

Multi-Agent Decision Making for Liquidity Risk Management under Market Stress

Liwia Nagel

Department of Computer Science, George Mason University, Fairfax, VA, USA.

liwia.nagel186@gmu.edu

Abstract

Liquidity risk management under market stress is increasingly a problem of distributed, adaptive, and interdependent decision making rather than a narrow treasury or balance-sheet function. During periods of market dislocation, liquidity evaporates through feedback loops among funding markets, dealer inventories, collateral calls, algorithmic execution, and regulatory constraints. This paper examines how multi-agent decision making can be designed, governed, and deployed to improve liquidity risk management under such conditions. It argues that the central challenge is not simply predictive accuracy but the alignment of heterogeneous agents, incentive structures, communication protocols, and human oversight within a socio-technical infrastructure. The paper develops a system-level framework that distinguishes perception, coordination, execution, and governance layers. It analyzes structural trade-offs involving centralization and decentralization, speed and accountability, resilience and efficiency, and fairness and systemic stability. It further considers governance mechanisms, regulatory reporting, auditability, stress testing, operational resilience, and sustainability. Particular attention is given to the risks of correlated agent behavior, liquidity mirages, procyclical automation, and opaque model interactions. The paper concludes that multi-agent liquidity risk management requires hybrid architectures that combine adaptive machine intelligence with circuit breakers, human judgment, regulatory transparency, and cross-institutional coordination. Without such safeguards, multi-agent systems may amplify stress even as they improve local decision quality. The analysis has implications for financial institutions, regulators, infrastructure providers, and policy makers seeking to build robust, fair, and sustainable liquidity risk systems.

Keywords

multi-agent systems; liquidity risk; market stress; financial stability; governance; robustness; fairness; socio-technical infrastructure.

1. Introduction

Liquidity risk under market stress is among the most difficult problems in contemporary finance because it sits at the intersection of market microstructure, funding fragility, institutional behavior, and public policy. A market may appear deep and resilient during normal conditions, yet rapidly transform into a fragmented environment in which traders withdraw quotes, collateral becomes contentious, and funding horizons collapse. The classical account of bank runs explains how coordination failures can turn illiquidity into insolvency, particularly when liabilities are short term and assets are difficult to liquidate [1]. Subsequent work on contagion demonstrates that shocks can spread across institutions through direct exposures and common asset holdings, producing outcomes that no single institution fully controls [2]. Liquidity therefore cannot be treated as a static stock held by an isolated firm; it

is a systemic property generated by the interaction of balance sheets, market conventions, information asymmetries, and decision rules.

The rise of algorithmic trading, electronic market making, and automated treasury functions has made liquidity risk management more dependent on multi-agent decision systems. Market making agents continuously quote prices, hedge inventories, and adjust risk limits. Treasury agents manage cash buffers, collateral, and funding schedules. Risk agents monitor exposures, while compliance agents enforce regulatory ratios and internal mandates. External agents, including central counterparties, custodians, and regulators, shape the feasible action space through margin requirements, reporting obligations, and supervisory expectations. In such an environment, liquidity outcomes emerge from the interaction of many decision makers rather than from a single optimizer. Multi-agent systems offer a conceptual and computational vocabulary for modeling this distributed decision landscape [12]. They also introduce new failure modes, including correlated behavior, communication bottlenecks, strategic manipulation, and the diffusion of responsibility across automated components.

This paper argues that multi-agent decision making for liquidity risk management should be approached as a problem of architecture and governance, not merely model performance. A system that predicts liquidity stress accurately but coordinates agents poorly may still intensify a crisis. A system that executes quickly but cannot explain its actions may undermine supervisory trust and accountability. A system that centralizes control may achieve speed and consistency but create single points of failure. A system that decentralizes control may improve adaptability but produce coordination failures, duplicated risk taking, or hidden interdependencies. The central task is therefore to design socio-technical infrastructures that balance competing values: robustness, efficiency, fairness, sustainability, transparency, and legitimacy.

The paper proceeds as follows. Section 2 reviews conceptual foundations linking liquidity risk, market stress, information asymmetry, and systemic feedback. Section 3 develops a multi-agent architecture for liquidity risk management. Section 4 analyzes coordination, market impact, and systemic trade-offs. Section 5 examines governance, accountability, and regulatory infrastructure. Section 6 considers robustness, fairness, and sustainability. Section 7 discusses deployment, evaluation, and operational resilience. Section 8 explores policy implications and forward-looking perspectives. Section 9 concludes.

2. Conceptual Foundations of Liquidity Risk and Market Stress

Liquidity has two intertwined dimensions: funding liquidity and market liquidity. Funding liquidity refers to the ability of an institution to meet cash and collateral obligations as they come due. Market liquidity refers to the ability to trade assets quickly with limited price impact. These dimensions reinforce one another. When funding becomes scarce, institutions sell assets, which depresses prices, increases margin calls, and further tightens funding conditions [3]. The relationship is especially fragile when assets are used as collateral in short-term financing arrangements. In repo markets, for example, haircuts and collateral valuation can shift abruptly, converting a seemingly stable funding source into a run-prone structure [4]. Liquidity risk is therefore endogenous: the actions taken by institutions to protect themselves can degrade the liquidity available to others.

Information asymmetry intensifies these dynamics. In markets with heterogeneous information, lenders and traders may ration credit or withdraw from trading when they cannot distinguish sound assets from impaired ones [7]. The classic lemons problem explains why

quality uncertainty can cause markets to unravel, because buyers discount all assets when they cannot observe quality [8]. During stress, opacity about counterparty exposures, collateral quality, and model assumptions can produce sharp reductions in market participation. Even well-capitalized institutions may face funding difficulties if others cannot assess their solvency or if they are perceived to be exposed to common risks. The limits of arbitrage further explain why sophisticated investors may fail to correct mispricing when capital and risk-bearing capacity are constrained [6]. Liquidity provision requires balance-sheet capacity, not only information, and that capacity can disappear precisely when it is most needed.

Systemic risk research has emphasized that financial systems behave less like mechanical equilibrium models and more like ecological networks with feedback, contagion, and tipping points [5]. Stress can propagate through direct exposures, fire sales, common lenders, and correlated strategies. The increasing presence of automated agents adds another layer of complexity because these agents can react at speeds beyond human intervention. Agent-based computational finance has shown that heterogeneous, adaptive agents can generate realistic market phenomena, including bubbles, crashes, clustered volatility, and liquidity withdrawal [11]. Market microstructure research likewise demonstrates that order flow, inventory management, and information frictions shape how prices form and how liquidity is supplied [9]. The modern trading environment is not merely a venue for human decisions; it is a computational ecology in which latency, order types, and risk controls influence systemic outcomes [10]. Under market stress, these micro-level mechanisms can aggregate into macro-level fragility.

Regime shifts are central to this problem. Liquidity conditions are not stable across time; they switch between normal, stressed, and crisis regimes, often with little warning. Models trained in one regime may fail in another, and agents that optimize for historical regularities may become destabilizing when correlations and liquidity relationships change. Trust-calibrated multi-agent decision intelligence has been proposed as one response to these regime shifts, emphasizing the need for agents to assess the reliability of signals, peers, and models as conditions evolve [13]. The broader conceptual lesson is that liquidity risk management must incorporate adaptive trust, regime awareness, and coordination rather than relying on static correlations or fixed thresholds.

3. Multi-Agent Architecture for Liquidity Risk Management

A multi-agent architecture for liquidity risk management can be understood through four layers: perception, coordination, execution, and governance. The perception layer collects market data, funding data, collateral information, counterparty exposures, and operational signals. It may include internal systems, external market feeds, regulatory data, and alternative data sources. The coordination layer interprets these signals, forms shared situational awareness, negotiates actions among agents, and resolves conflicts among objectives. The execution layer translates decisions into trades, funding operations, collateral movements, and reporting actions. The governance layer establishes constraints, audit trails, escalation paths, and human oversight. This layered view helps prevent the mistake of treating multi-agent systems as simply many independent algorithms; their value and risk depend on how layers interact.

Agents within such an architecture may represent different institutional functions. A treasury agent manages cash and funding diversification. A market-making agent manages quotes and inventory. A risk agent monitors liquidity ratios, exposure concentrations, and stress indicators. A compliance agent enforces regulatory constraints and internal limits. A

counterparty agent evaluates the reliability of lenders, dealers, and clearing members. An external monitoring agent may represent a regulator, central counterparty, or custodian. In a multi-agent system, these entities can communicate, share representations, and coordinate actions [12]. However, coordination does not necessarily mean consensus. Agents may have conflicting objectives, incomplete information, and different time horizons. The architecture must therefore support negotiation, priority setting, and conflict resolution.

Artificial intelligence research provides useful perspectives on agent design, planning, learning, and interaction, but financial deployment raises distinctive concerns [14]. Financial agents operate in adversarial, non-stationary environments with high stakes, regulatory constraints, and systemic externalities. A machine learning agent may optimize for local reward while ignoring market impact or correlated behavior. A reinforcement learning agent may discover strategies that exploit market structure in ways that regulators consider manipulative. A large language model agent may generate plausible but unverifiable rationales for decisions. The architecture must therefore include constraint enforcement, uncertainty representation, and explainability mechanisms that go beyond narrow performance metrics. The goal is not to build autonomous traders but to build accountable decision support systems that augment human and institutional judgment.

Centralization and decentralization present a fundamental trade-off. A centralized coordinator can maintain a global view, enforce consistent risk limits, and prevent conflicting actions. It can also become a bottleneck, a single point of failure, and a source of correlated error. Decentralized agents can respond quickly to local conditions and adapt to changing environments. They may also duplicate exposures, compete for the same liquidity, or fail to internalize systemic costs. Practical architectures are likely hybrid. A central coordination layer may set boundaries, maintain a common operating picture, and manage emergency protocols, while decentralized agents handle local execution and adaptation. The design challenge is to determine which decisions require global consistency and which benefit from local discretion. During stress, the balance may shift toward centralization to prevent panic and coordination failure, but this shift must be preplanned and legitimate.

4. Coordination, Market Impact, and Systemic Trade-Offs

Coordination among liquidity risk agents is difficult because individually rational actions can produce collectively harmful outcomes. If multiple agents simultaneously reduce risk by selling assets, withdrawing quotes, or hoarding cash, they can deepen the stress that prompted their actions. This procyclical behavior is a recurring feature of financial crises. Fire sales depress prices, mark-to-market losses increase, margin requirements rise, and further selling follows. Multi-agent systems can amplify this dynamic if agents share similar models, data feeds, or risk thresholds. Correlated automation may create a liquidity mirage in which apparent depth disappears once several agents attempt to execute similar trades. The resulting market impact is not merely a cost of trading; it is a channel through which liquidity risk becomes systemic.

Market microstructure research shows that liquidity supply depends on inventory capacity, information asymmetry, and competition among liquidity providers [9]. High-frequency trading has changed the speed and structure of liquidity provision, sometimes improving narrow spreads while also introducing fragility when liquidity providers withdraw rapidly [10]. Automated execution can reduce search costs and improve efficiency, but it can also create fleeting liquidity that is available only under benign conditions [20]. Historical and sociological studies of high-frequency trading highlight how material infrastructure, latency,

and competitive dynamics shape market behavior in ways that are not fully captured by price-based analysis [17]. For multi-agent liquidity risk management, this means that execution agents must be designed with market impact models, circuit breakers, and coordination protocols that prevent self-reinforcing withdrawals.

Mechanism design offers a way to align incentives, but it must contend with complexity and incomplete information. Agents may not truthfully reveal their liquidity needs or risk exposures if doing so weakens their bargaining position. A central coordinator may lack the information needed to allocate liquidity efficiently. Regulatory constraints may prevent certain transfers or require separate legal entities to act independently. In this setting, coordination mechanisms should combine transparency, incentives, and fallback rules. For example, agents might share stress signals under predefined conditions, while retaining autonomy over proprietary strategies. Prearranged liquidity auctions, collateral transformation facilities, and coordinated circuit breakers can provide public infrastructure without requiring full centralization. The design question is how to create enough coordination to prevent panic while preserving enough decentralization to avoid brittle single-point control.

Fairness is also at stake in coordination. If large institutions receive preferential access to liquidity or if automated agents disadvantage certain market participants, the legitimacy of the system erodes. Information asymmetry and credit rationing already create distributional effects [7]. Multi-agent systems can exacerbate these effects if data, compute, and model quality are concentrated among a few firms. A fair architecture should include transparent access rules, nondiscriminatory execution, and oversight of how liquidity is allocated during stress. Fairness does not necessarily mean identical treatment, because risk profiles differ, but it does require that differential treatment be justified, auditable, and consistent with public interest. Without fairness, coordination mechanisms may be perceived as private bailouts or as tools for market power, undermining the trust needed for crisis response.

5. Governance, Accountability, and Regulatory Infrastructure

Governance is the foundation of trustworthy multi-agent liquidity risk management. Because liquidity risk is systemic, individual firms cannot be relied upon to internalize all social costs. Regulatory frameworks such as the Basel III liquidity standards establish quantitative buffers and monitoring tools, including the liquidity coverage ratio and net stable funding ratio [15]. These standards provide a baseline, but they were designed primarily for human-managed institutions and may not fully address the speed, opacity, and interaction effects of multi-agent systems. Governance must therefore extend beyond compliance checklists to include model risk management, change control, auditability, and supervisory technology. Regulators need the ability to inspect not only final positions but also decision logic, communication patterns, and escalation histories.

Accountability in multi-agent systems is complicated by distributed agency. If several agents contribute to a harmful outcome, who is responsible? The institution that deployed the agents, the developers who trained the models, the operators who set parameters, or the regulators who approved the system? Clear allocation of responsibility is essential for both legal liability and public trust. Governance frameworks should require documented authority, decision rights, and escalation protocols. Agents should operate within explicit limits and provide explanations that human supervisors can understand. Human oversight should not be a formality; it must include the capacity to intervene, override, or shut down systems during stress. Kill switches, circuit breakers, and manual fallback procedures are not technological afterthoughts but core governance instruments.

The Financial Stability Board has examined the implications of artificial intelligence and machine learning for financial services, highlighting benefits such as efficiency and risk assessment alongside risks of opacity, herding, and new dependencies [16]. These concerns are directly relevant to multi-agent liquidity risk management. If many institutions rely on similar cloud providers, data vendors, or foundation models, operational and strategic dependencies may become systemic. Governance must therefore address third-party risk, concentration risk, and cross-border coordination. Regulatory reporting should be timely, standardized, and machine-readable, enabling supervisors to detect correlated behavior and emerging liquidity strains. Supervisory sandboxes and public-private data collaboratives can help test governance mechanisms before they are deployed at scale.

Cross-border governance is particularly challenging because liquidity moves across jurisdictions while regulatory authority remains largely national. A multi-agent system may execute trades in one market, hold collateral in another, and rely on cloud infrastructure in a third. During stress, national authorities may prioritize domestic stability, creating conflicts with global coordination. Governance frameworks should therefore include international standards, mutual recognition of supervisory practices, and crisis communication protocols. They should also protect against regulatory arbitrage, in which firms deploy opaque or lightly regulated agents in jurisdictions with weaker oversight. The goal is not to impose a single global architecture but to ensure interoperability, transparency, and accountability across diverse legal and institutional settings.

6. Robustness, Fairness, and Sustainability under Stress

Robustness in multi-agent liquidity risk management requires more than redundancy. A system with many identical agents is not robust if they fail in the same way. Diversity of models, data sources, strategies, and communication channels can reduce correlated failure. However, diversity also complicates coordination and may reduce efficiency. The challenge is to design architectures that preserve enough diversity to avoid common-mode failure while maintaining enough shared understanding to coordinate during stress. Normal accident theory suggests that tightly coupled, complex systems are prone to unexpected interactions that cannot be fully predicted [18]. Financial systems are tightly coupled in some dimensions and loosely coupled in others, so robustness must be assessed through scenario analysis, stress testing, and adversarial simulation. Black swan events remind us that rare, high-impact shocks can invalidate assumptions and overwhelm safeguards [19]. A robust system should therefore fail gracefully, preserve essential functions, and allow for rapid recovery.

Fairness in liquidity risk management involves both procedural and distributive dimensions. Procedural fairness concerns access, transparency, and voice in governance. Distributive fairness concerns how liquidity, losses, and risks are allocated across institutions, sectors, and society. Multi-agent systems can improve fairness by applying consistent rules and reducing certain human biases, but they can also encode historical discrimination, reinforce market power, or create new forms of exclusion. If agents are trained on data shaped by past crises, they may learn to withdraw from certain counterparties or asset classes in ways that reproduce inequality. Governance must therefore include fairness audits, impact assessments, and mechanisms for appeal or correction. Fairness is not only an ethical requirement; it is a condition for the legitimacy and stability of crisis response.

Sustainability adds temporal and resource dimensions. Computational sustainability concerns the energy, hardware, and data infrastructure required to run large-scale multi-agent systems. Financial sustainability concerns the long-term viability of institutions and markets under

repeated stress. If agents optimize for short-term liquidity at the expense of long-term resilience, they may create boom-bust cycles that undermine the system they depend on. Sustainable architectures should internalize environmental costs, avoid excessive complexity that cannot be maintained, and support long-horizon risk management. They should also be adaptable to changing regulations, technologies, and market structures. Sustainability is thus not a separate concern but an integral part of robustness and governance.

7. Deployment, Evaluation, and Operational Resilience

Deployment of multi-agent liquidity risk systems raises practical questions about infrastructure, latency, cybersecurity, and change management. Some agents may run on-premises for latency-sensitive execution, while others rely on cloud services for scalable analytics. Hybrid deployment can improve flexibility but also creates dependencies, data transfer risks, and inconsistent security postures. Operational resilience requires that critical functions continue during outages, cyberattacks, or infrastructure failures. This means designing for degraded modes, maintaining manual fallbacks, and testing recovery procedures under realistic conditions. Human factors are central: operators must understand system behavior, recognize anomalies, and intervene effectively under pressure. Complex automation can create complacency or confusion, especially when multiple agents provide conflicting recommendations [18].

Evaluation must go beyond backtesting and historical simulation. Backtests may overfit to past regimes and underestimate market impact. Multi-agent systems should be evaluated through stress scenarios, adversarial simulations, and digital twin environments that model counterparties, market microstructure, and regulatory responses. Evaluation should include coordination performance, not only individual agent accuracy. Can agents avoid correlated withdrawal? Can they share information without creating herd behavior? Can they preserve liquidity while respecting risk limits? Evaluation should also consider fairness outcomes, operational costs, and environmental impacts. Regime shift detection is essential, because agents that perform well in normal conditions may fail when correlations change [13]. Continuous monitoring, drift detection, and periodic retraining are necessary but must be governed to avoid instability.

Operational resilience also depends on clear escalation and communication. During stress, agents may generate many alerts and recommendations. Human supervisors need prioritization, context, and explanations. A well-designed system should distinguish between local anomalies and systemic signals, between model uncertainty and data errors, and between reversible and irreversible actions. It should provide audit trails that allow post-crisis reconstruction of decisions. It should also support coordination with external authorities, such as central banks and regulators, without compromising security or confidentiality. Deployment is therefore not a one-time technical project but an ongoing socio-technical process involving training, governance, and adaptation.

8. Policy Implications and Forward-Looking Perspectives

Policy makers face a dilemma. Overregulation may stifle innovation and push activity into less transparent venues, while underregulation may allow multi-agent systems to amplify systemic risk. A balanced approach should focus on outcomes, capabilities, and governance rather than specific technologies. Regulators should require that liquidity risk systems be auditable, controllable, and interoperable with supervisory reporting. They should promote standards for model documentation, stress testing, and incident reporting. They should also

invest in supervisory capacity, including data science expertise and market surveillance tools. International coordination is essential to prevent regulatory gaps and to manage cross-border liquidity crises [16]. Public-private collaboration can support shared infrastructure for stress testing and data sharing while protecting proprietary information.

A forward-looking agenda should explore several research and policy questions. How can multi-agent systems be designed to internalize systemic risk without centralizing all decision authority? How can trust be calibrated dynamically across agents, models, and institutions? How can fairness be measured and enforced in automated liquidity allocation? How can operational resilience be tested against correlated cyber and market shocks? How can sustainability concerns be integrated into financial infrastructure planning? These questions require interdisciplinary collaboration among computer science, finance, economics, law, and organizational behavior. They also require experimental governance, such as sandboxes and pilot programs, that allows learning without exposing the entire system to uncontrolled risk.

The broader implication is that multi-agent decision making for liquidity risk management is a form of critical infrastructure. Like other critical infrastructures, it must be designed for reliability, security, equity, and public accountability. Efficiency alone is an insufficient objective. The architecture should support human judgment, regulatory oversight, and democratic legitimation. It should be robust to failure, fair in access and outcome, and sustainable in resource and institutional terms. The technical choices made today will shape the distribution of power and risk in future crises. A thoughtful approach can harness the benefits of distributed intelligence while protecting against its most dangerous emergent behaviors.

9. Conclusion

Liquidity risk management under market stress is fundamentally a multi-agent problem. Funding liquidity and market liquidity interact through feedback loops, information asymmetries, and correlated behavior. Multi-agent systems can improve situational awareness, coordination, and execution, but they can also create new fragilities through herding, opacity, and correlated automation. This paper has argued that the central challenge is architectural and governance-oriented rather than merely predictive. Effective systems must combine perception, coordination, execution, and governance layers; balance centralization and decentralization; align incentives; and preserve human oversight. They must be robust to regime shifts, fair in access and allocation, and sustainable in technical and institutional terms. Regulatory frameworks should evolve to address model risk, third-party dependency, cross-border coordination, and operational resilience. The required reference for this discussion emphasizes trust-calibrated multi-agent decision intelligence under regime shifts, but the wider literature confirms that trust, governance, and infrastructure are systemic concerns [13]. Ultimately, multi-agent liquidity risk management should be treated as critical socio-technical infrastructure. Its design should serve not only efficiency but also stability, legitimacy, and the public interest.

References

1. Diamond, D. W., & Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3), 401–419. <https://doi.org/10.1086/261155>
2. Allen, F., & Gale, D. (2000). Financial contagion. *Journal of Political Economy*, 108(1), 1–33. <https://doi.org/10.1086/262109>

3. Brunnermeier, M. K., & Pedersen, L. H. (2009). Market liquidity and funding liquidity. *The Review of Financial Studies*, 22(6), 2201–2238. <https://doi.org/10.1093/rfs/hhn098>
4. Gorton, G., & Metrick, A. (2012). Securitized banking and the run on repo. *Journal of Financial Economics*, 104(3), 425–451. <https://doi.org/10.1016/j.jfineco.2011.03.016>
5. Haldane, A. G., & May, R. M. (2011). Systemic risk in banking ecosystems. *Nature*, 469(7330), 351–355. <https://doi.org/10.1038/nature09659>
6. Shleifer, A., & Vishny, R. W. (1997). The limits of arbitrage. *The Journal of Finance*, 52(1), 35–55. <https://doi.org/10.1111/j.1540-6261.1997.tb03807.x>
7. Stiglitz, J. E., & Weiss, A. (1981). Credit rationing in markets with imperfect information. *The American Economic Review*, 71(3), 393–410.
8. Akerlof, G. A. (1970). The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3), 488–500. <https://doi.org/10.2307/1879431>
9. Kyle, A. S., & Obizhaeva, A. A. (2016). Market microstructure invariance: Empirical hypotheses. *Econometrica*, 84(4), 1345–1404. <https://doi.org/10.3982/ECTA10769>
10. O’Hara, M. (2015). High frequency market microstructure. *Journal of Financial Economics*, 116(2), 257–270. <https://doi.org/10.1016/j.jfineco.2015.01.003>
11. LeBaron, B. (2006). Agent-based computational finance. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of computational economics* (Vol. 2, pp. 1187–1233). Elsevier. [https://doi.org/10.1016/S1574-0021\(05\)02024-1](https://doi.org/10.1016/S1574-0021(05)02024-1)
12. Wooldridge, M. (2009). *An introduction to multiagent systems* (2nd ed.). Wiley.
13. Liu, T. (2026). Trust-Calibrated Multi-Agent Decision Intelligence for Financial Risk Management under Market Regime Shifts. Available at SSRN 7307503.
14. Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
15. Basel Committee on Banking Supervision. (2013). *Basel III: The liquidity coverage ratio and liquidity risk monitoring tools*. Bank for International Settlements.
16. Financial Stability Board. (2017). *Artificial intelligence and machine learning in financial services: Market developments and financial stability implications*. FSB.
17. MacKenzie, D. (2018). Material signals: A historical sociology of high-frequency trading. *American Journal of Sociology*, 123(6), 1635–1683. <https://doi.org/10.1086/696860>
18. Perrow, C. (1999). *Normal accidents: Living with high-risk technologies*. Princeton University Press.
19. Taleb, N. N. (2007). *The black swan: The impact of the highly improbable*. Random House.
20. Gomber, P., Arndt, B., Lutat, M., & Uhle, T. (2011). High-frequency trading. *Business & Information Systems Engineering*, 3(2), 109–112. <https://doi.org/10.1007/s12599-011-0141-0>