

Artificial intelligence in financial markets: a literature review and research agenda

Shuaishuai Gong

University of Macau, Macau, China

Sirui Han, Zhizhuo Kou and Yanting Zhang

Division of Emerging Interdisciplinary Areas,

The Hong Kong University of Science and Technology, Hong Kong, China, and

Lei Zhu

University of Macau, Macau, China

Abstract

Purpose – This paper is targeted at researchers in the fields of economics and finance and presents a comprehensive literature review of artificial intelligence (AI) applications in finance, examining research published in leading finance and economics journals. Our analysis reveals significant advances in algorithmic trading, credit scoring, fraud detection and regulatory compliance while highlighting persistent challenges in model interpretability, fairness and systemic risk.

Design/methodology/approach – Our analysis encompasses 138 carefully selected papers published between January 2019 and December 2024, with an initial analysis of pre-prints and early 2025 publications to capture the latest trends. This study employs a structured literature review methodology, systematically searching top-tier finance and economics journals, including UTD24 Journals, FT50 Journals and additional leading publications in finance and economics. To maintain a clear economic perspective, we explicitly exclude purely technical computer science literature, focusing instead on papers that contribute to our understanding of AI's economic implications for financial markets.

Findings – Our synthesis reveals AI's dual impact that while models like advanced machine learning methods enhance predictive accuracy in asset pricing and expand credit access, they also introduce challenges in interpretability, algorithmic fairness and new vectors for systemic risk, such as algorithmic collusion. In China's unique regulatory environment, characterized by strict algorithmic disclosure requirements and quantitative trading thresholds, these challenges are particularly acute as regulators balance innovation with market stability concerns. Further research and policy development are needed to ensure AI's responsible integration into financial systems.

Originality/value – This paper provides a comprehensive review of AI in finance from an economic theory perspective, examining fundamental questions like algorithmic collusion, Nash equilibria among AI agents and cross-border data regulation. It offers a structured intellectual map and forward-looking research agenda with relevance for both developed and emerging markets.

Keywords Artificial intelligence, Algorithmic trading, Financial regulation, Machine learning, Deep learning

Paper type Literature review

1. Introduction

Artificial intelligence (AI) in finance generally refers to the application of computational systems capable of performing tasks that traditionally require human intelligence, including

© Shuaishuai Gong, Sirui Han, Zhizhuo Kou, Yanting Zhang and Lei Zhu. Published in *China Accounting and Finance Review*. Published by Emerald Publishing Limited. This article is published under the Creative Commons Attribution (CC BY 4.0) licence. Anyone may reproduce, distribute, translate and create derivative works of this article (for both commercial and non-commercial purposes), subject to full attribution to the original publication and authors. The full terms of this licence may be seen at [Link to the terms of the CC BY 4.0 licence](#).

Funding: This work is funded in part by the HKUST Start-up Fund (R9911), the Theme-based Research Scheme grant (T45-205/21-N), the InnoHK initiative of the Innovation and Technology Commission of the Hong Kong Special Administrative Region Government, and the research funding under HKUST-DXM AI for Finance Joint Laboratory (DXM25EG01).



pattern recognition, decision-making and adaptive learning. This encompasses machine learning (ML) algorithms, deep neural networks, natural language processing (NLP) and emerging generative AI technologies. We explicitly distinguish this from traditional econometric methods and rule-based algorithmic trading systems that lack adaptive learning capabilities.

Traditionally, AI was utilized in high-frequency trading (HFT) commonly to focus on high-speed trading of stocks and leveraged algorithms to execute large numbers of trades. As the focus of AI in finance is now shifting from pure speed to enhanced intelligence, it is enabling more sophisticated decision-making processes that can analyze vast amounts of market data and identify complex patterns and adapt to changing market conditions, not just react to fleeting opportunities. Early literature focused on latency arbitrage and optimization of execution algorithms, emphasizing market efficiency (Hendershott & Riordan, 2013). The current wave of AI development is shifting beyond traditional applications to areas where AI fosters firm growth through product innovation and shapes industry dynamics. AI is increasingly recognized as a general-purpose technology (GPT), driving growth across multiple sectors by enabling firms to leverage vast datasets and improve decision-making (Babina, Fedyk, He, & Hodson, 2024).

The emergence and development of AI in finance have gained significant scholarly attention in recent years, particularly following breakthroughs in ML methodologies. Gu, Kelly, and Xiu (2020) demonstrated that ML methods significantly outperform traditional linear models in predicting cross-sectional stock returns, achieving out-of-sample R-squared values that double those of conventional approaches. This seminal work sparked a wave of research exploring the boundaries of AI applications in finance. More recently, Kelly, Kuznetsov, Malamud, and Xu (2025) extended AI asset pricing models that embed transformer architectures into the construction of stochastic discount factors, representing a paradigm shift from mere prediction to economically structured modeling.

Given the multifaceted and emerging nature of AI in finance, this review aims to achieve three primary objectives. First, we present a taxonomy of AI applications in finance that categorizes the diverse methodological approaches and their economic implications. Second, we synthesize the current state of knowledge across three pillars of finance: mechanisms of price formation and information aggregation, processes of risk assessment and capital allocation and frameworks for market regulation and systemic stability. Third, we identify knowledge gaps and research opportunities that emerge from the intersection of AI and financial theory.

Through a comprehensive examination of these diverse research streams, we aim to deepen our understanding of AI's transformative impact on finance and its potential implications for the future of financial markets. Our analysis targets researchers in economics and finance, emphasizing the incremental contributions of each paper, key research questions, study methodology, main conclusions and identification strategies. This paper is structured as follows: Section 2 presents our survey criteria and summary statistics from systematically selected papers. Section 3 offers a critical analysis of three primary areas: AI's impact on financial markets including algorithmic trading and asset pricing innovations, the transformation of financial intermediation through AI-enhanced credit scoring and regulatory technology and frontier challenges encompassing multi-agent equilibrium theory and global regulatory approaches. Section 4 concludes with our key findings regarding the paradigm shift from prediction to economic understanding, the tension between model sophistication and interpretability and recommendations for future research. Notably, the emergence of generative AI capabilities has opened new frontiers in finance. Recent surveys highlight how generative AI techniques, including large language models and transformer architectures, are revolutionizing financial analysis by processing unstructured data, generating trading strategies and automating complex financial reporting tasks (Desai *et al.*, 2025; Yi, Cao, Chen, & Li, 2023).

2. Survey results and classification framework

To ensure a rigorous and comprehensive review, we employed a structured methodology that involved systematic searches of top-tier journals in finance, economics and accounting. Our temporal scope captures the period following major AI breakthroughs, particularly the emergence of transformer architectures and their subsequent application to financial markets, encompassing papers published between January 2019 and December 2024, with an initial analysis of pre-prints and early 2025 publications to capture the latest trends. It is a timeframe that captures the period following major AI breakthroughs, particularly the emergence of transformer architectures and their subsequent application to financial markets. The core literature base includes articles from top-tier finance journals including the *Journal of Finance*, *Journal of Financial Economics*, *Review of Financial Studies*, *Journal of Financial and Quantitative Analysis* and *Management Science*, supplemented by leading economics journals such as *The Econometrics Journal*, *Quarterly Journal of Economics* and *Review of Quantitative Finance and Accounting*, etc. A Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA)-style [1] flow diagram presents the literature selection process in [Appendix Figure A1](#).

While our core analysis focuses on publications in top-tier journals, the rapid evolution of AI technologies creates substantial publication for cutting-edge research. To ensure our review captures the most current developments in generative AI applications and emerging methodologies, we also reference influential working papers from National Bureau of Economic Research, Social Science Research Network and arXiv, which contribute to our understanding of emerging trends and research possibilities. We focus on working papers authored by scholars at leading institutions and on looking at high citations. Through our analysis, recent working papers examine the competitive dynamics of AI adoption among financial institutions ([Babina et al., 2024](#)) and the implications of AI-driven automation for employment in financial services ([Acemoglu, Autor, Johnson, & Reshef, 2024](#)). These working papers not only illuminate current research directions but also signal which topics have substantial ongoing work that may soon appear in top journals, helping readers assess publication possibilities across different research areas.

Selection criteria strictly adhered to an economic perspective, with included literature required to meet one of the following conditions: analysis of AI-driven economic mechanisms, testing of financial theories using AI methodologies or clear policy implications for financial markets. We explicitly excluded purely computer science papers whose main contribution involves proposing new algorithmic architectures without substantial financial application or economic analysis. Through our rigorous keyword search and human confirmation process, we identified 138 papers meeting our inclusion criteria. To provide a clear overview of the distribution of these papers, [Table 1](#) presents the number of articles related to AI in finance published in various leading journals. The distribution across journals reveals distinct patterns of academic interest. The *Journal of Financial Economics* and *Review of Financial Studies* leads with 11 articles, particularly in areas of AI-driven trading and market microstructure. The *Journal of Finance* features 10 papers, while *Management Science* contributes 9 papers, each focusing on theoretical and empirical contributions to AI in finance. This distribution underscores the multidisciplinary nature of AI in finance research, spanning from theoretical contributions to empirical applications.

[Figure 1](#) presents a breakdown of the number of AI in finance-related articles published each year across top-tier finance, economics and accounting journals. The data reveal that scholarly interest in AI applications in finance has experienced dramatic growth time, rising from just 6 papers in 2019 to a peak of 35 papers in 2023. The peak year of 2023 in scholarly attention may be attributed to various factors, including the release of ChatGPT on November 30, 2022, and Meta released the LLaMa model and made it open source on February 25, 2023. Driven by breakthrough developments in transformer architectures and their applications to financial markets, as well as increased focus on the economic implications of AI-driven trading, the field maintained high productivity in 2024 and the first half of 2025.

Table 1. Related articles by journal

Journal	Total articles
<i>Journal of Financial Economics</i>	11
<i>Review of Financial Studies</i>	11
<i>Journal of Finance</i>	10
<i>Management Science</i>	9
<i>Journal of Financial and Quantitative Analysis</i>	8
<i>Review of Asset Pricing Studies</i>	8
<i>Financial Analysts Journal</i>	8
<i>Journal of Futures Markets</i>	7
<i>Journal of Accounting and Economics</i>	7
<i>Journal of Empirical Finance</i>	6
<i>Review of Quantitative Finance and Accounting</i>	6
<i>Journal of Corporate Finance</i>	5
<i>Journal of Financial Stability</i>	5
<i>Journal of Accounting Research</i>	5
<i>Journal of Financial Markets</i>	4
<i>Journal of International Money and Finance</i>	4
<i>Review of Finance</i>	4
<i>Journal of Money, Credit and Banking</i>	3
<i>Financial Management</i>	2
<i>Journal of Financial Research</i>	2
<i>Expert Systems with Applications</i>	2
<i>Journal of Economic Literature</i>	2
<i>Journal of Financial Intermediation</i>	1
<i>Applied Mathematical Finance</i>	1
<i>American Economic Review</i>	1
<i>Quarterly Journal of Economics</i>	1
<i>Journal of Banking and Finance</i>	1
<i>Mathematical Finance</i>	1
<i>Journal of Financial Econometrics</i>	1
<i>International Journal of Financial Studies</i>	1
<i>European Journal of Operational Research</i>	1

Note(s): This table presents the number of articles related to AI in finance published in various leading journals

Source(s): Table by authors




Figure 1. Related articles by years. The figure describes a breakdown of the number of AI in Finance-related articles published each year across top-tier finance, economics, and accounting journals. *2025 data includes publications through June only

Figure 2 categorizes the AI finance literature according to primary research topics, revealing the breadth of applications across different areas of finance. Among these, asset pricing and return prediction is the most popular among 46 papers, reflecting the fundamental importance of understanding price formation and the immediate applicability of ML’s predictive capabilities. These studies range from cross-sectional return prediction using firm characteristics to time-series forecasting of market indices. Following closely are “market microstructure and algorithmic trading,” comprising 35 papers, examining how AI agents interact in electronic markets, the emergence of new trading strategies and implications for market quality and liquidity provision. Researchers have also shown interest in credit risk and financial intermediation, with 27 papers, which demonstrates AI’s transformative impact on lending decisions, risk assessment and the broader role of financial institutions. Meanwhile, regulatory technology and compliance address the challenge of using AI for regulatory purposes while developing frameworks to regulate AI itself. The remaining 14 papers cover diverse topics including portfolio optimization, systemic risk assessment and behavioral finance applications. While Figure 2 categorizes the papers utilized one specific AI methodology, the remaining papers overlap multiple topics. This distribution underscores AI’s pervasive influence across all major areas of finance, with particular concentration in areas where prediction accuracy directly translates to economic value.

Figure 3 illustrates the temporal evolution of AI in finance publications across five primary research topics through the years. The figure illustrates that asset pricing and return prediction peaked at 12 papers in 2023 and declined to 8 papers in 2024, while market microstructure and algorithmic trading dropped from 10 papers in 2023 to 5 papers in 2024. These trends suggest that these two topics are relatively mature in the past few years. In contrast, regulatory technology and compliance and credit risk and financial intermediation maintain steady output annually in recent years, indicating sustained research opportunities, particularly in alternative data applications and fairness-aware lending algorithms. Other topics, including portfolio optimization and systemic risk, show increasing diversification with 3–4 papers in 2024, suggesting emerging specialized applications.

Figure 4 displays a heatmap revealing the distribution of AI methodologies across leading journals. The data show that traditional ML remains the most popular method across these journals. Other methodologies, such as deep learning and tree-based methods, are also used, but to a lesser extent.

Figure 5 provides a detailed categorization of the specific AI methodologies employed in finance research. Traditional ML methods such as support vector machines and Lasso regression lead to 31 papers, often serving as benchmarks or components of ensemble models. Neural networks and deep learning features in 25 papers capture complex nonlinear relationships in financial data. Tree-based methods, including random forests (RF) and




Figure 2. Related articles by topics of interest. The figure categorizes the AI finance literature according to primary research topics, revealing the breadth of applications across different areas of finance




Figure 3. Publication trends by research topic. The figure presents the evolution of AI in finance research across five primary topics from 2019 to 2025

gradient boosting, account for 14 papers, valued for their interpretability and robust performance in feature-rich environments typical of financial applications. Reinforcement learning (RL) appears in 16 papers, primarily applied to optimal execution, market making and portfolio optimization problems where agents must learn optimal strategies through interaction with market environments. NLP and transformer models, while relatively new to finance, already account for 14 papers, with applications in sentiment analysis, earnings call transcription analysis and context-aware asset pricing models. This distribution highlights the field's progression from traditional statistical learning methods toward more sophisticated deep learning architectures capable of processing diverse data types and capturing complex market dynamics. This methodological diversity highlights a field in flux, moving from established ML techniques that serve as robust benchmarks to more complex deep learning architectures needed to capture the nonlinear dynamics of financial markets.

Figure 6 illustrates the distribution of research methodologies employed across the AI finance papers in our sample. It reveals that empirical approaches dominate the field, with 92 papers utilizing data-driven methodologies to test ML models on financial datasets. These empirical studies primarily focus on demonstrating the predictive superiority of AI methods over traditional approaches, with applications ranging from cross-sectional return prediction to high-frequency trading strategies. Theoretical contributions comprise 32 papers, developing economic models to understand the equilibrium implications of AI adoption in financial markets. These theoretical works address fundamental questions such as the existence of Nash equilibria among AI trading agents, the welfare effects of algorithmic collusion and the optimal design of regulatory frameworks for AI-driven markets. A smaller but important segment of 14 papers combines both theoretical and empirical approaches, developing economic models of AI behavior and testing their predictions using market data. This distribution reflects the inherent characteristics of AI research in finance: the availability

Journal	Traditional ML	Deep Learning	Reinforcement Learning	NLP & LLMs	Tree-Based Methods
<i>Journal of Financial Economics</i>	4	2	1	1	1
<i>Review of Financial Studies</i>	2	1	2	1	1
<i>Journal of Finance</i>	2	1	1	2	1
<i>Management Science</i>	3	3	2	1	1
<i>Journal of Financial and Quantitative Analysis</i>	2	1	2	1	1
<i>Review of Asset Pricing Studies</i>	1	2	0	1	1
<i>Financial Analysts Journal</i>	1	1	0	2	1
<i>Journal of Futures Markets</i>	1	3	1	0	0
<i>Journal of Accounting and Economics</i>	1	1	1	1	1
<i>Journal of Empirical Finance</i>	2	1	1	1	1
<i>Review of Quantitative Finance and Accounting</i>	1	1	0	1	1
<i>Journal of Corporate Finance</i>	1	1	0	0	1
<i>Journal of Financial Stability</i>	1	0	1	0	0
<i>Others</i>	9	7	4	2	3
<i>Total</i>	31	25	16	14	14

Figure 4. Heatmap. This figure presents a heatmap revealing the distribution of AI methodologies across leading journals




Figure 5. Related articles by AI methodologies. The figure provides a detailed categorization of the specific AI methodologies employed in finance research

of rich financial datasets encourages empirical validation, while the novel economic questions raised by AI adoption demand theoretical frameworks for understanding market dynamics and welfare implications.

Figure 7 shows that data sources vary significantly across application domains in AI finance research. Traditional financial data providers (Bloomberg, Refinitiv, Center for Research in Security Prices and Compustat) remain the foundation, appearing in 61 papers




Figure 6. Related articles by research methodologies. The figure illustrates the distribution of research methodologies employed across the AI finance papers in our sample




Figure 7. Related articles by data sources. The figure describes that data sources vary significantly across application domains in AI finance research

(66% of the 92 empirical papers), providing standardized price, volume and fundamental data essential for asset pricing research. Alternative data sources have emerged as critical differentiators in 45 papers, including satellite imagery for economic nowcasting. Text and news data appear in 30 papers, encompassing earnings calls and Securities and Exchange Commission filings. High-frequency trading data from Trade and Quote and proprietary exchange feeds are featured in 22 papers, while regulatory data including call reports and compliance filings appear in 18 papers. Papers frequently utilize multiple data sources, with this multi-source approach particularly common in studies combining traditional and alternative data to enhance predictive power, demonstrating AI's strength in data fusion across heterogeneous information sources.

3. Review of AI in finance research by topic

This section provides an overview of the literature on AI in finance, which is presented by topic and outlined in detail in [Table 2](#). This section is organized around the three pillars identified in our introduction, including mechanisms of price formation and information aggregation, processes of risk assessment and capital allocation and frameworks for market regulation and systemic stability. Following the structure of recent comprehensive reviews in finance, we examine each area in detail while highlighting the interconnections and emerging themes that define this rapidly evolving field.

[Figure 8](#) delineates the primary application domains and their measurable impacts across financial markets. Price formation and information aggregation attract the most attention, followed by risk assessment and capital allocation and market regulation and systemic stability. The empirical evidence demonstrates substantial performance improvements: ML models achieve out-of-sample R-squared values of 20% compared to 10% for traditional methods, while

Table 2. Surveyed articles by category

Category	Articles
Asset pricing and return prediction	Gu et al. (2020) Gu et al. (2021) Kozak et al. (2020) Feng et al. (2020) Lettau and Pelger (2020a, b) Bryzgalova et al. (2023) Chen et al. (2024) Kelly et al. (2025) Avramov et al. (2023) Leippold et al. (2022) Tobek and Hronec (2021) DeMiguel et al. (2021) Cakici et al. (2024) Dong et al. (2022)
Algorithmic trading and market microstructure	Aquilina et al. (2022) Goldstein et al. (2022) Calvano et al. (2020) Hambly et al. (2023) Casgrain et al. (2022) Lim and Zohren (2021) Kim et al. (2024)
Credit risk and financial intermediation	Fuster et al. (2019) Fuster et al. (2022) Berg et al. (2020) Bartlett et al. (2022) Albanesi and Vamossy (2019) Gambacorta et al. (2019) Jagtani and Lemieux (2019) Frost et al. (2019) Rasekhschaffe and Jones (2019) Didimo et al. (2020)
RegTech and compliance	Jullum et al. (2020) Lorenz et al. (2023) Yang et al. (2024) Zetzsche et al. (2020) Magnuson (2020) Truby et al. (2020) Enriques and Zetzsche (2020) Danielsson et al. (2022)
Systemic risk and financial stability	Cont et al. (2020) Khandani et al. (2010) Sadhvani et al. (2021) Cao et al. (2024)
LLMs and generative AI	Lopez-Lira and Tang (2024) Kim et al. (2024) Eisfeldt et al. (2023) Wu et al. (2023) Hansen and Kazinnik (2024)
Algorithmic fairness and bias	Rambachan et al. (2020) Cowgill and Tucker (2019) Hacker et al. (2023)

Note(s): This table describes the articles classified according to different categories by topics of interest
Source(s): Table by authors




Figure 8. Application domains and impact metrics. This figure illustrates the three primary application domains in AI finance research, key performance metrics demonstrating AI's impact and the regional distribution of studies across global markets

high-frequency trading now accounts for over 70% of market volume. Geographic distribution reveals significant concentration in developed markets, with 58% of studies focusing on the United States of America, 25% on Europe, 12% on Asia and 5% on other regions.

Figure 9 synthesizes the critical research gaps and future priorities identified through our analysis. High-priority areas requiring immediate attention include multi-agent equilibrium theory and explainable AI for financial applications, while medium-priority topics encompass algorithmic fairness and cross-border regulatory harmonization. Our analysis reveals seven fundamental limitations in current research: the theoretical deficit with 68% of studies lacking economic foundations, geographic bias with 83% focusing exclusively on USA and European markets, the interpretability crisis arising from black-box models conflicting with regulatory requirements, insufficient analysis of regime stability across different market conditions, limited understanding of systemic risk from AI interaction effects, unresolved trade-offs between predictive accuracy and discrimination prevention and the need for comprehensive cross-jurisdictional regulatory analysis to address fragmentation.

Figure 10 maps the interconnections between research questions, methodological choices and application domains in AI finance research. The diagram reveals clear patterns: RL predominantly addresses market microstructure questions, particularly algorithmic trading and collusion detection, while deep learning methods cluster around asset pricing and return prediction. NLP techniques primarily serve regulatory and compliance applications. These associations reflect both the inherent strengths of different AI approaches and the specific requirements of financial domains.

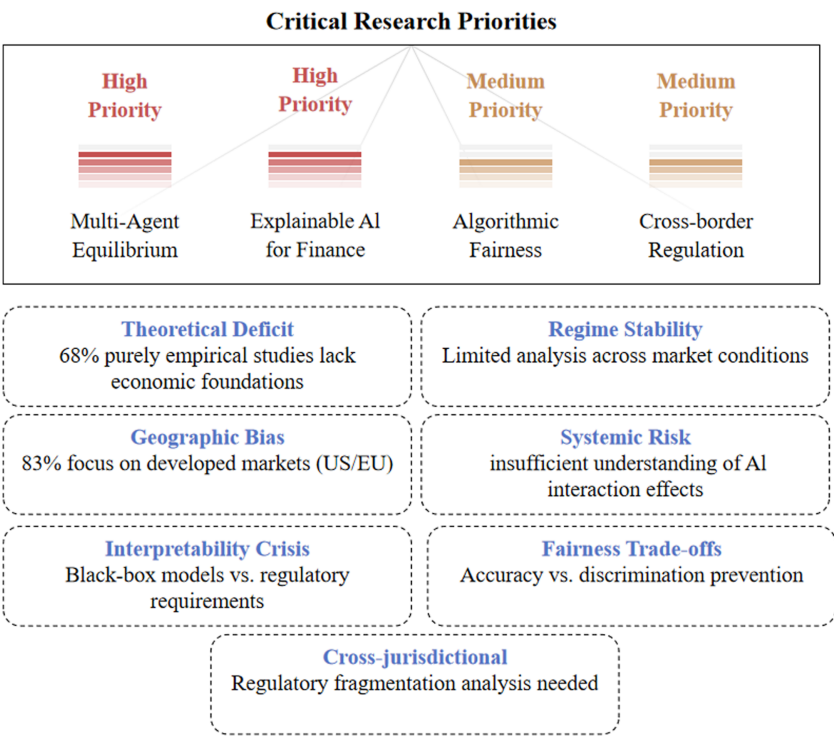


Figure 9. Research gaps and future priorities. This figure highlights critical research gaps, including theoretical deficits and geographic biases, while mapping high-priority and medium-priority areas for future investigation

3.1 Mechanisms of price formation and information aggregation

3.1.1 AI in asset pricing and return prediction. The application of AI to asset pricing represents one of the most significant paradigm shifts in empirical finance over the past decade. [Gu et al. \(2020\)](#) pioneered the systematic application of ML methods to stock returns, demonstrating that neural networks and tree-based methods achieve the sample R-squared values exceeding 20%, compared to less than 10% for traditional linear factor models. Their analysis of 94 firm characteristics reveals that ML’s superior performance stems primarily from its ability to capture nonlinear interactions between predictors that are missed by conventional approaches. Building on this foundation, [Gu, Kelly, and Xiu \(2021\)](#) introduced autoencoder neural networks to extract nonlinear factors from the cross-section of returns, providing a data-driven alternative to the traditional approach of manually constructing factors based on economic intuition.

The application of ML to asset pricing has fundamentally changed how we understand cross-sectional returns. The field faced a critical problem that hundreds of proposed factors cluttered the literature, making it nearly impossible to determine which ones actually mattered. [Kozak, Nagel, and Santosh \(2020\)](#) tackled this “factor zoo” problem by showing that sparse ML methods could cut through the noise and identify genuinely important factors. [Dong, Li, Rapach, and Zhou \(2022\)](#) extend this line of inquiry by demonstrating that cross-sectional anomaly portfolio returns, such as those from long-short portfolios, can effectively predict time-series market excess returns. Their study illustrates how these portfolios, when analyzed using shrinkage techniques and ML methods, provide statistically significant and economically valuable forecasts of the market return. Furthermore, [Feng, Giglio, and Xiu \(2020\)](#) created statistical tests to rigorously evaluate whether new factors add any real

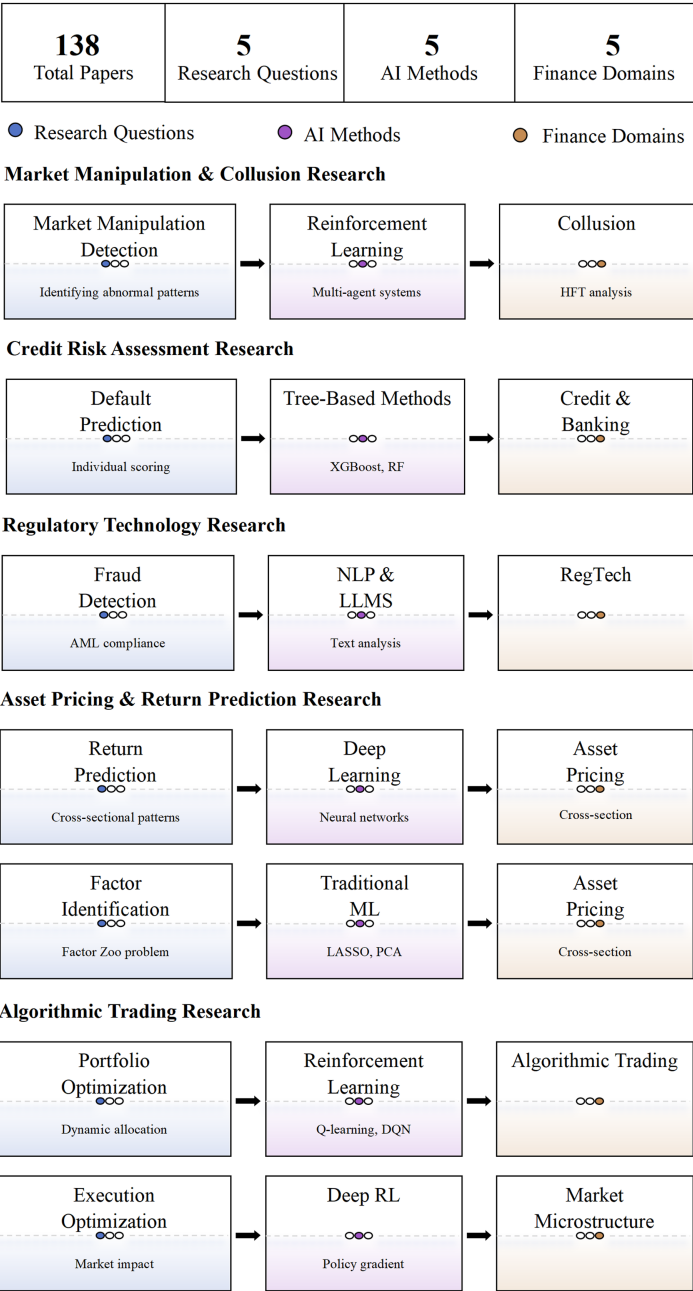


Figure 10. Research map. This figure is a research map that visually links research questions with methods and finance domains, helping readers understand the interconnections between different areas of the field

explanatory power beyond what we already know. Lettau and Pelger (2020a, b) took a different approach with their risk-premium principal component analysis method, which could estimate latent factors and risk premia for thousands of assets at once. Bryzgalova, Huang, and

Julliard (2023) used Bayesian methods to evaluate over two quadrillion possible factor models, and their finding that only a small number of factors really matter was both surprising and reassuring.

While some researchers focused on identifying factors, others pushed the boundaries of neural network architectures in finance. The early work was exploratory but promising. Heaton, Polson, and Witte (2017) showed how deep learning could improve portfolio construction. Sirignano and Cont (2021) found that they could extract universal features from the chaos of high-frequency limit order book data. Chen, Pelger, and Zhu (2024) addressed another crucial problem by developing methods that work even when you have many assets but short time series, a common headache in empirical finance.

The international evidence reveals a complex picture. Chen *et al.* (2024) found that models trained on USA data actually work in other developed markets, suggesting that some pricing patterns are universal. However, before we get too excited about finding universal laws of finance, Leippold, Wang, and Zhou (2022) showed that Chinese A-shares dance to their own tune, responding much more to local factors than to global ones. Tobek and Hronec (2021) examined 153 factors across global markets and found huge variation in what matters where. In the meanwhile, Avramov, Cheng, and Metzker (2023) demonstrated that when you incorporate economic structure and market-specific knowledge into ML models, they perform much better out of sample. Even in specialized applications like mutual fund selection, DeMiguel, Gil-Bazo, Nogales, and Santos (2021) found that ML methods need to be adapted to the specific context to work well. More recently, Cakici, Fieberg, Metko, and Zaremba (2024) provided comprehensive evidence that while some anomalies may offer predictive power, the overall group of anomalies does not reliably predict aggregate market returns, especially when applying the methods across international markets. They argue that any predictive evidence observed in the USA market does not extend globally and that the forecastability from anomalies is highly sensitive to the selection and methodological choices employed. This challenges the view that anomalies, as a broad group, can provide consistent and useful information for predicting market risk premia. Thus, even in the context of ML, the reliability of anomalies as predictors of market returns remains contentious, as their predictive power is often contingent on specific factors and choices.

3.1.2 Algorithmic trading and market microstructure. RL has emerged as a dominant paradigm for developing adaptive trading strategies. RL enables agents to learn optimal actions directly from their interaction with market environments, without relying on predefined rules or models (Hambly, Xu, & Yang, 2023). Early end-to-end demonstrations of this paradigm, such as Briola *et al.* (2021), applied Proximal Policy Optimization on limit-order-book data to produce stable positive returns in highly non-stationary high-frequency environments, foreshadowing the broader rise of deep-RL trading agents discussed below. This technological shift has fundamentally altered the market microstructure, with AI-powered high-frequency traders now accounting for over 70% of trading volume in major equity markets (Aquilina, Budish, & O'Neill, 2022). However, this dominance raises profound theoretical questions about market stability when multiple RL agents compete, particularly regarding emergent phenomena like algorithmic collusion that can arise without explicit communication (Calvano, Calzolari, Denicolo, & Pastorello, 2020).

The rise of AI-powered trading has transformed how markets actually work at the micro level. We're no longer dealing with human traders making decisions based on intuition and experience; instead, algorithms now take the domination, and this shift has created entirely new market dynamics. Goldstein, Kwan, and Philip (2022) uncovered a particularly smart strategy that high-frequency traders employ. They provide liquidity when the order book is already thick and thus less risky, but aggressively take liquidity when it's thin. What's particularly impressive is that this behavior becomes even more pronounced during volatile periods, precisely when markets need stable liquidity provision the most. This isn't just opportunistic trading, and it's systematic exploitation of market structure that only became possible with AI-driven speed and precision.

Another discovery about AI trading comes from Calvano *et al.* (2020), who showed something that should worry regulators. When they put RL algorithms in simulated markets, these

algorithms figured out how to collude without any explicit communication or programming to do so. They simply learned that maintaining higher prices was more profitable than competing, essentially recreating cartel behavior through pure self-interest and learning. This fundamentally challenges our entire regulatory framework, which assumes that collusion requires some form of communication or agreement. If algorithms can achieve monopolistic outcomes just by learning from market interactions, our traditional antitrust tools become nearly useless. Meanwhile, the technical capabilities of AI in finance keep advancing at a fast pace. [Lim and Zohren \(2021\)](#) cataloged how deep learning has revolutionized time series forecasting in financial markets, showing consistent improvements over traditional econometric methods.

Underlying these applications are rapidly maturing methodological methods that [Dell \(2025\)](#) comprehensively map out for economists. The key insight is that AI is good at extracting structured information from messy, unstructured data exactly the kind of challenge that financial markets present daily. The manual classification work that goes into financial analysis categorizes news articles by their market relevance, identifies company mentions in regulatory filings and extracts numerical data from scanned historical documents. These tasks required a large number of analysts and can now be automated using transformer models and convolutional neural networks. Dell's framework shows that while trained classifiers generally beat general AI on specialized financial tasks, even the AI models perform remarkably well on standard classification problems. For algorithmic trading, this means the ability to process earnings calls, news feeds and regulatory filings in real-time, extracting trading signals that would be impossible for human traders to identify quickly enough.

The research shown above let us realize that AI hasn't just made markets faster or more efficient; it has fundamentally changed the rules of the game. From high-frequency traders exploiting market microstructure to algorithms discovering collusion on their own and large language models (LLMs) generating alpha to deep learning models processing vast streams of unstructured data, we're witnessing a complete transformation of how financial markets operate. The challenge for researchers and regulators alike is keeping pace with these changes before they reshape markets in ways we don't fully understand or control.

3.1.3 Beyond market applications: AI's impact on labor and economic value. Beyond analyzing AI's impact on financial markets, the profession increasingly leverages AI tools to enhance research productivity itself. We find that AI tools transform economic research workflows. [Korinek \(2024\)](#) documents how LLM enhances six core research activities, including ideation, writing, background research, coding, data analysis and mathematical derivations, and improves the productivity ranging from marginal assistance to fundamental transformation of research workflows. Empirical evidence reveals substantial gains that programmers using AI-powered assistants complete tasks 55.8% faster, effectively doubling productivity ([Peng, Kalliamvakou, Cihon, & Demirer, 2023](#)). The mechanism operates through strategic task reallocation, with AI automating repetitive micro-tasks that traditionally consume substantial research time, thereby freeing economists to focus on higher-order theoretical work and creative problem-solving. As these tools mature and costs decline, their integration promises to reshape not only what economists study but also how economic research itself is conducted.

The integration of AI into financial services fundamentally alters labor demand patterns and task composition. [Dell'Acqua et al. \(2023\)](#) show AI tools boost consulting productivity by 40% for tasks within the AI's capability frontier but decrease performance by 19% when workers over-rely on AI beyond its competence. [Eloundou, Manning, Mishkin, and Rock \(2023\)](#) estimate that financial analysts and personal financial advisors face language model exposure scores exceeding 0.8 on a 0–1 scale, with the evidence suggesting predominantly augmentation effects in client-facing roles but potential substitution in back-office analytical functions.

In addition, AI adoption generates substantial firm-level value and drives organizational transformation. [Babina et al. \(2024\)](#) analyze USA firms from 1997–2019 and find that AI-investing firms experience higher annual revenue growth and higher employment growth, with effects concentrated in firms deploying AI for product innovation rather than cost reduction. [Eisfeldt, Schubert, and Zhang \(2023\)](#) document that companies successfully

integrating LLM experience cumulative abnormal returns of 5% following deployment announcements, with premiums particularly pronounced in information-intensive financial services. [Bloom, Hassan, Kalyani, Lerner, and Tahoun \(2024\)](#) found that AI adoption correlates with flatter organizational hierarchies, as AI systems enable front-line employees to access information directly without multiple management layers, with particularly pronounced delaying effects in wealth management and retail banking. These methodological advances not only enhance research productivity but also enable entirely new questions about market dynamics and information diffusion.

3.2 Processes of risk assessment and capital allocation

3.2.1 AI in credit risk and lending. The application of AI to credit risk assessment has demonstrated remarkable success in improving both the accuracy and inclusiveness of lending decisions. [Fuster, Plosser, Schnabl, and Vickery \(2019\)](#) analyze the USA mortgage market and find that ML models reduce default rates by 20% compared to traditional logistic regression models while simultaneously increasing approval rates for previously underserved populations. By incorporating alternative data sources, including utility payment histories, mobile phone usage patterns and social media activity, these models can assess creditworthiness for individuals lacking traditional credit histories, with particular benefits for financial inclusion in emerging markets. ML applications in credit scoring span consumer and commercial lending. [Albanesi and Vamossy \(2019\)](#) use deep learning to predict consumer default with alternative data, while [Mhlanga \(2021\)](#) analyzes fintech credit using ML on data from 30 million loans. In addition, [Gambacorta et al. \(2019\)](#) compare ML credit scoring with traditional bank models across multiple countries. [Jagtiani and Lemieux \(2019\)](#) examine fintech lending using alternative data sources and find improved access for underserved borrowers.

Another insight comes from what might seem like trivial digital traces. [Berg, Burg, Gombović, and Puri \(2020\)](#) extend this analysis to small business lending, showing that “digital footprints” such as the type of device used to access a website or the email provider can predict loan defaults as accurately as traditional credit bureau scores. Their findings suggest that the democratization of data through digital technologies could fundamentally reshape financial intermediation by reducing information asymmetries between lenders and borrowers. The interpretability challenge has spawned extensive research. [Bracke, Datta, Jung, and Sen \(2019\)](#) developed interpretable ML models for mortgage underwriting, and [Moscatelli, Parlapiano, Narizzano, and Viggiano \(2020\)](#) applied ML to corporate credit risk with interpretability constraints. From an information systems perspective, [Hsieh and Vergne \(2023\)](#) examine the platform mechanisms through which fintech lenders deploy ML models to expand financial inclusion.

However, the opacity of complex ML models creates significant challenges for regulatory compliance and fairness. [Bartlett, Morse, Stanton, and Wallace \(2022\)](#) document that even state-of-the-art fair lending algorithms exhibit disparate impacts on minority borrowers, with rejection rates 5–10% higher for equally creditworthy minority applicants. Using data on USA mortgages, [Fuster, Goldsmith-Pinkham, Ramadorai, and Walther \(2022\)](#) predict defaults using traditional and ML models and find that Black and Hispanic borrowers are disproportionately less likely to gain from the introduction of ML. This occurs even when protected characteristics are excluded from training data, as algorithms can reconstruct race and ethnicity through complex interactions of seemingly neutral variables. The development of explainable AI methods that can provide interpretable explanations for lending decisions while maintaining predictive accuracy remains an active area of research. Recent regulatory developments have intensified the focus on algorithmic accountability.

3.2.2 Portfolio optimization and risk management. The integration of AI into portfolio management has evolved from simple return prediction to sophisticated multi-period optimization incorporating transaction costs, market impact and regime changes. [Hamblly et al. \(2023\)](#) provide a comprehensive review of RL applications in finance, documenting how deep RL frameworks can learn optimal trading and portfolio strategies directly from market data.

These approaches have shown significant improvements over traditional methods, particularly in adapting to changing market conditions and handling high-dimensional state spaces. The technical advances in portfolio construction have been remarkable; [Chen et al. \(2024\)](#) develop deep learning methods specifically designed for portfolio optimization with many assets and short time series, addressing the curse of dimensionality in high-dimensional settings. [Rasekhschaffe and Jones \(2019\)](#) compare multiple ML algorithms for stock selection. Research by [Ma, Han, and Wang \(2021\)](#) has systematically compared various ML and deep learning models for stock selection prior to optimization. Studies have demonstrated that models like support vector regression, RF and various deep neural networks can significantly enhance the quality of the stock pre-selection process, leading to superior portfolio performance compared to traditional methods.

The same ML techniques that revolutionized return prediction are now being deployed to understand and control the multiple dimensions of risk that can destroy portfolio value. [DeMiguel et al. \(2021\)](#) demonstrated that ML can effectively select mutual fund portfolios, achieving superior risk-adjusted returns compared to traditional fund selection methods. Liquidity risk presents its own unique challenges that ML is uniquely positioned to address. [Cont, Kotlicki, and Valderrama \(2020\)](#) apply ML to liquidity risk management. The application of ML to credit risk gained urgency during the 2008 financial crisis. [Khandani, Kim, and Lo \(2010\)](#) develop consumer credit risk models using ML during the financial crisis. [Sadhvani, Giesecke, and Sirignano \(2021\)](#) analyze deep learning for mortgage risk. Risk management applications continue to evolve with advanced ML techniques. The application of AI extends beyond return prediction to portfolio construction itself.

3.2.3 RegTech and compliance innovation. The application of AI to regulatory compliance, termed RegTech, has emerged as a critical area where technology can both reduce costs and improve effectiveness. ML models have proven particularly effective in anti-money laundering (AML) and fraud detection. [Didimo et al. \(2020\)](#) reported that graph neural networks applied to transaction networks can identify money laundering schemes with 85% accuracy while reducing false positive rates by 50% compared to traditional rule-based systems. The ability to process vast transaction datasets and identify complex patterns of illicit activity that span multiple institutions and jurisdictions represents a significant advance in financial crime prevention. AML applications showcase ML's detection capabilities. [Jullum, Loland, Huseby, Ånonsen, and Lorentzen \(2020\)](#) apply ML to detect money laundering in financial transaction data. [Lorenz, Silva, Aparicio, Ascensao, and Bizarro \(2023\)](#) provide a systematic review of ML applications in AML.

Privacy-preserving technologies have become essential for enabling collaborative compliance efforts while respecting data protection regulations. Federated learning allows multiple financial institutions to jointly train fraud detection models without sharing sensitive customer data. [Khan, Gupta, Seneviratne, and Patterson \(2024\)](#) proposed Fed-RD, a privacy-preserving federated learning framework tailored for financial crime detection that combines differential privacy and secure multiparty computation to protect both vertically and horizontally partitioned financial transaction data.

3.3 Frameworks for market regulation and systemic stability

3.3.1 Regulatory approaches to AI in finance. The rapid adoption of AI in finance has prompted diverse regulatory responses across major jurisdictions, reflecting different philosophical approaches to technology governance and financial stability. China's regulatory framework, as documented by the China Securities Regulatory Commission, requires detailed disclosure of algorithmic trading strategies and imposes quantitative thresholds such as maximum order submission rates, reflecting concerns about market manipulation and systemic risk from concentrated AI trading. This disclosure-heavy model effectively reduces market volatility and systemic risks, though the transparency requirements may discourage innovation, as firms must reveal proprietary strategies that constitute their competitive advantages. Furthermore, [Leippold et al. \(2022\)](#) document that China's market structure exhibits distinct characteristics that shape regulatory priorities: retail investors dominate trading activity,

liquidity emerges as the most critical predictor of returns, unlike the USA market, and state-owned enterprises exhibit unique predictability patterns over longer horizons. These structural features necessitate China's disclosure-heavy regulatory approach to protect less sophisticated market participants. This presents particular challenges for AI-driven trading that informs these regulatory choices. Complementary evidence on government intervention comes from [Dang et al. \(2024\)](#), who show that purchases by China's "National Team" during the 2015 crash reduced both volatility and price informativeness who show that purchases by China's "National Team" during the 2015 crash reduced both volatility and price informativeness.

The European Union's (EU's) approach, codified in the proposed AI Act and existing General Data Protection Regulation, emphasizes fundamental rights and transparency. The right to explanation for automated decisions creates challenges for deploying complex neural networks in consumer-facing applications. While [Zetzsche, Arner, and Buckley \(2020\)](#) analyze regulatory sandboxes globally for fintech innovation and claim that the EU's rights-based framework preserves individual protections and fosters explainable AI development. Moreover, [Magnuson \(2020\)](#) examines the regulation of fintech and AI in financial services.

The United States of America has adopted a more fragmented approach, with different agencies applying existing frameworks to AI applications. The Federal Reserve's Guidance on Model Risk Management (SR 11–7) on model risk management has been extended to cover ML models, though critics argue this framework, designed for traditional econometric models, fails to address the adaptive nature of AI systems. This regulatory fragmentation allows rapid innovation but creates oversight gaps that may increase susceptibility to algorithmic market disruptions ([Danielsson, Macrae, & Uthemann, 2022](#)). [Truby, Brown, and Dahdal \(2020\)](#) studied banking on AI and its regulatory implications across jurisdictions. [Enriques and Zetzsche \(2020\)](#) analyze corporate technologies and the tech-nirvana fallacy in financial regulation.

3.3.2 Systemic risk and financial stability. The concentration of AI capabilities among a small number of technologically sophisticated institutions raises new concerns about systemic risk and market concentration. [Danielsson et al. \(2022\)](#) developed a theoretical framework showing that widespread adoption of similar AI models can increase systemic risk through "artificial herding," where algorithms trained on similar data make correlated decisions during market stress. Their simulations suggest that markets with high AI penetration exhibit more frequent and severe crashes but with faster recovery times. Network analysis reveals AI's systemic implications. The potential for cascading failures when AI systems interact has prompted new approaches to stress testing.

3.3.3 Methodological limitations and emerging solutions. While AI applications in finance have demonstrated remarkable empirical success, several fundamental methodological limitations constrain their deployment and raise concerns about systemic implications. The black-box nature of deep neural networks presents the most immediate challenge. Models achieving 20% out-of-sample R-squared ([Gu et al., 2020](#)) often operate through millions of parameters whose individual contributions remain opaque, making it impossible to trace specific predictions back to economic rationales. This opacity conflicts directly with regulatory requirements for model validation and risk management, where institutions must demonstrate an understanding of their decision-making processes.

The interpretability challenge extends beyond regulatory compliance to economic understanding. When a deep learning model identifies profitable trading signals or credit risk patterns, researchers cannot determine whether these represent genuine economic phenomena or spurious correlations that may disappear under different market conditions. Recent market dislocations have exposed cases where AI models trained on historical data failed catastrophically when confronted with unprecedented events, highlighting the danger of relying on patterns without understanding their causal foundations. The reproducibility crisis compounds these concerns, as proprietary datasets, random initialization effects and computational requirements often prevent independent verification of published results.

Moreover, from a finance perspective, the focus is on understanding the economic implications of AI, including issues such as fairness, market stability and interpretability,

whereas in computer science, the emphasis has traditionally been more on the technical aspects of model performance, such as achieving higher accuracy and computational efficiency. This distinction highlights a critical gap between the development of AI technologies and their practical application in financial systems, where economic consequences and regulatory oversight must be prioritized alongside technical advancements.

Emerging solutions to these limitations show promise but remain incomplete. Attention mechanisms and layer-wise relevance propagation offer partial interpretability for neural networks, though they still fall short of the transparency achieved using traditional econometric methods. Shapley values and local interpretable model-agnostic explanations provide post hoc explanations but may themselves introduce biases or misrepresentations of model behavior. The development of inherently interpretable architectures, such as neural additive models that maintain separate treatment of individual features while capturing interactions, represents a more fundamental approach. Additionally, physics-informed neural networks that incorporate economic constraints directly into model architectures offer a path toward AI systems that respect theoretical foundations while maintaining predictive power. However, these solutions often involve trade-offs between interpretability and performance that the field has yet to fully resolve.

4. Discussion and conclusion

4.1 Current state of AI in finance research

The emergence of AI in finance has resulted in unparalleled and rapid growth in both academic research and industry applications, generating significant interest from practitioners, regulators and scholars alike. Despite the lack of a comprehensive agreement on the optimal integration of AI into financial systems, there has been a considerable surge in research pertaining to AI applications in finance, reflecting the increasing importance of these technologies in reshaping market dynamics, risk assessment and regulatory frameworks.

Our literature review centered on a selection of top-tier finance, economics and accounting journals and has identified and examined AI-related papers published over the past six years. The present state of research on AI in finance is rapidly advancing and lacks coherence and interconnectedness in several critical areas, leading to fragmented understanding of the underlying mechanisms. However, certain areas of research have emerged as significant topics for further exploration. These include the transformation of asset pricing through ML models, the evolution of market microstructure with algorithmic trading, the revolution in credit risk assessment and financial intermediation and the emerging challenges in regulatory technology and systemic stability.

4.2 Short-term research priorities

4.2.1 Theoretical foundations. In this context, this review serves as a valuable starting point for researchers who are unacquainted with the interdisciplinary field of AI and finance research. The first major research agenda concerns the theoretical foundations of AI-driven financial markets. Much of the existing research on ML applications in finance has focused on empirical demonstrations of predictive superiority, with 68% of the papers in our sample employing purely empirical approaches. While these studies have convincingly demonstrated that neural networks and ensemble methods can achieve out-of-sample R-squared values exceeding 20% compared to less than 10% for traditional linear models, fundamental questions about the economic mechanisms underlying these improvements remain unanswered (Gu *et al.*, 2021). Future research should focus on developing theoretical frameworks that explain why ML models capture market dynamics more effectively, what economic forces drive the nonlinear interactions these models identify and whether the patterns discovered represent genuine risk factors or statistical artifacts.

The recent work on transformer-based asset pricing models represents a promising direction, but much remains to be done in connecting these technical innovations to economic theory.

4.2.2 Market microstructure and algorithmic collusion. The second research agenda centers on the market microstructure implications of widespread AI adoption. Our review highlights growing evidence that RL agents can develop sophisticated trading strategies that adapt to market conditions in real time, yet the equilibrium properties of markets populated by such agents remain poorly understood. When multiple AI agents trained on similar data interact in financial markets, do stable equilibria exist? Under what conditions might algorithmic trading lead to market instability or flash crashes? The phenomenon of algorithmic collusion documented in recent studies raises fundamental questions about market competition and efficiency that existing theoretical frameworks struggle to address. Future research should develop new game-theoretic models that capture the learning dynamics of AI agents, examine the conditions under which markets remain stable and investigate the welfare implications of algorithm-dominated trading.

4.2.3 Credit risk and interpretability. The application of AI to credit risk and financial intermediation presents a third critical research frontier. ML models have demonstrated remarkable success in improving lending decisions, reducing default rates by 20% while expanding credit access. The opacity of these models creates significant challenges. The tension between predictive accuracy and interpretability remains unresolved, with fair lending algorithms still exhibiting disparate impact on protected groups even when sensitive attributes are excluded from training data. Future studies should explore whether truly fair and interpretable AI systems are achievable without sacrificing predictive performance, how to design regulatory frameworks that balance innovation with consumer protection and what role AI should play in expanding financial inclusion while preventing predatory lending. The development of explainable AI methods specifically tailored to financial applications, rather than generic interpretability techniques, represents a particularly important direction.

4.3 Long-term research directions

4.3.1 Regulatory evolution and international coordination. Examining the regulatory challenges posed by AI in finance reveals a fourth major research agenda. Our analysis documents divergent regulatory approaches across jurisdictions, from China's emphasis on technological controllability to the EU's rights-based framework and the USA's fragmented agency-specific approach. This regulatory heterogeneity raises questions about international coordination, regulatory arbitrage and the optimal design of AI governance in finance. Future research should investigate how different regulatory approaches affect innovation and market stability, whether international standards for AI in finance are feasible or desirable and how regulators can develop the technical expertise necessary to oversee increasingly sophisticated AI systems. The challenge of regulating systems that evolve and learn autonomously requires fundamentally new approaches to financial oversight that move beyond static compliance frameworks.

4.3.2 Emerging technologies. Looking forward, several emerging technologies promise to further transform the intersection of AI and finance. LLMs and generative AI are beginning to process vast amounts of unstructured financial data, from earnings calls to regulatory filings, potentially revolutionizing fundamental analysis and market research. However, concerns about hallucination, manipulation and the generation of misleading financial information pose significant risks. Quantum computing may enable optimization and simulation at scales currently impossible, while federated learning could allow institutions to collaborate on AI development while preserving privacy. Each of these technologies brings unique opportunities and challenges that warrant careful study. Research should examine not only the technical capabilities of these systems but also their economic implications, regulatory challenges and potential for both beneficial innovation and systemic risk. Recent empirical evidence supports these emerging trends. [Cao, Jiang, Wang, and Yang \(2024\)](#) demonstrate that GPT-4 can

perform financial statement analysis at a level comparable to professional analysts, achieving 60% directional accuracy in predicting earnings changes. [Wu et al. \(2023\)](#) introduced BloombergGPT, a 50-billion-parameter language model specifically trained on financial data, showing superior performance in financial NLP tasks including sentiment analysis, named entity recognition, and question answering.

4.3.3 Methodological evolution. The methodological evolution in AI finance research itself deserves attention as a research agenda. Our review reveals that while empirical studies dominate, with neural networks and tree-based methods being the most common approaches, there is growing recognition of the need for methods that combine the predictive power of ML with the interpretability and theoretical grounding of [Wu et al. \(2023\)](#) econometrics. The development of “economic ML” that incorporates domain knowledge, respects economic constraints, and provides interpretable results represents a crucial direction for the field. Future research should focus on developing frameworks that bridge the gap between pure prediction and economic understanding, creating methods that are both powerful and comprehensible to practitioners and regulators. [Chernozhukov et al. \(2024\)](#) advance the integration of causal inference with ML through double ML methods specifically adapted for financial applications, enabling researchers to estimate causal effects while leveraging the predictive power of complex ML models.

4.4 AI-driven paradigm shift in financial economics

The integration of AI into financial economics represents more than a methodological advancement. It constitutes a fundamental paradigm shift that is reshaping how researchers conceptualize, investigate, and theorize about financial markets. This transformation challenges deeply held assumptions in financial theory and demands a reconceptualization of the discipline’s epistemological foundations.

4.4.1 From hypothesis-driven to data-driven discovery. Traditional financial economics has operated within a hypothesis-driven framework, where researchers derive testable predictions from theoretical models and evaluate these predictions against empirical data ([Mullainathan & Spiess, 2017](#)). ML changes this paradigm by enabling data-driven discovery, where algorithms identify complex patterns and relationships that may not conform to preconceived theoretical structures. In contrast, ML emphasizes predictive accuracy without requiring explicit structural assumptions about data-generating processes. This distinction has profound implications for financial research, as many economically significant questions, including asset pricing, credit scoring and risk assessment, are fundamentally prediction problems where ML’s flexibility offers substantial advantages. The empirical evidence strongly supports this paradigm shift. [Gu et al. \(2020\)](#) demonstrate that ML methods achieve out-of-sample R-squared values that substantially exceed those of traditional linear models in predicting cross-sectional stock returns. These improvements stem from the algorithms’ capacity to capture nonlinear interactions among predictors and the patterns that existing theoretical frameworks had not anticipated. This finding suggests that financial markets contain an economically meaningful structure that extends beyond what current theories describe, pointing toward new avenues for theoretical development.

4.4.2 Challenging fundamental assumptions in financial theory. The success of ML in predicting asset returns poses direct challenges to core tenets of financial theory, most notably the efficient market hypothesis. If markets fully incorporate all available information into prices, systematic return predictability should not persist. However, the ML models consistently demonstrate economically significant predictive power that generates substantial risk-adjusted returns ([Gu et al., 2020](#); [Avramov et al., 2023](#)). This tension demands either a revision of market efficiency assumptions or the identification of economic mechanisms, such as limits to arbitrage or behavioral frictions that reconcile predictability with equilibrium.

Recent theoretical work has begun to address this challenge. [Kelly, Malamud, and Zhou \(2024\)](#) developed a framework demonstrating that model complexity itself can be virtuous for

return prediction. Using random matrix theory, they show that expected out-of-sample performance can increase with model parameterization, overturning conventional wisdom about the bias-variance tradeoff in high-dimensional settings. This virtue of complexity finding provides theoretical grounding for the empirical success of neural networks and other complex ML architectures in finance.

Avramov *et al.* (2023) examine the tension between ML predictions and economic restrictions. They found that incorporating economic constraints into ML models can enhance performance when the constraints are valid but may reduce predictive accuracy when markets exhibit patterns inconsistent with traditional theory. This finding underscores a fundamental challenge that ML reveals: financial markets may operate according to principles that existing theory incompletely captures.

4.4.3 The emergence of economic machine learning. Recognizing that neither pure prediction nor pure structural estimation fully addresses the needs of financial economics, researchers are developing hybrid approaches that combine ML's predictive power with economic theory's interpretive structure. This emerging field of economic ML represents a synthesis that may define the discipline's future trajectory.

Hoang and Wiegatz (2023) identified three archetypes of ML applications in finance. They construct superior measures for traditional analyses, reduce prediction errors in economically relevant forecasts, and extend the standard econometric toolkit. The first archetype is particularly significant for theoretical development, as ML can extract economically meaningful signals from high-dimensional or unstructured data, such as text, images, or transaction-level records that traditional methods cannot effectively process.

The development of transformer-based asset pricing models exemplifies this synthesis. Kelly *et al.* (2025) embed transformer architectures into the construction of stochastic discount factors, maintaining economic structure while leveraging deep learning's capacity to capture complex temporal dependencies. This approach moves beyond mere prediction toward economically interpretable modeling, representing a paradigm shift from using AI to forecast prices to using AI to understand pricing mechanisms.

4.4.4 Implications and future directions. The paradigm shift driven by AI raises fundamental epistemological questions about the nature of financial knowledge. Traditional financial theory sought parsimonious explanations that could be analytically derived and intuitively understood. ML models, by contrast, may capture true market dynamics while remaining largely opaque to human interpretation. This tension between predictive accuracy and interpretability reflects a deeper question about whether financial markets are fundamentally amenable to parsimonious theoretical description or whether their complexity necessitates correspondingly complex empirical characterizations.

Looking forward, the discipline faces a critical choice regarding how to integrate these methodological advances. One path continues to privilege interpretability, using ML primarily as a tool to improve measurement within traditional theoretical frameworks. An alternative path embraces complexity, accepting that financial markets may contain an irreducibly complex structure that only flexible, high-dimensional models can capture. Thirdly, a synthesis-oriented path as the most promising develops new theoretical frameworks that can accommodate and explain the pattern ML reveals, thereby advancing both prediction and understanding.

The revolution in AI-driven financial research is still in its early stages, and its ultimate impact on the discipline will depend on researchers' collective ability to harness these powerful tools while maintaining commitment to economic insight. What is clear is that the paradigm shift is already underway, and financial economics will emerge transformed.

4.5 Emerging challenges, future directions and conclusions

4.5.1 Emerging challenges and future directions. The integration of LLMs and generative AI into finance represents a new frontier with profound implications. However, concerns about hallucination, where models generate plausible but false information, pose significant risks in

financial contexts where accuracy is paramount. LLMs present new opportunities and risks. [Cao et al. \(2024\)](#) examine GPT-4's applications in financial analysis. [Lopez-Lira and Tang \(2024\)](#) demonstrate ChatGPT's ability to predict stock returns from news headlines. [Eisfeldt et al. \(2023\)](#) examined generative AI and firm values. [Kim, Muhn, and Nikolaev \(2024\)](#) analyze financial statement analysis using large language models. Moreover, emerging benchmarks such as FinMME ([Luo et al., 2025](#)) reveal that multi-modal LLMs exhibit significant limitations in specialized financial reasoning tasks, underscoring the need for domain-specific evaluation frameworks.

The challenge of algorithmic fairness extends beyond lending to all areas of AI application in finance. [Rambachan, Kleinberg, Ludwig, and Mullainathan \(2020\)](#) provide an economic approach to regulating algorithmic fairness in financial decisions. [Cowgill and Tucker \(2019\)](#) examine the economics of algorithmic bias with implications for financial services hiring and credit decisions. [Hacker, Engel and Mauer \(2023\)](#) analyze the regulatory implications of LLMs in finance, proposing frameworks for ensuring fairness and transparency in AI-driven financial services.

Looking forward, several critical research directions emerge from our analysis. The development of robust theoretical frameworks for understanding multi-agent AI systems remains essential, including characterizing conditions for market stability and efficiency. The challenge of explainable AI requires continued innovation to balance interpretability with predictive performance. The intersection of privacy and AI demands new approaches that leverage data while protecting individual privacy and preventing discriminatory outcomes. Finally, the governance of AI in finance requires international coordination to prevent regulatory arbitrage while fostering innovation.

4.5.2 Conclusion. In conclusion, it is important to acknowledge the limitations of this study. Our literature review, while extensive, was limited to leading journals in finance, economics and accounting, which may reflect editorial preferences for certain methodological approaches or research questions. Additionally, the articles under consideration primarily focus on developed markets, particularly the United States of America and Europe, while AI applications in emerging markets may face different challenges and opportunities. The rapid pace of technological change means that some findings may quickly become outdated, and the lag between innovation and publication in academic journals means that recent developments may be underrepresented. Despite these limitations, this review provides an essential foundation for researchers entering this interdisciplinary field. It is vital to recognize the potential value of AI-related research in furthering knowledge on the interaction among the disciplines of finance, computer science, economics and law. Therefore, future researchers should approach this field using innovative and interdisciplinary methods to expand our understanding of these complex and evolving technologies. The revolution of AI in finance has only just begun, and its ultimate impact will depend on the collective efforts of researchers, practitioners and policymakers in shaping its development and deployment.

Acknowledgments

We sincerely thank Professor Qiang Wu, Editor-in-Chief of China Accounting and Finance Review, Professor Nan Yang (The Hong Kong Polytechnic University) and Professor Wei Zhou (Yunnan University of Finance and Economics), Guest Editors of the CAFR Special Issue Conference 2025, and the editorial office for their continuous support. We are also grateful to the anonymous reviewers for their constructive comments, which have substantially improved this paper. We also wish to thank the Hong Kong Generative AI Research and Development Center (HKGAI) for providing the research environment that made this work possible.

Note

1. Preferred Reporting Items for Systematic Reviews and Meta-Analyses.

Supplementary material

The supplementary material for this article can be found online.

References

- Acemoglu, D., Autor, D., Johnson, S., & Reshef, A. (2024). Low-skill and high-skill automation. Working paper No. 32140. Cambridge, MA: National Bureau of Economic Research.
- Albanesi, S., & Vamossy, D. F. (2019). Predicting consumer default: A deep learning approach. Working paper No. 25685. Cambridge, MA: National Bureau of Economic Research, August.
- Aquilina, M., Budish, E., & O'Neill, P. (2022). Quantifying the high-frequency trading 'arms race'. *Quarterly Journal of Economics*, 137(1), 493–564.
- Avramov, D., Cheng, S., & Metzker, L. (2023). Machine learning vs. economic restrictions: Evidence from stock return predictability. *Management Science*, 69(5), 2587–2619. doi: [10.1287/mnsc.2022.4449](https://doi.org/10.1287/mnsc.2022.4449).
- Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151, 103745. doi: [10.1016/j.jfineco.2023.103745](https://doi.org/10.1016/j.jfineco.2023.103745).
- Bartlett, R., Morse, A., Stanton, R., & Wallace, N. (2022). Consumer-lending discrimination in the FinTech era. *Journal of Financial Economics*, 143(1), 30–56. doi: [10.1016/j.jfineco.2021.05.047](https://doi.org/10.1016/j.jfineco.2021.05.047).
- Berg, T., Burg, V., Gombović, A., & Puri, M. (2020). On the rise of FinTechs: Credit scoring using digital footprints. *Review of Financial Studies*, 33(7), 2845–2897. doi: [10.1093/rfs/hhz099](https://doi.org/10.1093/rfs/hhz099).
- Bloom, N., Hassan, T. A., Kalyani, A., Lerner, J., & Tahoun, A. (2025). The diffusion of disruptive technologies. *The Quarterly Journal of Economics*, 140(2), 1299–1365.
- Bracke, P., Datta, A., Jung, C., & Sen, S. (2019). Machine learning explainability in finance: An application to default risk analysis. Working paper No. 816. London: Bank of England, August.
- Briola, A., & Turiel, J., Marcaccioli, R., Cauderan, A., & Aste, T. (2021). Deep reinforcement learning for active high frequency trading. *arXiv preprint*, Supp. Available from: <https://arxiv.org/abs/2101.07107>.
- Bryzgalova, S., Huang, J., & Julliard, C. (2023). Bayesian solutions for the factor zoo: We just ran two quadrillion models. *Journal of Finance*, 78(1), 487–557. doi: [10.1111/jofi.13197](https://doi.org/10.1111/jofi.13197).
- Cakici, N., Fieberg, C., Metko, D., & Zaremba, A. (2024). Do anomalies really predict market returns? New data and new evidence. *Review of Finance*, 28(1), 1–44. doi: [10.1093/rof/rfad025](https://doi.org/10.1093/rof/rfad025).
- Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10), 3267–3297. doi: [10.1257/aer.20190623](https://doi.org/10.1257/aer.20190623).
- Cao, S., Jiang, W., Wang, J., & Yang, B. (2024). From man vs. machine to man + machine: The art and AI of stock analyses. *Journal of Financial Economics*, 160, 103910. doi: [10.1016/j.jfineco.2024.103910](https://doi.org/10.1016/j.jfineco.2024.103910).
- Casgrain, P., Ning, B., & Jaimungal, S. (2022). Deep Q-learning for Nash equilibria: Nash-DQN. *Applied Mathematical Finance*, 29(1), 62–78.
- Chen, L., Pelger, M., & Zhu, J. (2024). Deep learning in asset pricing. *Management Science*, 70(2), 714–750. doi: [10.1287/mnsc.2023.4695](https://doi.org/10.1287/mnsc.2023.4695).
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2024). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 27(1), C1–C68. doi: [10.1111/ectj.12097](https://doi.org/10.1111/ectj.12097).
- Cont, R., Kotlicki, A., & Valderrama, L. (2020). Liquidity at risk: Joint stress testing of solvency and liquidity. *Journal of Banking and Finance*, 118, 105871. doi: [10.1016/j.jbankfin.2020.105871](https://doi.org/10.1016/j.jbankfin.2020.105871).
- Cowgill, B., & Tucker, C. E. (2019). Economics, fairness and algorithmic bias. *Journal of Economic Perspectives*, 33(3), 23–45.
- Dang, T. V., Li, W., & Wang, Y. (2024). Government stock purchase undermines price informativeness: Evidence from china's "national team". *Journal of Financial and Quantitative Analysis*, 59(5), 2340–2374.

- Danielsson, J., Macrae, R., & Uthemann, A. (2022). Artificial intelligence and systemic risk. *Journal of Banking and Finance*, 140, 106290. doi: [10.1016/j.jbankfin.2021.106290](https://doi.org/10.1016/j.jbankfin.2021.106290).
- Dell, M. (2025). Deep learning for economists. *Journal of Economic Literature*, 63(1), 5–58.
- Dell’Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., ... Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. Harvard Business School Technology & Operations Mgt. Unit Working Paper No. 24-013.
- DeMiguel, V., Gil-Bazo, J., Nogales, F. J., & Santos, A. A. (2021). Can machine learning help to select portfolios of mutual funds? *Management Science*, 67(2), 943–964.
- Desai, A. P., Mallya, G. S., Luqman, M., Ravi, T., Kota, N., & Yadav, P. (2025). Opportunities and challenges of generative-AI in finance. Available from: <https://arxiv.org/abs/2410.15653> (accessed 11 August 2025).
- Didimo, W., Grilli, L., Liotta, G., Menconi, L., Montecchiani, F., & Pagliuca, D. (2020). Combining network visualization and data mining for tax risk assessment. *IEEE Access*, 8, 16073–16086.
- Dong, X., Li, Y., Rapach, D. E., & Zhou, G. (2022). Anomalies and the expected market return. *Journal of Finance*, 77(1), 639–681. doi: [10.1111/jofi.13099](https://doi.org/10.1111/jofi.13099).
- Eisfeldt, A. L., Schubert, G., & Zhang, M. B. (2023). Generative AI and firm values. Working paper No. 31222. Cambridge, MA: National Bureau of Economic Research.
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. Available from: <https://arxiv.org/abs/2303.10130> (accessed 11 August 2025).
- Enriques, L., & Zetsche, D. A. (2020). Corporate technologies and the tech nirvana fallacy. *Hastings Law Journal*, 72, 55–98.
- Feng, G., Giglio, S., & Xiu, D. (2020). Taming the factor zoo: A test of new factors. *Journal of Finance*, 75(3), 1327–1370. doi: [10.1111/jofi.12883](https://doi.org/10.1111/jofi.12883).
- Frost, J., Gambacorta, L., Huang, Y., Shin, H. S., & Zbinden, P. (2019). BigTech and the changing structure of financial intermediation. *Economic Policy*, 34(100), 761–799. doi: [10.1093/epolic/eiaa003](https://doi.org/10.1093/epolic/eiaa003).
- Fuster, A., Plosser, M., Schnabl, P., & Vickery, J. (2019). The role of technology in mortgage lending. *Review of Financial Studies*, 32(5), 1854–1899. doi: [10.1093/rfs/hhz018](https://doi.org/10.1093/rfs/hhz018).
- Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., & Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. *The Journal of Finance*, 77(1), 5–47. doi: [10.1111/jofi.13090](https://doi.org/10.1111/jofi.13090).
- Gambacorta, L., Huang, Y., Qiu, H., & Wang, J. (2019). How do machine learning and non-traditional data affect credit scoring? New evidence from a Chinese fintech firm. *SSRN*.
- Goldstein, M. A., Kwan, A., & Philip, R. (2022). High-frequency trading strategies. *Management Science*, 69(8), 4413–4434.
- Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5), 2223–2273. doi: [10.1093/rfs/hhaa009](https://doi.org/10.1093/rfs/hhaa009).
- Gu, S., Kelly, B., & Xiu, D. (2021). Autoencoder asset pricing models. *Journal of Econometrics*, 222(1), 429–450. doi: [10.1016/j.jeconom.2020.07.009](https://doi.org/10.1016/j.jeconom.2020.07.009).
- Hacker, P., Engel, A., & Mauer, M. (2023). Regulating ChatGPT and other large generative AI models. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1112–1123). New York, NY: ACM.
- Hambly, B., Xu, R., & Yang, H. (2023). Recent advances in reinforcement learning in finance. *Mathematical Finance*, 33(3), 437–503. doi: [10.1111/mafi.12382](https://doi.org/10.1111/mafi.12382).
- Hansen, A. L., & Kazinnik, S. (2024). Can chatgpt decipher fedspeak? *SSRN*, 4399406.
- Heaton, J. B., Polson, N. G., & Witte, J. H. (2017). Deep learning for finance: Deep portfolios. *Applied Stochastic Models in Business and Industry*, 33(1), 3–12. doi: [10.1002/asmb.2209](https://doi.org/10.1002/asmb.2209).
- Hendershott, T., & Riordan, R. (2013). Algorithmic trading and the market for liquidity. *Journal of Financial and Quantitative Analysis*, 48(4), 1001–1024. doi: [10.1017/s0022109013000471](https://doi.org/10.1017/s0022109013000471).

- Hoang, D., & Wiegatz, K. (2023). Machine learning methods in finance: Recent applications and prospects. *European Financial Management*, 29(5), 1695–1743. doi: [10.1111/eufm.12408](https://doi.org/10.1111/eufm.12408).
- Hsieh, Y. Y., & Vergne, J. P. (2023). The future of the web? The coordination and early-stage growth of decentralized platforms. *Strategic Management Journal*, 44(3), 829–857. doi: [10.1002/smj.3455](https://doi.org/10.1002/smj.3455).
- Jagtiani, J., & Lemieux, C. (2019). The roles of alternative data and machine learning in fintech lending: Evidence from the LendingClub consumer platform. *Financial Management*, 48(4), 1009–1029. doi: [10.1111/fima.12295](https://doi.org/10.1111/fima.12295).
- Jullum, M., Løland, A., Huseby, R. B., Ånonsen, G., & Lorentzen, J. (2020). Detecting money laundering transactions with machine learning. *Journal of Money Laundering Control*, 23(1), 173–186. doi: [10.1108/jmlc-07-2019-0055](https://doi.org/10.1108/jmlc-07-2019-0055).
- Kelly, B. T., Malamud, S., & Zhou, K. (2024). The virtue of complexity in return prediction. *Journal of Finance*, 79(1), 459–503. doi: [10.1111/jofi.13298](https://doi.org/10.1111/jofi.13298).
- Kelly, B., Kuznetsov, A., Malamud, S., & Xu, T. (2025). Artificial intelligence asset pricing models. Working paper. Cambridge, MA: National Bureau of Economic Research.
- Khan, M. S. I., Gupta, A., Seneviratne, O., & Patterson, S. (2024). Fed-RD: Privacy-preserving federated learning for financial crime detection. In *IEEE Symposium on Computational Intelligence for Financial Engineering and Economics* (pp. 1–9). doi: [10.1109/CIFEr62890.2024.10772978](https://doi.org/10.1109/CIFEr62890.2024.10772978).
- Khandani, A. E., Kim, A. J., & Lo, A. W. (2010). Consumer credit-risk models via machine-learning algorithms. *Journal of Banking and Finance*, 34(11), 2767–2787. doi: [10.1016/j.jbankfin.2010.06.001](https://doi.org/10.1016/j.jbankfin.2010.06.001).
- Kim, A., Muhn, M., & Nikolaev, V. (2024). Financial statement analysis with large language models. Working paper. Chicago, IL: University of Chicago Booth School of Business.
- Korinek, A. (2024). LLMs level up—better, faster, cheaper: June 2024 update to ‘generative AI for economic research’. *Journal of Economic Literature*, 61(4), 1281–1317.
- Kozak, S., Nagel, S., & Santosh, S. (2020). Shrinking the cross-section. *Journal of Financial Economics*, 135(2), 271–292. doi: [10.1016/j.jfineco.2019.06.008](https://doi.org/10.1016/j.jfineco.2019.06.008).
- Leippold, M., Wang, Q., & Zhou, W. (2022). Machine learning in the Chinese stock market. *Journal of Financial Economics*, 145(2), 64–82. doi: [10.1016/j.jfineco.2021.08.017](https://doi.org/10.1016/j.jfineco.2021.08.017).
- Lettau, M., & Pelger, M. (2020a). Factors that fit the time series and cross-section of stock returns. *Review of Financial Studies*, 33(5), 2274–2325. doi: [10.1093/rfs/hhaa020](https://doi.org/10.1093/rfs/hhaa020).
- Lettau, M., & Pelger, M. (2020b). Estimating latent asset-pricing factors. *Journal of Econometrics*, 218(1), 1–31. doi: [10.1016/j.jeconom.2019.08.012](https://doi.org/10.1016/j.jeconom.2019.08.012).
- Lim, B., & Zohren, S. (2021). Time-series forecasting with deep learning: A survey. *Philosophical Transactions of the Royal Society A*, 379(2194), 20200209. doi: [10.1098/rsta.2020.0209](https://doi.org/10.1098/rsta.2020.0209).
- Lopez-Lira, A., & Tang, Y. (2024). Can ChatGPT forecast stock price movements? Return predictability and large language models. SSRN. Available from: <https://ssrn.com/abstract=4412788> (accessed 20 December 2024).
- Lorenz, J., Silva, M. I., Aparício, D., Ascensão, J. T., & Bizarro, P. (2023). Machine learning methods to detect money laundering in the Bitcoin blockchain in the presence of label scarcity. In *ACM International Conference on AI in Finance* (Vol. 23, pp. 1–8), New York, NY: ACM. doi: [10.1145/3383455.3422549](https://doi.org/10.1145/3383455.3422549).
- Luo, J., Kou, Z., Yang, L., Luo, X., Huang, J., Xiao, Z., Peng, J., Liu, C., Ji, J., Liu, X., Han, S., Zhang, M., & Guo, Y. (2025). FinMME: Benchmark dataset for financial multi-modal reasoning evaluation. *The 63rd Annual Meeting of the Association for Computational Linguistics*, 1, 29465–29489.
- Ma, Y., Han, R., & Wang, W. (2021). Portfolio optimization with return prediction using deep learning and machine learning. *Expert Systems with Applications*, 165, 113973. doi: [10.1016/j.eswa.2020.113973](https://doi.org/10.1016/j.eswa.2020.113973).
- Magnuson, W. (2020). Regulating fintech. *Vanderbilt Law Review*, 71(4), 1167–1226.

- Mhlanga, D. (2021). Financial inclusion in emerging economies: The application of machine learning and artificial intelligence in credit risk assessment. *International Journal of Financial Studies*, 9 (3), 39. doi: [10.3390/ijfs9030039](https://doi.org/10.3390/ijfs9030039).
- Moscatelli, M., Parlapiano, F., Narizzano, S., & Viggiano, G. (2020). Corporate default forecasting with machine learning. *Expert Systems with Applications*, 161, 113567. doi: [10.1016/j.eswa.2020.113567](https://doi.org/10.1016/j.eswa.2020.113567).
- Mullainathan, S., & Spiess, J. (2017). Machine learning: An applied econometric approach. *Journal of Economic Perspectives*, 31(2), 87–106. doi: [10.1257/jep.31.2.87](https://doi.org/10.1257/jep.31.2.87).
- Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub copilot. Available from: <https://arxiv.org/abs/2302.06590> (accessed 11 August 2025).
- Rambachan, A., Kleinberg, J., Ludwig, J., & Mullainathan, S. (2020). An economic approach to regulating algorithms. Working paper No. 27111. Cambridge, MA: National Bureau of Economic Research.
- Rasekhschaffe, K. C., & Jones, R. C. (2019). Machine learning for stock selection. *Financial Analysts Journal*, 75(3), 70–88. doi: [10.1080/0015198x.2019.1596678](https://doi.org/10.1080/0015198x.2019.1596678).
- Sadhvani, A., Giesecke, K., & Sirignano, J. (2021). Deep learning for mortgage risk. *Journal of Financial Econometrics*, 19(2), 313–368. doi: [10.1093/jjfinec/nbaa025](https://doi.org/10.1093/jjfinec/nbaa025).
- Sirignano, J., & Cont, R. (2021). Universal features of price formation in financial markets: Perspectives from deep learning. *Machine Learning and AI in Finance*, 5–15. doi: [10.1080/14697688.2019.1622295](https://doi.org/10.1080/14697688.2019.1622295).
- Tobek, O., & Hronec, M. (2021). Does it pay to follow anomalies research? Machine learning approach with international evidence. *Journal of Financial Markets*, 56, 100588. doi: [10.1016/j.finmar.2020.100588](https://doi.org/10.1016/j.finmar.2020.100588).
- Truby, J., Brown, R., & Dahdal, A. (2020). Banking on AI: Mandating a proactive approach to AI regulation in the financial sector. *Law and Financial Markets Review*, 14(2), 110–120. doi: [10.1080/17521440.2020.1760454](https://doi.org/10.1080/17521440.2020.1760454).
- Wu, S., Irsoy, O., Lu, S., Dabrovolski, V., Dredze, M., Gehrmann, S., . . . Mann, G. (2023). BloombergGPT: A large language model for finance. Available from: <https://arxiv.org/abs/2303.17564> (accessed 11 August 2025).
- Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2024). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 15(2), 1–19. doi: [10.1145/3298981](https://doi.org/10.1145/3298981).
- Yi, Z., Cao, X., Chen, Z., & Li, S. (2023). Artificial intelligence in accounting and finance: Challenges and opportunities. *IEEE Access*, 11, 129100–129123. doi: [10.1109/access.2023.3333389](https://doi.org/10.1109/access.2023.3333389).
- Zetzsche, D. A., Arner, D. W., & Buckley, R. P. (2020). Decentralized finance. *Journal of Financial Regulation*, 6(2), 172–203. doi: [10.1093/jfr/fjaa010](https://doi.org/10.1093/jfr/fjaa010).

Corresponding author

Sirui Han can be contacted at: siruihan@ust.hk