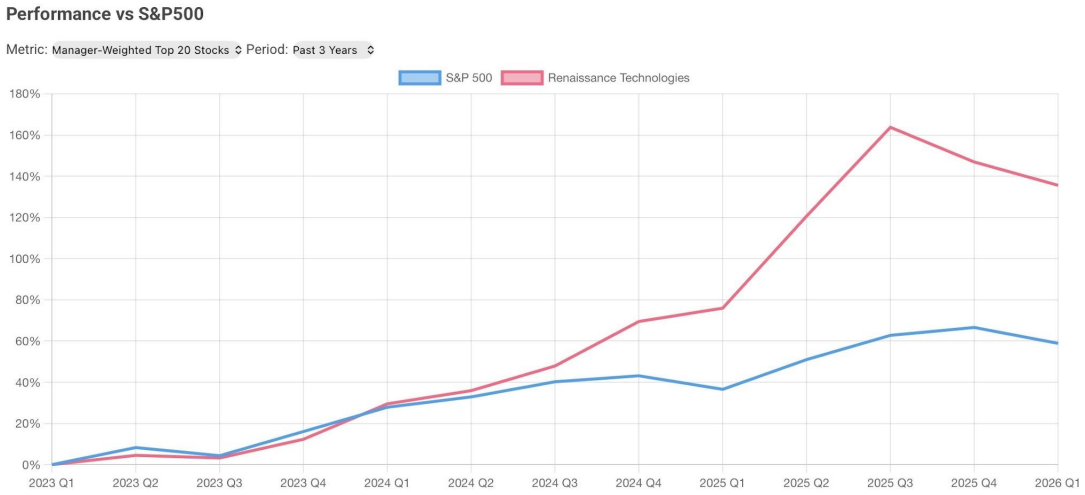

Figure 1. Renaissance Technologies performance vs. S&P 500, manager-weighted top-20 holdings, Q1 2023–Q1 2026. Source: HedgeFollow n.d.-d.

3.2 Two Sigma

3.2.1 AI Disclosure and Strategies

Two Sigma is another prominent quantitative investment firm. Founded in 2001, Two Sigma utilizes quantitative research and computation techniques to identify investment opportunities. Its approach incorporates statistical modeling, machine learning and data analysis to inform trading decisions (TrendSpider n.d.-c). When looking at their levels of disclosure in comparison to other large investment firms, Two Sigma discloses relatively more about its technological approach, displaying a machine learning section on its website where researchers publish data science articles and insights (Two Sigma 2026).

However, this transparency masks significant gaps in the core model governance disclosure, as its public filings provide very limited detail on retraining protocols, risk parameter controls, or how they resolve conflicting model signals. And separately, in 2025, the SEC charged Two Sigma with breaching its fiduciary duties by failing to address known vulnerabilities in its investment models (U.S. Securities and Exchange Commission 2025). According to the SEC's order, these vulnerabilities were identified as early as March 2019, but were not addressed for over four years, during which a single employee made unauthorized parameter modifications across fourteen live trading models (Dakak 2025). The firm also violated the SEC's whistleblower protection rule by requiring departing employees to sign separation agreements which stated that they had not filed any complaints with any government agency. Resulting from these model governance and supervisory failures, Two Sigma voluntarily repaid $165 million to affected funds and agreed to pay $90 million in civil penalties to settle the charges (U.S. Securities and Exchange Commission 2025). The Two Sigma enforcement action illustrates the consequences of

inadequate internal model governance. Although the action did not specifically establish a failure of public AI disclosure, it provides evidence supporting the importance of model governance and supervisory controls.

3.2.2 Alpha Generation vs. Benchmarks

Over the period analyzed, Two Sigma's publicly disclosed top-20 equity holdings generally performed similarly to the S&P 500 with certain quarters exhibiting outperformance by the S&P 500 benchmark. In fact, from 2023 Q2 to 2025 Q1, the S&P 500 demonstrated stronger growth. However, the absence of consistent outperformance among the firm's publicly disclosed top-20 equity holdings does not directly mean that the firm's AI driven strategies generate value across its investment portfolio. Nevertheless, these observations highlight the importance of evaluating both investment performance and model governance.

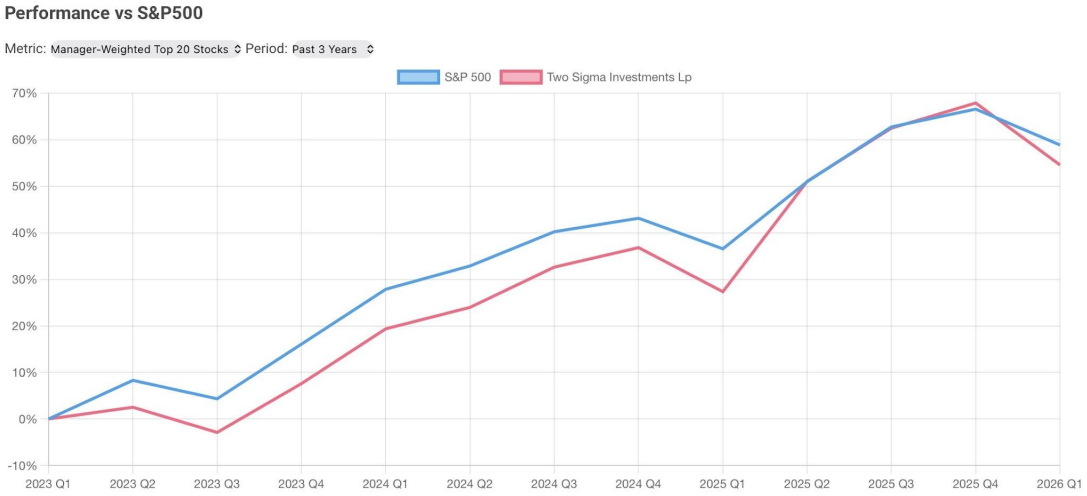

Figure 2. Two Sigma performance vs. S&P 500, manager-weighted top-20 holdings, Q1 2023–Q1 2026. Source: HedgeFollow n.d.-e.

3.3 D. E. Shaw

3.3.1 AI Disclosure and Strategies

Another major quantitative firm is D. E. Shaw & Co., which has long been recognized as another trailblazer in systematic investing. According to TrendSpider, the firm employs a wide range of quantitative strategies, relying on mathematical modeling, computer science, and statistical analysis to identify inefficiencies across global markets. These models process vast amounts of market data and generate trading signals that guide portfolio construction and risk management (TrendSpider n.d.-a). Unlike Renaissance Technologies, D. E. Shaw publicly acknowledges using a combination of quantitative tools and human-driven insights, and the firm publicly states that it works to "uncover independent, hard-to-find sources of return across global public and private markets" (D. E. Shaw Group n.d.). However, this acknowledgement reveals very little about its actual practices, as one industry analyst notes the firm is "one of those places where you just don't know exactly what [it is] they do, except that it is some mix of quantitative and discretionary investing" (Cheung 2020).

Furthermore, D. E. Shaw publicly describes its investment approach as combining quantitative tools with human-driven insights, with strategies ranging from systematic to discretionary investing (D. E. Shaw Group n.d.). However, its public disclosures do not fully describe how investment decisions are divided between the quantitative models and human decisions. This makes it difficult for outside investors to fully evaluate the role each plays in the ultimate performance outcomes, which further underlines the importance of requesting greater transparency.

In addition, D. E. Shaw's hybrid model, combining highly automated quantitative systems with active human oversight, creates a fundamental attribution issue. Regardless of the firm's performance, it is

difficult for outside investors to determine whether or not the success or failure was driven by human discretion or purely algorithmic execution. This lack of clarity further highlights the necessity of a framework that ensures transparency for investors.

3.3.2 Alpha Generation vs. Benchmarks

Similar to Renaissance Technologies, D. E. Shaw demonstrates clear outperformance relative to the S&P 500, reaching a disparity of nearly 80 percentage points during 2025 Q4. However, it is difficult to confidently attribute D. E. Shaw & Co.'s success to either research-based alpha generation or human decision-making, as both elements are deeply integrated within its investment process.

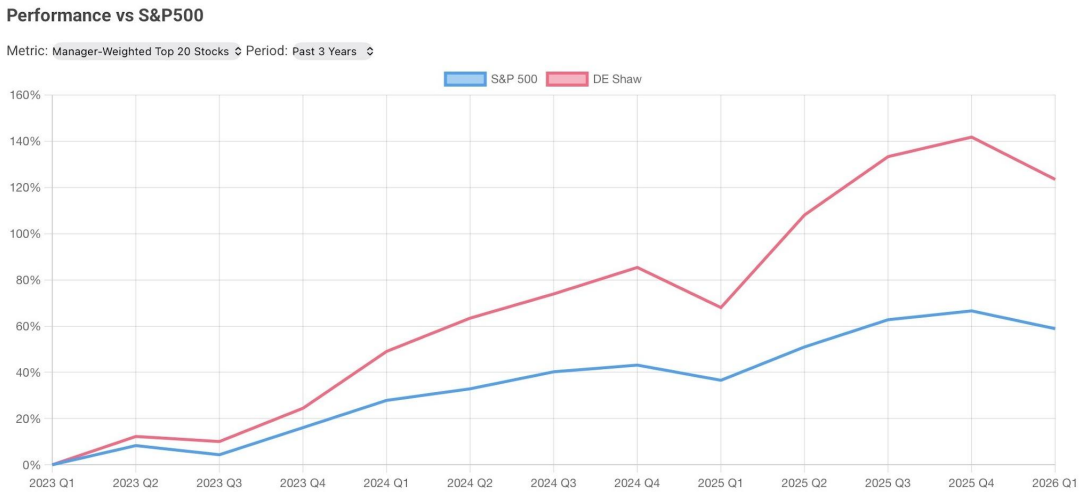

Figure 3. D. E. Shaw performance vs. S&P 500, manager-weighted top-20 holdings, Q1 2023–Q1 2026. Source: HedgeFollow n.d.-b.

3.4 Citadel

3.4.1 AI Disclosure and Strategies

Regarded as another prominent hedge fund, Citadel is a major investment firm that invests across global asset classes using a combination of quantitative models and human decision-making. Although the firm publicly talks about their investment strategies and risk-management practices, it provides limited information about the specific AI models and mechanisms used in its investment operations (Citadel n.d.). Similar to the firms mentioned above, Citadel provides little meaningful public disclosure of how AI is integrated into its trading and risk management operations. The information that is known publicly comes mostly from public statements by Citadel CEO Ken Griffin, which have become more and more contradictory. In July 2024, Griffin expressed skepticism that AI would replace most human jobs in the near future, emphasizing the limitations of machine learning models when adapting to changing conditions (Li 2024). However, by May 2026, Griffin changed his perspective. Speaking at Stanford, he acknowledged that AI had become much more capable, explaining how work that previously required teams of trained professionals over weeks or months could now be completed by AI agents in hours or days (Varanasi 2026). With the consistent progression of AI capabilities, there is more emphasis again on the necessity of a firm regulatory framework.

3.4.2 Alpha Generation vs. Benchmarks

The chart depicts a very similar performance between the S&P 500 and Citadel, with the largest gap being approximately 13 percentage points during 2025 Q4. Throughout this time horizon, the S&P 500 actually outperforms Citadel in the most recent quarter, 2026 Q1. This pattern is significant as it places Citadel much lower than the other quantitative firms mentioned when compared in regards to sheer performance. However, the results do not establish whether differences in performance are attributable to differences in technological usage, investment strategy, or other factors. This combination of observed performance and

limited public disclosure illuminates the difficulty of determining from publicly available data whether technological systems contribute to measurable alpha.

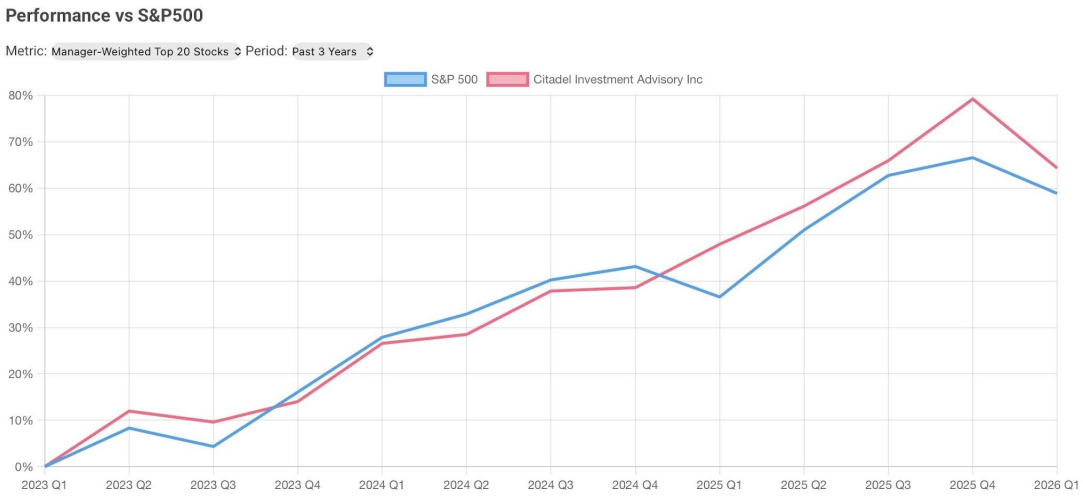

Figure 4. Citadel performance vs. S&P 500, manager-weighted top-20 holdings, Q1 2023–Q1 2026. Source: HedgeFollow n.d.-a.

3.5 Jane Street

3.5.1 AI Disclosure and Strategies

The final firm that will be analyzed in this paper is the Jane Street Group. Jane Street is a global proprietary trading firm founded in 2000 that specializes in quantitative trading and market making across equities, bonds, exchange-traded funds (ETFs), and derivatives. The firm uses mathematical models, machine learning, and large-scale data analysis to develop trading strategies and identify opportunities across financial markets (Jane Street n.d.). Unlike traditional asset managers, Jane Street trades primarily with its own capital rather than managing external client funds, and it plays a major role in providing liquidity across electronic markets (Barabanov 2026).

Jane Street has publicly acknowledged its use of machine learning in trading since at least 2017, including the utilization of deep learning and other machine learning techniques in live trading strategies (Jane Street 2017). Furthermore, the firm's blog mentions that even Jane Street's researchers acknowledge the complexity surrounding AI, even describing it as "a large, exciting and occasionally confusing area" (Jane Street 2017). Although the blog discusses model validation and precautions against costly prediction errors, it does not provide detailed information about the firm's specific risk limits, stress-response procedures, or oversight responsibilities.

In addition, Jane Street's scale makes its role in financial markets especially important. In the second quarter of 2025, the firm reported approximately $10.1 billion in net trading revenue, illustrating the size of its trading operations (Bloomberg News 2025). Because of Jane Street's significant trading activity across financial markets, potential weaknesses in automated trading systems raise concerns for parties beyond the firm itself. However, it is important to consider that the available evidence does not establish that Jane Street's models have behaved uncontrollably or that any resulting disruptions have affected broader financial markets. In another case, in July 2025, India's Securities and Exchange Board (SEBI) accused Jane Street of running an "intentional, well-planned and sinister scheme" of market manipulation, highlighting the regulatory examination surrounding the firm's trading activities (Balakrishner 2025). SEBI further noted that Jane Street had continued this even after receiving an explicit advisory from India's National Stock Exchange in February 2025, underscoring the danger of inadequate regulatory oversight (Balakrishner 2025). However, it is important to keep in mind that these allegations do not mean that Jane Street's AI systems behaved uncontrollably.

3.5.2 Alpha Generation vs. Benchmarks

Over the period analyzed, Jane Street's manager-weighted top 20 equity holdings generally showed higher cumulative returns than the S&P 500, with the gap between the two reaching approximately 40 percentage points in Q4 2025. This performance draws many parallels to Renaissance as both firms outperformed the benchmark while not fully disclosing how that success was generated. As a result, it is imperative to implement regulatory practices that mandate transparency while preserving the capacity for strong alpha generation.

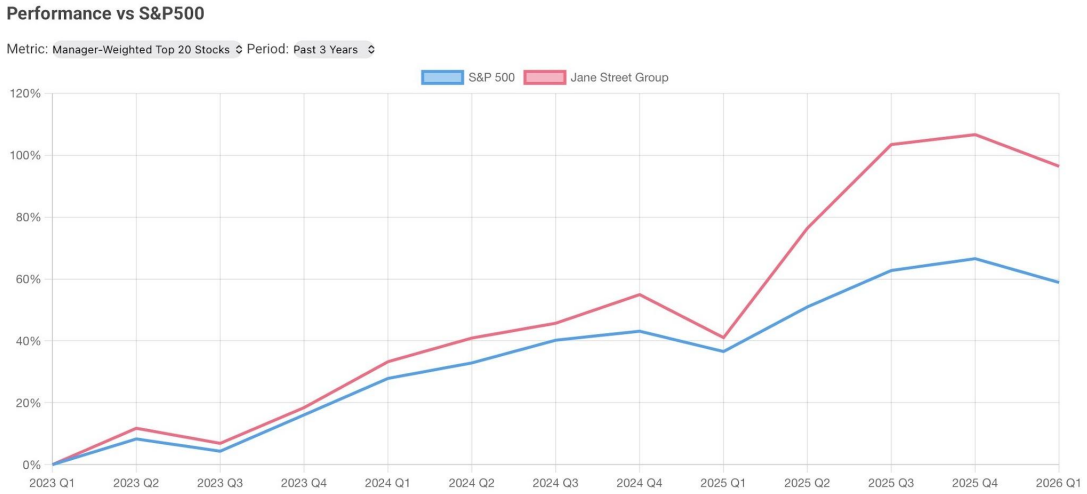

Figure 5. Jane Street performance vs. S&P 500, manager-weighted top-20 holdings, Q1 2023–Q1 2026. Source: HedgeFollow n.d.-c.

3.6 Comparative Analysis Summary

Across the five firms analyzed, a clear pattern emerges. Many of the firms demonstrated relatively strong performances compared to the S&P 500, yet each provides limited public disclosure about the specific AI systems, model governance structures, or risk constraints behind these performances. These observations raise questions about transparency, as the available disclosures provide limited insight into the relationship between quantitative investment methods, benchmark-relative returns, and the management of potential risks.

To better visualize this pattern across all firms, this paper utilizes two original figures generated by using information derived from HedgeFollow's manager-weighted top 20 holdings metric, and the disclosure scores represent a qualitative assessment based on a variety of public sources. To further elaborate on the qualitative assessment, the graph evaluates firm transparency through a disclosure rubric. The rubric contains ten categories with each category being scored as clearly disclosed, partially disclosed/implied, or not publicly disclosed from the sources reviewed.

The disclosure rubric in Figure 6 makes the pattern discussed evident. Across all five firms, public AI-related disclosure remains limited, with converted scores ranging from 5 out of 100 for Citadel to 35 out of 100 for Jane Street. These numbers emphasize that no firm achieved more than 35% of the available disclosure points, demonstrating that none of the firms provide a complete public AI governance framework. The most significant gaps appear in the categories closely associated with regulatory oversight: model governance, retraining procedures, risk limits, stress-response behavior, and incident-reporting protocols. Although these findings do not establish that the firms lack reasonable and adequate internal controls, they highlight the difficulty outside investors and regulators face when evaluating how they are governed.

The rubric also reveals that firms are generally more willing to acknowledge the use of AI, machine learning, or quantitative models than to explain how those systems are controlled. Renaissance Technologies, for example, receives credit for quantitative model use, but discloses little about data inputs, model architecture, human oversight, or crisis-response protocols. Citadel scores the lowest because its

public materials provide almost no structured explanation of how AI is integrated into trading or risk management. Two Sigma and Jane Street score highest in the sample because they provide more public discussion of machine learning and model use, yet both still fail to disclose several of the categories most important for evaluating model risk.

The cumulative return chart in Figure 7 reinforces the significance of these disclosure gaps over time. Renaissance Technologies, despite receiving only 15 out of 100 on the disclosure rubric, peaked at approximately 162% cumulative return in Q3 2025 before settling at roughly 136% by Q1 2026. D. E. Shaw, with a disclosure score of 25 out of 100, reached approximately 125% by Q1 2026. Jane Street, the highest-scoring firm on the rubric at 35 out of 100, reached approximately 96%. Meanwhile, Two Sigma, which scored 30 out of 100, produced the lowest cumulative return out of the group.

Together, Figures 6 and 7 show that strong benchmark-relative performance can coexist with limited public disclosure about AI governance and model-risk controls. This does not prove that AI directly caused the return gaps, nor does it prove that any firm's systems are unsafe. Instead, the figures suggest that current public disclosure is insufficient for outside observers to determine whether strong returns are produced within transparent, explainable, and stable risk boundaries. This supports the paper's broader argument that regulation should focus on functional transparency. Furthermore, without proper standards, there is a strong possibility that the gap between what these firms know about their own complex strategies and what regulators can assess will only widen as the technology advances.

Disclosure Category	Citadel	Renaissance Technologies	D.E. Shaw	Two Sigma	Jane Street
Acknowledges AI, ML, or quantitative model use	1	2	2	2	2
Explains specific AI/model use case in investing or trading	0	1	1	1	1
Identifies data input or training data categories	0	0	0	1	1
Identifies model architecture or model type	0	0	0	1	2
Discloses model governance or validation framework	0	0	1	0	1
Discloses retraining or update procedures	0	0	0	0	0
Discloses human oversight or decision attribution	0	0	1	1	0
Discloses risk limits or risk parameters	0	0	0	0	0
Discloses stress-response or kill-switch behavior	0	0	0	0	0
Discloses incident reporting or remediation protocols	0	0	0	0	0
Raw score (out of 20)	1/20	3/20	5/20	6/20	7/20
Converted disclosure score (out of 100)	5/100	15/100	25/100	30/100	35/100

Figure 6. AI Disclosure Rubric Across Five Leading Quantitative Firms (2 = clearly disclosed, 1 = partially disclosed or implied, 0 = not publicly disclosed. The rubric only uses publicly disclosed sources)

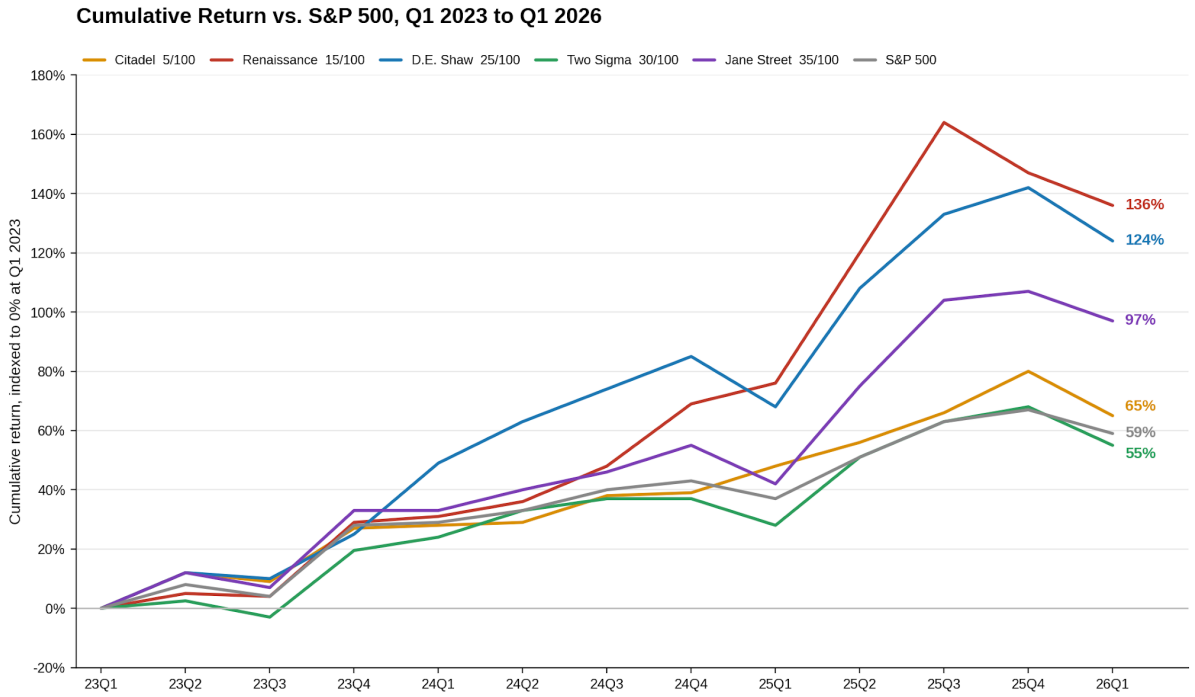

Figure 7. Cumulative Return of Publicly Disclosed Top-20 Holdings Compared with the S&P 500.

4. Discussion

The evidence from the five quantitative firms studied suggests that effective regulation of AI in portfolio construction must be both performance-based and disclosure-based. The purpose of regulation should not be to restrict quantitative innovation or punish firms for generating superior returns. Instead, regulation should ensure that when firms use AI to produce significant alpha, regulators can understand whether that outperformance comes from legitimate predictive skill, superior data infrastructure, or potentially hidden sources of systemic risk. This distinction is particularly relevant for firms such as Renaissance Technologies and D. E. Shaw, whose publicly disclosed equity holdings showed benchmark-relative returns that hint at the possibility of sophisticated quantitative investment methods. However, the available data cannot fully establish whether these methods, especially the AI-driven strategies, were responsible for the observed returns. Their limited public disclosures further limit the ability of outside researchers to evaluate the systems and risk parameters underlying their investment decisions.

This regulatory gap makes performance data incomplete. If a firm outperforms the broader market, that success may reflect model effectiveness; however, it may also reflect concentrated exposure, excessive leverage, or black-box decision-making that cannot be evaluated under current disclosure standards. Therefore, an empirical regulatory framework must examine not only which quantitative firms perform best, but also why they perform best, and whether those methods remain consistent with fiduciary duties, market stability, and investor protection.

Accordingly, this paper proposes four targeted reforms to address the disclosure gaps identified across the firms studied: mandatory AI disclosure standards, explainability and interpretability requirements, systemic risk monitoring, and human oversight attribution rules.

4.1 Mandatory AI Disclosure Standards

The most fundamental gap identified across all five firms in this study is the absence of any requirement to disclose how AI systems are actually built and governed. As established in Section 2.3, Form ADV—the SEC's primary registration mechanism for investment advisers—currently contains no requirement to disclose model architecture, training data sourcing, retraining protocols, or risk parameters.

This paper proposes that the SEC require quantitative investment advisers to file structured AI disclosure supplements alongside their Form ADV, detailing model categories and primary use cases, the general nature of training data sources, and the internal controls governing model deployment and modification. This proposal does not require firms to disclose full proprietary algorithms or competitive methodology—it requires them to document that governance structures exist and are functioning, a standard already applied to risk management frameworks in banking regulation.

4.2 Explainability and Interpretability Requirements

Beyond documenting that AI governance frameworks exist, regulators must be able to assess how material investment decisions are reached. As identified in Section 1.3, advanced neural networks are often highly complex and difficult to interpret even by the firms that build them. Jane Street's own technical blog acknowledges this directly, noting that AI is difficult to fully interpret and understand (Jane Street 2017). When a firm accounting for a meaningful portion of U.S. ETF trading recognizes this notion, the fiduciary and systemic implications extend well beyond its own investors (Wigglesworth 2024). This paper proposes that firms deploying black-box architectures in material investment decisions be required to demonstrate, upon regulatory request, how those decisions were reached.

4.3 Systemic Risk Monitoring for AI Concentration

The five firms studied in this paper compete in overlapping markets and potentially process many of the same data signals. If multiple institutions utilize similar types of algorithmic strategies simultaneously, their interactions could contribute to risks beyond the firm-level. The 2010 Flash Crash and the 2014 U.S. Treasury market disruption both demonstrated how automated trading systems can amplify instability when feedback loops develop across multiple firms responding to the same signals (Svetlova 2022). This paper proposes that the SEC, in coordination with the Financial Stability Oversight Council, establish cross-firm surveillance mechanisms requiring periodic reporting by quantitative firms of their model categories and primary signal types—not full proprietary algorithms, but sufficient disclosure to identify when multiple large institutions are running structurally similar strategies in correlated asset classes.

4.4 Human Oversight Attribution Rules

D. E. Shaw presents an important example of the challenges associated with evaluating quantitative and human-driven investment strategies (D. E. Shaw Group n.d.). Although the firm publicly discloses that it utilizes both systematic and discretionary aspects, its disclosures do not fully explain how responsibility is divided between models and human decision-makers within the individual strategies. This paper proposes that firms operating with investment strategies that contain automated models that function alongside human judgment should be required to clearly identify in their regulatory filings which categories of investment decisions are algorithmically driven, which are human-directed, and which involve a combination of both. These additional disclosures would help regulators establish accountability without having firms reveal their full trading strategies.

5. Future Directions

Although this paper provides a foundational framework for understanding the regulatory challenges of AI in portfolio management, future research should expand the study in different aspects of the issue.

More specifically, future analysis should examine a larger and more diverse group of financial institutions. This paper focuses on five major quantitative trading and investment firms as they represent some of the most influential and advanced players in the modern financial ecosystem; however, a broader sample including traditional asset managers and smaller hedge funds would help better determine whether the disclosure gaps identified in this paper are unique to the most prominent quantitative firms or if they are an accurate representation of the wider industry.

Second, future research should incorporate more data to obtain a more comprehensive outlook on the situation. This paper heavily relies on a single publicly available metric derived from HedgeFollow's

manager-weighted top 20 holdings approach. While this method provides a useful tool to compare visible equity performance against the S&P 500, it cannot fully capture the complexity of the strategies employed by quant firms. Many quantitative firms use short positions, complex derivatives, private investments, and other alternative assets that are not completely displayed through public holdings data. Future studies should therefore include broader measures such as Sharpe ratios, factor-adjusted alpha, and liquidity risk. These additional metrics would help distinguish AI outperformance and returns that may result from other hidden risk-taking.

Third, future research should develop a more standardized AI disclosure index. The disclosure scores used in this paper are based on qualitative analysis of publicly available sources, which is helpful for visualizing broad patterns, but remains susceptible to error due to the somewhat subjective nature of the evaluation. Developing a reproducible scoring rubric would make it easier for researchers, investors, and regulators to compare firms consistently over time.

Finally, future research should focus more directly on how AI systems behave during periods of market distress. One of the central concerns of this paper is that AI systems may appear effective under normal conditions but become unstable during volatility and other market shocks. Future studies should use historical stress events and other forms of simulations to evaluate whether AI-driven systems reduce exposure during crises or amplify instability through feedback loops. This would be especially vital for firms involved in market making and liquidity provision. In addition, a future endeavor should be the establishment of a standardized AI risk stress test that can be universally applied across major quantitative firms. This stress test would not require firms to reveal proprietary trading strategies, but it would require them to demonstrate how their systems respond to extreme but very realistic market conditions. As a result, there would be a consistent benchmark to evaluate whether AI systems are properly constrained.

Overall, future research should move beyond asking whether AI improves portfolio performance and instead investigate whether that performance is produced within transparent, explainable, and stable risk boundaries. By expanding the dataset, improving performance measurement, standardizing disclosure scoring, and developing a universal AI risk stress test, future work can help regulators design rules that preserve innovation while strengthening investor protection and market stability.

6. Conclusion

This study highlights a transparency gap between the advances of quantitative investment strategies and the information available to evaluate their use of AI. Across the five firms analyzed, limited public disclosure of model governance, risk controls, and oversight practices exist alongside periods of strong benchmark-relative performance. Although these findings do not directly establish that AI generated these returns or that firms lack adequate internal safeguards, they indicate that publicly available information alone may be insufficient to fully evaluate the governance and risk characteristics of advanced quantitative investment systems.

To address these concerns, this paper proposes four regulatory reforms: mandatory AI disclosure standards, explainability requirements, systemic risk monitoring, and human oversight attribution rules. Ultimately, stronger regulation should not restrict financial innovation, but ensure that as AI becomes increasingly utilized in investments, transparency and accountability develop alongside it.

Funding Sources
The author declares that there were no funding sources related to this work.

Conflicts of Interest
The author declares that there are no conflicts of interest related to this work.

Acknowledgements
The author would like to thank Mr. Wen for serving as an adviser throughout the research process. His guidance and feedback were essential for the refinement of this study.

AI Usage Acknowledgement
The author used Claude to generate a visual depicting the cumulative return chart. This use was limited only to display and presentation. The insertion of data was done by the author.

References

Balakrishner, A. 2025. "Jane Street Barred from Indian Markets as Regulator Freezes $566 Million over Nifty 50 Manipulation Claims." *CNBC*, July 4, 2025. https://www.cnbc.com/2025/07/04/indian-regulator-bars-us-trading-firm-jane-street-from-accessing-securities-market.html.

Barabanov, M. 2026. "Jane Street Explained: How Wall Street's Secretive Trading Firm Works." *Elite CurrenSea*, June 25, 2026. https://elitecurrensea.com/education/how-jane-street-works/.

Bloomberg News. 2025. "Jane Street's $10.1 Billion Trading Haul Sets Wall Street Record." *NDTV Profit*, September 2, 2025. https://www.ndtvprofit.com/markets/jane-streets-101-billion-trading-haul-sets-wall-street-record-9205305.

CFA Institute. 2019. "Flash Crashes." *CFA Institute Research and Policy Center*, October 29, 2019. https://rpc.cfainstitute.org/policy/positions/flash-crashes.

Cheung, KC. 2020. "Pioneer AI Hedge Fund—DE Shaw." *Algorithm-X Lab*, May 27, 2020. https://algorithmxlab.com/blog/pioneer-ai-hedge-fund-de-shaw/.

Citadel. "Global Quantitative Strategies." Accessed May 27, 2026. https://www.citadel.com/what-we-do/global-quantitative-strategies/.

Dakak, JP. 2025. "Investment Adviser Settles SEC Case on Model Security." *RIA Compliance Blog*, February 13, 2025. https://www.riacomplianceblog.com/investment-adviser-settles-sec-case-on-model-security/.

D. E. Shaw Group. "What We Do." Accessed April 29, 2026. https://www.deshaw.com/what-we-do.

Financial Stability Board. 2017. *Artificial Intelligence and Machine Learning in Financial Services: Market Developments and Financial Stability Implications*. November 1, 2017. https://www.fsb.org/uploads/P011117.pdf.

Ghatak, S., Khaledian, A., Parvini, N., and Khaledian, N. 2025. "Increase Alpha: Performance and Risk of an AI-Driven Trading Framework." *arXiv*, September 20, 2025. https://doi.org/10.48550/arXiv.2509.16707.

Government Office for Science. 2012. *The Future of Computer Trading in Financial Markets: An International Perspective*. Foresight, October 23, 2012. https://www.cftc.gov/sites/default/files/idc/groups/public/@aboutcftc/documents/file/tacfuturecomputertrading1012.pdf.

HedgeFollow. n.d.-a. "Citadel Advisors Fund Performance History & Annualized Return." Accessed June 12, 2026. https://hedgefollow.com/funds/Citadel+Advisors/Performance-History.

HedgeFollow. n.d.-b. "DE Shaw Fund Performance History & Annualized Return." Accessed June 12, 2026. https://hedgefollow.com/funds/DE+Shaw/Performance-History.

HedgeFollow. n.d.-c. "Jane Street Group Fund Performance History & Annualized Return." Accessed June 12, 2026. https://hedgefollow.com/funds/Jane+Street+Group/Performance-History.

HedgeFollow. n.d.-d. "Renaissance Technologies Fund Performance History & Annualized Return." Accessed June 12, 2026. https://hedgefollow.com/funds/Renaissance+Technologies/Performance-History.

HedgeFollow. n.d.-e. "Two Sigma Investments LP Fund Performance History & Annualized Return." Accessed June 12, 2026. https://hedgefollow.com/funds/Two+Sigma+Investments+Lp/Performance-History.

Jane Street. 2017. "Real World Machine Learning (Part 1)." *Jane Street Blog*, August 28, 2017. https://blog.janestreet.com/real-world-machine-learning-part-1/.

Jane Street. "Overview." Accessed June 8, 2026. https://www.janestreet.com/what-we-do/overview/.

Li, Y. 2024. "Citadel's Ken Griffin Says He's Not Convinced That AI Will Replace Human Jobs in the Near Future." *CNBC*, July 1, 2024. https://www.cnbc.com/2024/07/01/ken-griffin-says-hes-not-convinced-ai-will-replace-human-jobs-in-near-future.html.

Quantified Strategies. 2026. "The Culture of Secrecy: Why Renaissance Technologies Keeps Its Cards Close." July 28, 2026. https://www.quantifiedstrategies.com/the-culture-of-secrecy-why-renaissance-technologies-keeps-its-cards-close/.

Scrivner, D. 2025. "Renaissance Technologies: Breaking Down Renaissance's Quantitative Model, Returns, Medallion Fund Economics, Moat, Culture, and Origin Story. Plus Lessons Learned from Jim Simons." *Daniel Scrivner*, August 15, 2025. https://www.danielscrivner.com/renaissance-technologies-business-breakdown/.

Stanley, C. 2025. "AI Compliance: Applying Existing SEC Regulatory Frameworks to Fast-Moving Technologies." *Kitces.com*, November 19, 2025. https://www.kitces.com/blog/artificial-intelligence-compliance-considerations-investment-advisers-sec-securities-exchange-commission-legal-regulation-framework/.

Sutiene, K., Schwendner, P., Sipos, C., et al. 2024. "Enhancing Portfolio Management Using Artificial Intelligence: Literature Review." *Frontiers in Artificial Intelligence* 7: 1371502. https://doi.org/10.3389/frai.2024.1371502.

Svetlova, E. 2022. "AI Ethics and Systemic Risks in Finance." *AI and Ethics* 2 (4): 713–25. https://doi.org/10.1007/s43681-021-00129-1.

TrendSpider. n.d.-a. "D. E. Shaw & Company." *TrendSpider Learning Center*. Accessed April 29, 2026. https://trendspider.com/learning-center/d-e-shaw-company/.

TrendSpider. n.d.-b. "Renaissance Technologies." *TrendSpider Learning Center*. Accessed April 20, 2026. https://trendspider.com/learning-center/renaissance-technologies/.

TrendSpider. n.d.-c. "Two Sigma Investments." *TrendSpider Learning Center*. Accessed May 20, 2026. https://trendspider.com/learning-center/two-sigma-investments/.

Two Sigma. 2026. "AI in Investment Management: 2026 Outlook (Part II)." January 21, 2026. https://www.twosigma.com/articles/ai-in-investment-management-2026-outlook-part-ii/.

U.S. Securities and Exchange Commission. 2025. "SEC Charges Two Sigma for Failing to Address Known Vulnerabilities in Its Investment Models." Press release, January 16, 2025. https://www.sec.gov/newsroom/press-releases/2025-15.

Varanasi, L. 2026. "Citadel CEO Ken Griffin Was a Prominent AI Skeptic. Now He Says, 'AI Is Real.' " *Business Insider*, May 17, 2026. https://www.businessinsider.com/ken-griffin-citadel-ai-skeptic-changes-tune-hype-jobs-finance-2026-5.

Wigglesworth, R. 2024. "Jane Street Is Big. Like, Really, Really Big." *Financial Times*, April 29, 2024. https://www.ft.com/content/54671865-4c7f-4692-a879-867ef68f0bde.