

## Algorithmic stewardship: Institutional investors, artificial intelligence and systemic risk

Kabir Oyewale <sup>1,\*</sup> and Hosea Kipchumba <sup>2</sup>

<sup>1</sup> Department of Finance, Sanders College of Business, University of North Alabama, Florence, One Harrison Plaza, Florence, AL 35632, Alabama, United States.

<sup>2</sup> Department of Accounting, Willie A. Deese College of Business and Economics, North Carolina A & T State University, Greensboro, NC 27411, North Carolina, United States.

World Journal of Advanced Research and Reviews, 2026, 29(01), 934-941

Publication history: Received on 08 December 2025; revised on 12 January 2026; accepted on 15 January 2026

Article DOI: <https://doi.org/10.30574/wjarr.2026.29.1.0110>

### Abstract

Institutional investors have become the primary owners of public equities, fundamentally transforming corporate governance and market dynamics. This paper explores how the rise of artificial intelligence (AI) in investment management introduces new systemic risks and challenges traditional fiduciary duties. We define "algorithmic stewardship" as the governance of AI-driven decision-making within fiduciary institutions. Our framework connects investor constraints, AI decision rules, and market outcomes, highlighting that while AI can enhance efficiency and risk management, it may also synchronize behavior, amplify procyclical feedback loops, and obscure accountability. The paper discusses implications for regulators, suggesting the need for interaction-based oversight and AI-aware stress tests, as well as responsibilities for institutional investors. We conclude with future research directions on accounting disclosure and assurance in an AI-driven financial ecosystem.

**Keywords:** Institutional Investors; Algorithmic Stewardship; Systemic Risk; Fiduciary Duty; Artificial Intelligence

### 1. Introduction

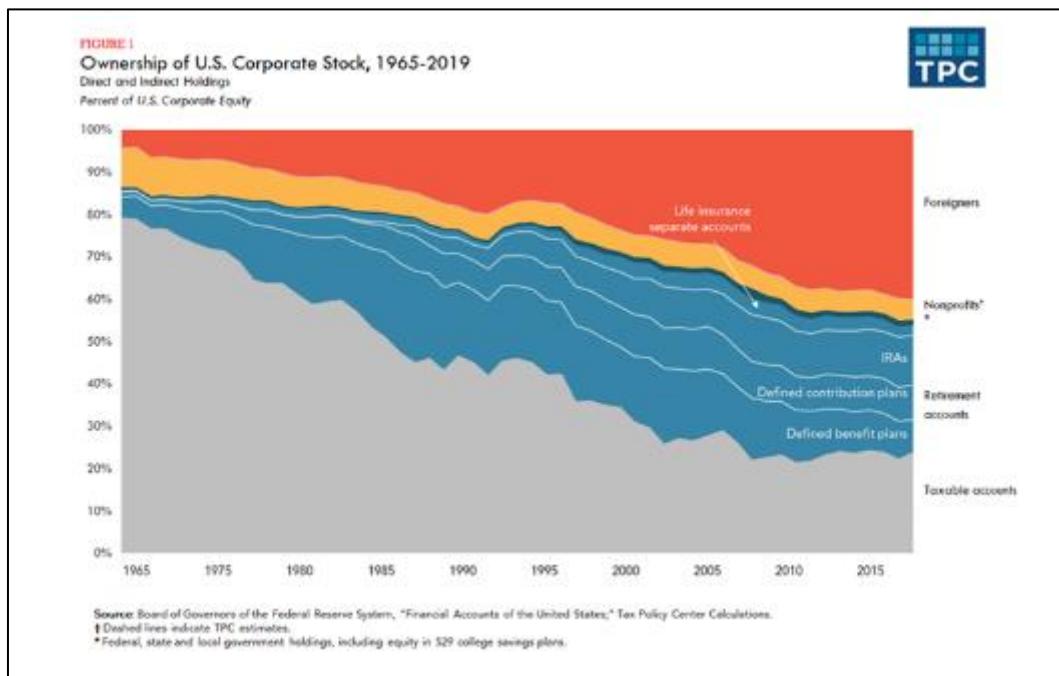
Over the past few decades, capital markets have undergone a dramatic transformation. In 1950, individuals directly owned approximately 90% of U.S. corporate equities; by 2015, institutional investors held 76% of public stocks. This concentration continues: the "Big Three" asset managers (BlackRock, Vanguard, State Street) are now the largest shareholders in 88% of S&P 500 companies, collectively controlling 20-25% of those firms' outstanding shares. This shift from dispersed individual ownership to concentrated institutional ownership fundamentally alters governance dynamics and market behavior.

**Table 1** Institutional Ownership Share of the U.S. Equity Market, 1950-2015

Year	Estimated Institutional Share (%)	Key Market Characteristic
1950	10	Individual investors dominate the market.
1980	29	Rising institutional presence amid the growth of pension and mutual funds.
2015	76	Institutional investors became the dominant market participants.

\* Corresponding author: Kabir Oyewale

This long-term realignment in ownership is further highlighted in Figure 1, which depicts the composition of U.S. corporate equity holdings by investor type from 1965 to 2019. This figure highlights the decline in direct household (taxable account) ownership, alongside the concurrent rise of various institutional categories, including retirement accounts, life insurance vehicles, and foreign holdings. By 2019, foreign investors accounted for approximately 40% of U.S. equity ownership, underscoring the globalized and institutional character of contemporary capital markets.



**Figure 1** Evolution of U.S. Corporate Equity Ownership by Investor Type, 1965–2019

This figure shows the distribution of U.S. corporate equity ownership across major investor categories from 1965 to 2019. Data are from the Federal Reserve's Financial Accounts of the United States, with calculations by the Tax Policy Center. Dashed lines denote Tax Policy Center estimates.

Institutional investors function as critical information intermediaries, governance actors, and market participants. They digest financial reports, exercise voting power, and direct capital allocation. Yet these stewards of capital now operate through an increasingly critical new channel: artificial intelligence and machine learning systems. Investment decisions are increasingly made by algorithms that predict returns, allocate assets, execute trades, monitor risks, and manage compliance at speeds and scales beyond human capability. This intersection of institutional dominance and AI-driven decision-making creates novel governance, and stability challenges that existing frameworks do not fully address.

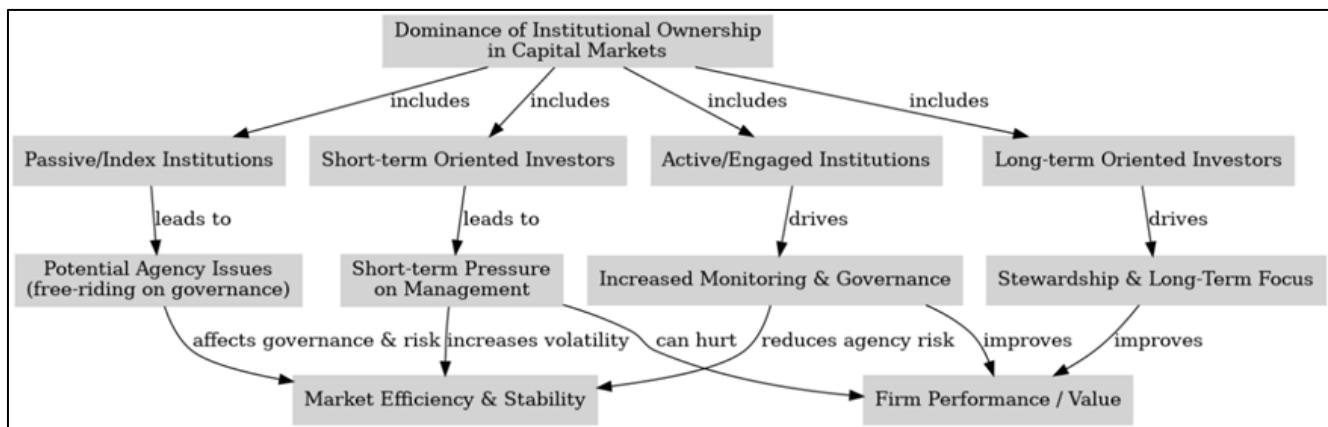
We introduce *algorithmic stewardship* to describe the governance of AI systems within fiduciary institutions. This is not simply adding another tool to the investor's toolkit; rather, AI has become an autonomous agent embodying institutional objectives and constraints. Our framework identifies three key dimensions of algorithmic stewardship: institutional characteristics (liquidity needs, leverage, horizon), primary AI function (alpha-seeking vs. risk-control), and degree of autonomy (human-in-the-loop vs. fully autonomous). These elements jointly determine how institutions behave in markets and what systemic risks emerge.

This paper contributes by: (1) extending institutional investor literature to incorporate AI as a central behavioral factor; (2) highlighting implications for accounting and disclosure when algorithms consume financial information; and (3) identifying new systemic risk transmission channels, such as synchronized trading and model homogeneity. Our analysis is conceptual and theoretical, synthesizing existing evidence to map a research agenda. We propose testable propositions and discuss regulatory implications.

## 2. Literature Review: Institutional Investors and AI in Finance

### 2.1. Institutional Investor Roles and Heterogeneity

Classical agency theory suggests that large shareholders mitigate principal-agent problems by monitoring management. Empirical evidence supports this: higher institutional ownership correlates with better governance outcomes and improved firm performance. However, institutional investors are heterogeneous. Bushee (1998) distinguishes between "transient" institutions with short horizons (generating pressure for short-term results) and "dedicated" institutions pursuing long-term strategies. Transient institutions may exacerbate short-termism and procyclical behavior during downturns, while stable long-term investors provide countercyclical liquidity. This heterogeneity is critical: institutional impacts depend on their specific constraints, incentives, and characteristics rather than being uniformly benign.



**Figure 2** Heterogeneity in Institutional Ownership and Its Impact on Market Outcomes

This diagram illustrates the diverse behaviors of institutional investors, ranging from passive and short-term actors to engaged and long-term stewards, and their differentiated effects on agency issues, governance, volatility, and firm value.

### 2.2. Institutional investors and systemic risk

Recent research shows that institutional investor behavior can propagate shocks across financial markets. De George et al. (2019) found that higher institutional ownership is positively associated with future systemic risk, particularly when dominated by transient institutions with common constraints (e.g., stop-loss rules, leverage limits). When many institutions face similar regulatory or contractual requirements, they may all react identically to market stress, producing systemic amplification even without explicit coordination. This suggests systemic risk increasingly emerges from behavioral synchronization rather than balance-sheet interconnections alone.

### 2.3. AI in financial decision-making

AI and machine learning now pervade financial activities: algorithmic trading, credit scoring, risk management, and quantitative investing. Benefits include improved efficiency and potentially more accurate predictions. However, significant vulnerabilities exist. Many AI systems (e.g., deep neural networks) operate as "black boxes" where decision rationales are opaque. This creates risks: if an institution experiences unexpected losses, diagnosis becomes difficult. Additionally, models often overfit to historical data and may fail in novel market conditions. Perhaps most critically for systemic risk, institutions adopting similar AI models or data sources may converge on comparable strategies, reducing diversification and increasing correlation. The May 2010 Flash Crash exemplifies how algorithmic feedback loops can trigger rapid market dislocations.

### 2.4. Fiduciary duty in an algorithmic era

Institutional investors operate under fiduciary duties to clients: duties of care (prudent, well-informed decisions) and loyalty (putting client interests first). AI challenges these duties: fiduciaries must ensure algorithms align with clients' best interests and maintain "algorithmic competence" by understanding how algorithms operate. They cannot fully abdicate responsibility to opaque systems. Additionally, while fiduciary duty traditionally focuses on client interests,

scholars increasingly question whether fiduciaries should account for systemic externalities, and uncontrolled algorithmic behavior could destabilize markets that clients ultimately depend on.

### 3. Conceptual framework: Algorithmic Stewardship

#### 3.1. Defining algorithmic stewardship

Algorithmic stewardship is the governance of AI systems within institutional investing. It treats AI-driven decision processes as integral components of institutional behaviour requiring guidance, monitoring, and, when necessary, restraint. Unlike purely technological views, we situate algorithms in an institutional context: the goals they execute, constraints they operate under, and feedback loops they create. AI encodes institutional incentives into decision rules, potentially amplifying their effects. For example, an institution mandating portfolio risk limits might program an AI to automatically reduce exposure once volatility exceeds a threshold. This operationalizes institutional constraints with precision and speed, making institutional behaviour a direct function of both the constraint and the algorithm's design.

#### 3.2. Core dimensions of algorithmic stewardship

We propose three key dimensions characterizing AI-augmented institutional investors:

- Intrinsic institutional characteristics: Structural features such as liability structure, liquidity needs, investment horizon, leverage, and regulatory constraints. AI intensifies these tendencies: liquidity-sensitive institutions become even more reactive when algorithms enforce constraints rapidly and rigidly.
- Primary AI function: Whether AI seeks alpha (excess returns through return prediction and arbitrage) or focuses on risk control and asset-liability optimization. Alpha-seeking systems may be opportunistic and potentially diverse in response; risk-control systems tend to be procyclical, all selling when volatility spikes.
- Degree of AI autonomy: The spectrum from human-in-the-loop (recommendations requiring human approval) to fully autonomous (automatic execution based on predefined parameters). Higher autonomy yields faster, more consistent responses but loses human judgment and discretion. Flash crash phenomena correlate with highly autonomous systems.

#### 3.3. Typology of algorithmic stewards

These dimensions yield distinct archetypes:

- Algorithmic herder (e.g., momentum or flow-driven hedge funds): Short-horizon, alpha-seeking, high autonomy. Under stress, likely to sell in sync, amplifying downturns.
- Synthetic liquidity provider (e.g., high-frequency trading firms): Risk-control focused, very high autonomy. Provides liquidity in calm periods; may withdraw abruptly under stress, creating air pockets.
- Risk-obsessed absorber (e.g., pension funds, insurers): Long-horizon, risk-control via AI, moderate autonomy. Generally countercyclical but may simultaneously de-risk with peers if key metrics are breached.
- Black box contrarian (e.g., quant hedge funds with deep learning): Flexible mandate, alpha-seeking, high autonomy, opaque strategy. Diversifies in normal times but risks unforeseen correlation in crises.

**Table 2** Typology of Algorithmic Stewards

Archetype	Investor Type	Financial Characteristics	AI Role / Primary Function	Systemic Risk Profile
Algorithmic Herder	Index funds, ETF sponsors	Short-term horizon, liquidity-driven, benchmark sensitivity	Momentum trading, flow forecasting, volatility targeting	High herding risk; synchronous de-risking in downturns
Synthetic Liquidity Provider	High-frequency traders, stat-arb hedge funds	High leverage, millisecond execution sensitivity	Order book modelling, arbitrage, and latency arbitrage	Volatility amplification; liquidity evaporation under stress
Risk-Obsessed Absorber	Pension funds, insurers, LDI managers	Long horizon, regulatory capital constrained	Liability-driven investment, hedging optimization	Stabilizing in calm periods; fragility in correlated unwind

Black Box Contrarian	AI-native hedge funds, quant strategies	Opaque mandates, high turnover, unobservable signals	Deep learning-based alpha generation, sentiment models	Unknown unknowns; correlation risk; contagion from model failure
----------------------	---	--	--	--

### 3.4. Core propositions

Proposition 1 – AI Intensification: AI adoption intensifies existing institutional tendencies by converting discretionary behavior into rule-based responses. Intrinsic characteristics like risk aversion or investment horizon have stronger, more immediate impacts when mediated by AI

Proposition 2 – Interaction effects: Systemic outcomes depend on interactions between institutional constraints and AI design. Neither factor alone determines systemic risk; a large leveraged fund with autonomous stop-loss systems poses greater risk than a small patient fund using AI.

Proposition 3 – Cybernetic Cliff Effects: When many institutions employ similar AI systems, markets exhibit threshold-based fragility: stable until critical tipping points, then subject to abrupt discontinuous adjustments. This reflects phase transitions inherent in complex networked systems.

---

## 4. Systemic Risk Implications of Algorithmic Stewardship

### 4.1. Synchronization Beyond Balance-Sheet Contagion

Traditional systemic risk models emphasize contagion via balance-sheet linkages or panic-induced runs. Algorithmic stewardship introduces a new mechanism: synchronization of decision rules across institutions. Because many institutional investors employ similar algorithms responding to the same signals, actions become synchronized even without direct financial connections. Consider a volatility spike: human traders might panic-sell, but others (contrarians) might buy, and many might hold—a range of behaviors that cushions shocks. With AI synchronization, if dozens of institutions have volatility-control algorithms triggering at the same threshold, all sell simultaneously. Volatility begets selling, which begets more volatility—a self-reinforcing spiral. This synchronization is algorithmic and behavioral, not due to financial exposure.

### 4.2. Procyclicality and Feedback Loops

Many AI systems, particularly for risk management, are inherently procyclical. Value-at-risk (VaR) models trigger selling when volatility rises, which increases volatility further. Algorithms using recent market data as inputs respond immediately without deliberation. An overnight shock could cause algorithmic trading at market open before humans intervene. The 2010 Flash Crash exemplified how feedback loops among algorithms create rapid price dislocations. AI can turn moderate shocks into severe dislocations through these feedback mechanisms. Markets may appear stable in benign periods (as algorithms smooth noise) but become fragile in extremes—characteristic of tightly coupled high-speed systems.

### 4.3. Homogenization of Risk Perception and Model Convergence

Markets benefit from diversity of opinion: when some investors see value where others see risk, they provide liquidity and balance. However, AI/ML models across firms are often trained on similar data (market prices, economic indicators) and use comparable architectures. This leads to convergence on similar signals. If prevailing models all learn that certain patterns predict downturns, they all de-risk simultaneously when those patterns emerge. The natural buyers disappear because everyone's AI "agrees" on risk. Markets lose the diversity of views vital for price discovery and liquidity. This homogenization paradoxically increases volatility because everyone acts together. If all algorithms sell, there are few contrarians to counteract it. This reduces market entropy and amplifies extreme outcomes.

### 4.4. Opacity, Model Risk, and Threshold Effects

Algorithmic stewardship increases opacity at the firm and system levels. Complex AI models often remain inscrutable even to firms' own risk managers. At the system level, with many opaque models interacting, the network of interdependencies (who will sell what when) becomes extremely difficult to map. This complicates crisis diagnosis: is a plunge driven by fundamental news (requiring one response) or self-reinforcing algos (requiring trading halts or liquidity injection)? Opacity itself becomes a systemic risk factor.

Additionally, many algorithms embed discrete decision rules: "trade normally until losses reach 5%, then liquidate half the portfolio." These threshold effects create "cliffs" where behavior changes discontinuously. If many institutions have similar thresholds, the market appears stable until crossing one critical point, then experiences a sudden cascade. Unlike continuous human behavior (investors trickling out as sentiment worsens), algorithmic cliff effects create flash instability. Markets may drop 2% without incident, then another 2% hits a threshold and triggers cascading algorithmic sales, producing a 10% plunge. Circuit breakers mitigate this by forcing timeouts, but they must match algorithmic trading speeds.

## 5. Implications for Regulation, Fiduciary Duty, and Accounting

### 5.1. Limitations of Entity-Based Regulation

Financial regulation traditionally assesses individual entities' safety, using capital requirements, liquidity ratios, and single-firm stress tests. Algorithmic stewardship reveals that individual stability does not ensure systemic stability. A large manager meeting all regulatory requirements and showing minimal balance-sheet risk might, with peers, create systemic events through synchronized algorithmic behavior. No single firm might be "systemically important" by size, but a cohort acting together could be. This suggests supplementing entity-based oversight with interaction-based oversight monitoring patterns across institutions: concentrations in model usage, correlation in triggers, etc.

Practically, regulators could require disclosures about algorithmic strategies in broad terms without revealing proprietary details. Institutions might report how they would respond to a 10% market drop; if all say, "we'd sell heavily," that signals systemic risk. Alternatively, regulators could identify and potentially regulate key third-party AI providers, treating them as systemically important nodes similar to how cloud providers are now supervised.

### 5.2. Systemic Risk Assessment and Stress Testing

Traditional stress tests ask "What if GDP falls X% or rates spike Y%?" Algorithmic stewardship requires adding scenario questions: "What if volatility doubles overnight and algorithmic triggers activate simultaneously?" Useful scenarios include:

- Simultaneous Risk Limit Activation: Volatility spikes; how many funds hit de-risk thresholds simultaneously? What aggregate selling would this imply?
- Algorithmic Liquidity Withdrawal: Market-making algorithms pull out as spreads widen. What would happen to trading volumes and bid-ask spreads?
- Correlated Model Error: All AI models, trained on recent data, underestimate a particular risk (e.g., geopolitical shock). All simultaneously misprice, leading to wrong-way bets.

These scenarios help identify transmission channels and system vulnerabilities. Recent Financial Stability Board (2024) recommendations explicitly call for incorporating AI-driven feedback scenarios in stress tests.

### 5.3. Fiduciary Duty and Algorithmic Competence

Prudence now arguably includes algorithmic competence: trustees and advisers should either understand their AI tools or hire experts who do. Fiduciaries cannot blindly rely on algorithms without understanding them—that would violate the duty of care. Best practice suggests documenting AI testing and validation as part of fiduciary processes. Investment committees should review AI strategy summaries at least annually. This parallel audit committee oversight of accounting includes oversight of algorithmic decision-making. Advisers should inform clients if substantial portfolio decisions are algorithm-driven and explain general approaches and risks.

Additionally, the duty of loyalty requires avoiding conflicts. Future conflicts might arise when AI routes trades to exchanges paying for order flow versus pursuing best execution for clients. Fiduciaries must ensure AI decisions don't inadvertently favor the manager's interests over clients.

### 5.4. Accounting Disclosure and Assurance

Current financial statements and risk disclosures do not typically cover decision processes. Perhaps new sections could address algorithmic strategies and governance in annual reports or regulatory filings. While firms might hesitate (fearing competitive exposure), the analogy is disclosure of accounting policies: firms routinely disclose revenue recognition or valuation methods without revealing proprietary data. Similarly, firms could disclose high-level

information about model governance or data dependencies. The goal is improving transparency so regulators and markets aren't blindsided by algorithmic risks building invisibly.

Accounting assurance might evolve to cover AI systems. Just as auditors opine on financial statement integrity and internal controls, future assurance might verify that institutions' AI models have appropriate controls and align with stated policies. This could mirror SOC reports for service organizations, with "AI SOC reports" covering model governance controls. For major institutions, an annual audit opinion might include a section: "We have assessed the key algorithms impacting portfolio decisions and believe they operate effectively under established controls."

### **5.5. Internal Governance and AI Risk Management**

Within institutions, governance structures must evolve. Traditional investment and risk committees may need technologists and data scientists. An "AI oversight subcommittee" of the board or management committee could review model performance, detect biases, assess incidents, and ensure resources for model risk management. Some firms are appointing Chief AI Officers, indicating board-level focus. Internal audit checklists must include AI: sampling decisions, backtesting outputs, and monitoring drift. Internal audit might also ensure alignment between algorithmic behavior and risk appetite, verifying that a fund claiming low risk hasn't secretly embedded tail risks in its algorithms.

---

## **6. Conclusion and Future Research Directions**

Algorithmic stewardship reframes how we understand financial markets and systemic risk. The stewards of capital, institutional investors, now include algorithms processing vast data and executing strategies in microseconds. This evolution offers both promise and peril: AI can augment human decision-making and operate efficiently, yet it introduces complexities challenging our current governance frameworks. As algorithms grow more sophisticated, regulatory frameworks, disclosure standards, and oversight must evolve in tandem.

### **6.1. Accounting Research Directions**

For accounting scholars, algorithmic stewardship opens several avenues: (1) How do algorithms process financial information differently than humans? Do they focus on quantitative data, ignoring narrative nuance? Does machine consumption of disclosures change their informational role? (2) How does accounting shape algorithmic design? Risk disclosures and capital charges influence algorithm optimization. Does enhanced disclosure correlate with reduced volatility or improved market stability? (3) New empirical questions: Do institutions using similar AI exhibit more correlated trading than those without? Does algorithmic risk management amplify volatility? How do disclosures influence algorithmic decisions? (4) Assurance implications: Do independent AI audits reduce perceived risk and lower institutional cost of capital? Can accounting frameworks guide AI governance standards?

### **6.2. Broader implications**

For regulators, this analysis underscores the urgency in adapting oversight to AI-mediated finance. Traditional entity-level tools remain necessary but insufficient. Regulators need interaction-based monitoring, AI-aware stress testing, and expertise to evaluate algorithmic systems. International coordination is critical as AI effects transcend borders. Organizations like the FSB are developing AI governance principles; our framework provides conceptual support for data sharing on AI incidents and governance standards development.

For practitioners, asset managers, pension trustees, and risk officers, the takeaway is clear: govern AI systems as fiduciary agents. Institutions managing AI thoughtfully can differentiate through reliability and trustworthiness. Those failing to invest in understanding and oversight risk unintended consequences harming clients and the system. Good algorithmic stewardship could become a competitive advantage and reputational capital. Practical steps include internal war games (modeling what happens if many institutions do what your AI does), contingency planning, and industry collaboration through forums developing standards and sharing insights on algorithmic anomalies.

### **6.3. Final thoughts**

The financial system's "brains" are being partly rewired toward AI. We must update the "nervous system", regulations, oversight, and disclosures, to ensure accountability and ethical standards keep pace. Whether algorithmic stewardship ultimately enhances or undermines financial stability depends on choices by fiduciaries, regulators, and academics. By bringing transparency to algorithmic processes and assuring information and control integrity, accounting can align AI-driven finance with principles of accountability and trust underlying well-functioning markets. The task ahead is translating these conceptual insights into concrete research and practical actions guiding stewardship evolution in the age of algorithms.

## Compliance with ethical standards

### *Disclosure of conflict of interest*

We declare that there are no financial, professional, or personal relationships with any organizations, institutions, or commercial entities mentioned in this manuscript that could influence or bias the research presented. Neither author has received funding, grants, or compensation related to the subject matter of this paper, and there are no existing or potential conflicts of interest that could compromise the objectivity or integrity of this work.

---

## References

- [1] P. Berle, A. A., & Means, G. C. (1932). *The modern corporation and private property*. Macmillan.
- [2] Bushee, B. J. (1998). The influence of institutional investors on myopic R&D investment behavior. *The Accounting Review*, 73(3), 305–333.
- [3] Coval, J. D., & Stafford, E. (2007). Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics*, 86(2), 479–512.
- [4] De George, E. T., He, X., & Martineau, C. (2019). The role of institutional investors in propagating stale information. *Journal of Financial Economics*, 133(3), 713–738.
- [5] Financial Stability Board. (2017). Artificial intelligence and machine learning in financial services. FSB Report.
- [6] Financial Stability Board. (2024). AI and financial stability. FSB Report.
- [7] Grossman, S. J., & Hart, O. D. (1980). Takeover bids, the free-rider problem, and the theory of the corporation. *Bell Journal of Economics*, 11(1), 42–64.
- [8] Jensen, M. C., & Meckling, W. H. (1976). Theory of the firm: Managerial behavior, agency costs, and ownership structure. *Journal of Financial Economics*, 3(4), 305–360.
- [9] Shleifer, A., & Vishny, R. W. (1986). Large shareholders and corporate control. *Journal of Political Economy*, 94(3), 461–488.