

# Artificial Intelligence in Financial Decision-Making: Forecasting, Portfolio Optimization, and ESG-Related Corporate Finance Analysis

Adrian Lim<sup>1</sup>, Putri Rahayu<sup>2,\*</sup>

Received 16 January 2026

Accepted 23 March 2026

Published 31 March 2026

**Abstract:** Artificial intelligence has become a major methodological force in financial decision-making, but the literature remains fragmented across at least three partially connected domains: financial time-series forecasting, portfolio construction, and firm-level sustainability analysis. This review argues that these domains should be interpreted as parts of a broader decision architecture in which algorithms extract signals from noisy data, transform those signals into investment or financing choices, and then evaluate outcomes under multiple objectives that increasingly include environmental, social, and governance criteria. The review first synthesizes the evolution of forecasting methods from classical econometric models to recurrent neural networks, transformers, and hybrid architectures. It then examines how predictive outputs are translated into allocation rules, with emphasis on mean-variance optimization, shrinkage-based risk estimation, risk parity, hierarchical allocation, and reinforcement-learning-based dynamic rebalancing. The third substantive line concerns corporate finance and sustainable finance, where AI is used not only to predict ESG ratings and financial constraints but also to identify firm heterogeneity, financing frictions, and disclosure-based signals. Across these streams, the article compares predictive and explanatory models, clarifies the role of structured, textual, and alternative data, and evaluates major methodological risks including overfitting, regime instability, interpretability deficits, and institutional dependence. The central conclusion is that the next stage of research should not treat forecasting, allocation, and ESG-related corporate finance as separate literatures. Instead, future work should build integrated frameworks in which market prediction, portfolio design, and firm-level sustainable finance analysis are jointly modeled under explicit assumptions about data quality, decision frequency, and accountability.

**Keywords:** Artificial intelligence; Financial forecasting; Portfolio optimization; ESG; Financial constraints; Sustainable finance



ISSN 2759-7830 (Online)  
ISSN 2760-2508 (Print)

© 2026 The Author(s)  
Published by Jandoo Press Co., Ltd.

This article is licensed under the terms and conditions of the Creative Commons Attribution (CC BY) license: <http://creativecommons.org/licenses/by/4.0/>

## Introduction

The contemporary AI literature in finance has expanded rapidly, but its conceptual organization is less clear than its technical growth. Bibliometric reviews show a sustained shift from early applications in credit scoring and bankruptcy prediction toward market prediction, portfolio construction, fraud detection, textual analytics, and, more recently, explainable finance and sustainable finance (Ahmed et al., 2022; Goodell et al., 2021; Chen et al., 2023). Even within this broad

field, however, many studies still operate inside narrow technical silos. Forecasting papers focus on directional accuracy, regression loss, or benchmark beating. Portfolio studies emphasize Sharpe ratios or drawdown control. ESG and corporate finance studies examine ratings, financing frictions, or disclosure quality. What is often missing is a unifying interpretation of AI as a decision infrastructure rather than a collection of isolated algorithms.

This review starts from the proposition that AI in financial decision-making has two foundational research lines. The

<sup>1</sup> National University of Singapore, Singapore 117592; <sup>2</sup> Nanyang Technological University, Singapore 639798.

\* Corresponding author. Email: [p.rahayu@ntu.edu.sg](mailto:p.rahayu@ntu.edu.sg)

first is market prediction: identifying exploitable patterns in prices, returns, volatility, sentiment, and cross-sectional relations. The second is firm-level feature identification: using structured financial variables, disclosures, and sustainability indicators to infer latent firm characteristics such as financial constraints, governance quality, default risk, or ESG standing. The distinction matters because the first line usually targets tradable signals under short decision cycles, whereas the second line often targets slower-moving attributes that affect capital allocation, cost of capital, financing access, and strategic valuation (Sezer et al., 2020; Gupta et al., 2020; Mai et al., 2019; Lin & Hsu, 2023).

The two lines are analytically distinct but operationally connected. Asset pricing and portfolio construction depend not only on return forecasts but also on information about firms' balance sheets, financing capacity, exposure to sustainability risks, and disclosure credibility. Likewise, ESG analysis is no longer confined to normative screening. It increasingly enters expected-return models, downside-risk analysis, and strategic asset allocation because investors now treat sustainability information as a component of risk transmission, resilience, and regulatory exposure (Friede et al., 2015; Berg et al., 2022; Christensen et al., 2022; Amel-Zadeh & Serafeim, 2018). From this perspective, AI does not merely predict prices; it also transforms corporate attributes into investable signals.

The first problem that pushed finance toward AI was the persistent mismatch between the statistical complexity of financial data and the simplifying assumptions of classical models. Financial time series are noisy, weakly nonstationary, heteroskedastic, heavy-tailed, and exposed to regime shifts. Linear models remain useful because of their transparency and parsimonious structure, but they often struggle when signal generation depends on nonlinear interactions, temporal memory, or multimodal data. Deep learning entered finance largely by promising richer function approximation for such environments, especially in tasks where historical prices, market microstructure information, or textual sentiment interact in nonlinear ways (Fischer & Krauss, 2018; Krauss et al., 2017; Bao et al., 2017). Yet the record is mixed. AI models can improve benchmark performance in specific samples, but they also amplify model risk, overfitting, and sensitivity to regime change. The implication is that better prediction is not synonymous with better decision-making.

The second problem is that finance is not only about predicting the next return. Most important decisions are constrained optimization problems. Investors must translate beliefs into weights under transaction costs, turnover limits, factor exposures, risk budgets, liquidity constraints, and governance requirements. Corporate finance researchers face a parallel challenge: they must translate noisy indicators into inferences about financing constraints, sustainability performance, and capital market access. In both contexts, AI sits upstream of the final decision rule. A forecasting model is therefore only one component of a longer chain that includes target

definition, state representation, optimization, rebalancing, and ex post evaluation (Markowitz, 1952; Kolm et al., 2014; Ban et al., 2018).

A third reason to reconsider the field at a higher level is the rise of sustainable finance. The empirical literature now shows that ESG information affects investor demand, access to finance, cost of capital, bank lending conditions, firm risk, and crisis-period resilience, although the magnitude and direction of effects depend heavily on materiality, rating methodology, and institutional setting (Cheng et al., 2014; El Ghoul et al., 2011; Goss & Roberts, 2011; Albuquerque et al., 2019; Lins et al., 2017; Albuquerque et al., 2020). This shift has created a new role for AI. Algorithms are being used to predict ESG ratings, extract sustainability signals from disclosures, reconcile rating divergence, and identify heterogeneity across firms and markets. The practical significance is clear: once ESG moves from peripheral screening to core financial analysis, AI applications in finance must connect return prediction with firm-level sustainability inference.

The objective of this review is therefore not to juxtapose two unrelated papers on LSTM-based portfolio optimization and financing constraints with ESG ratings. Instead, it uses them as entry points into a larger analytical structure. The paper by Li and Liu illustrates a common AI-finance workflow: generate forecasts with a recurrent network, then feed predictions into portfolio optimization (Li & Liu, 2023). The paper by Liu illustrates the firm-side line: construct a financing-constraint indicator and estimate its association with ESG ratings in the Chinese stock market (Liu, 2022). Taken together, these papers point to a more general question: how should AI be used when financial decision-making requires both short-horizon market prediction and slower-moving inference about firm quality, sustainability, and financing frictions? These two studies are retained throughout the review as focal anchor papers: Li and Liu (2023) on LSTM-based portfolio optimization, and Liu (2022) on financing constraints and ESG ratings in the Chinese stock market.

The review proceeds in eight sections. Section 2 synthesizes forecasting methods, from ARIMA-style baselines to recurrent networks, transformers, and hybrid models. Section 3 discusses how predictions are translated into portfolio decisions, emphasizing the gap between predictive accuracy and out-of-sample allocation performance. Section 4 reviews the AI literature on corporate finance and sustainable finance, focusing on financial constraints, ESG ratings, and firm heterogeneity. Section 5 examines data sources and measurement issues, including market data, financial statements, ESG scores, texts, and alternative data. Section 6 compares predictive and explanatory models, especially under high-stakes financial decisions. Section 7 assesses limitations and risks, including overfitting, sample dependence, interpretability, and institutional context. Section 8 concludes by proposing a more integrated research agenda in which forecasting, allocation, and ESG-related firm analysis are jointly modeled.

## Financial Time-Series Forecasting Methods: from ARIMA to LSTM, Transformers, and Hybrid Models

The forecasting literature in finance predates contemporary AI by decades, and any serious review must begin by recognizing why classical econometric models remain relevant. Linear autoregressive frameworks, ARIMA-type models, and volatility models endure because they encode explicit assumptions about persistence, differencing, and conditional variance. They provide interpretable baselines, strong diagnostics, and relatively low estimation variance. In noisy financial environments, these virtues are substantial. Many AI studies still benchmark against linear time-series models because the real question is not whether neural networks can fit complex functions in-sample, but whether they deliver decision-relevant gains after accounting for instability, transaction costs, and changing regimes (Sezer et al., 2020; Patel et al., 2015).

Nevertheless, the attraction of machine learning in forecasting is easy to understand. Financial variables often exhibit nonlinear relations, threshold effects, and interactions across prices, volume, sentiment, and macro signals. Early machine-learning forecasting studies therefore used support vector machines, feedforward neural networks, and tree-based methods to relax linear restrictions while preserving manageable training costs (Kara et al., 2011; Patel et al., 2015). These models often improved directional classification or short-horizon return prediction, especially when the problem was framed as movement classification rather than precise point forecasting. Yet their limitations were also clear. Static machine-learning models do not naturally represent long-range temporal dependence, and their performance deteriorates when hand-engineered features fail to capture evolving market structure.

Recurrent neural networks, and especially long short-term memory models, became influential because they addressed this temporal limitation. LSTM architectures introduced gated memory mechanisms that can, in principle, retain relevant information over longer sequences while filtering noise. In finance, this made them attractive for predicting prices, returns, and volatility from historical windows of market data (Fischer & Krauss, 2018; Bao et al., 2017). The main empirical appeal of LSTM is not that financial series become suddenly predictable, but that sequence models can absorb richer lag structures and nonlinear dependence than conventional regressions. The paper by Li and Liu fits squarely within this tradition: it uses LSTM-generated forecasts as inputs to portfolio choice, representing the widespread belief that better temporal representation can improve downstream allocation (Li & Liu, 2023).

Even here, however, caution is required. The strongest results in LSTM-based finance often arise in specific samples, universes, or forecast horizons. Sequence models can exploit artifacts of market microstructure, training-window selection, or benchmark construction. They also require careful choices about sequence length, normalization, retraining frequency, and label definition. The literature therefore shows

both promise and fragility. Fischer and Krauss report that LSTM models can outperform traditional classifiers in cross-sectional stock selection, while Bao and colleagues show gains in time-series forecasting when LSTM is combined with stacked autoencoders (Fischer & Krauss, 2018; Bao et al., 2017). At the same time, review studies emphasize that performance heterogeneity across datasets is large, and apparent gains often shrink once costs, turnover, or different market conditions are considered (Sezer et al., 2020).

A second development concerns ensemble and hybrid models. Rather than treating a single algorithm as sufficient, many studies now combine predictive modules with optimization or feature-selection layers. Ma and colleagues integrate machine-learning and deep-learning return prediction with mean-variance and omega portfolio models, while Chen and colleagues combine an improved XGBoost prediction stage with downstream mean-variance selection (Ma et al., 2021; Chen et al., 2021). Ta and colleagues similarly use LSTM-based prediction as an upstream component in quantitative trading and portfolio formation (Ta et al., 2020). The logic behind such hybrids is straightforward. Prediction and decision are different tasks, and separating them may improve overall system performance. A dedicated prediction block can identify candidate assets or expected returns, while an optimization block imposes diversification and risk control.

Transformers introduced a new stage in this literature. Originally developed for natural language processing, they replaced recurrent computation with attention mechanisms that can model long-range dependence more efficiently and flexibly. Time-series variants such as Informer were designed specifically for long-sequence forecasting, which is especially appealing when financial prediction uses multivariate histories, irregular patterns, or multimodal inputs (Zhou et al., 2021). In theory, transformer models can identify relevant time points or features without compressing all temporal information into a fixed hidden state. In practice, they may outperform recurrent models when the series is sufficiently long and when cross-variable relations matter. Yet their financial advantage is still conditional. Transformers are data hungry, computationally costly, and harder to regularize in low signal-to-noise environments. Their theoretical flexibility may exceed the information content of the data they are given.

Another major extension of forecasting is the incorporation of text and sentiment. Financial prices do not move only because of lagged returns. News, earnings calls, annual reports, social media sentiment, and management disclosures all shape expectations. Tetlock's study on media content remains foundational because it shows that textual negativity helps explain market behavior beyond price history (Tetlock, 2007). Bollen and colleagues later extend this line to social media mood, while more recent reviews document a broad field of finance-oriented text mining (Bollen et al., 2011; Gupta et al., 2020). The methodological implication is that forecasting should be viewed as a representation problem. Market data provide one representation of state; text provides another. AI becomes valuable not merely through non-

linear fitting, but through the ability to combine heterogeneous data sources.

This multimodal logic is especially important for firm-level prediction tasks. Mai and colleagues demonstrate that textual disclosures improve bankruptcy prediction, suggesting that narrative information contains signals about corporate quality that structured ratios alone cannot capture (Mai et al., 2019). The same lesson applies to sustainability analysis. ESG controversies, governance practices, and financing frictions may be only partially observable in accounting variables but more fully reflected in disclosure patterns, tone, and unstructured corporate communication. Thus, the forecasting literature and the firm-feature literature are already converging at the level of data representation, even when they are still treated as separate subfields.

What, then, is the correct benchmark for forecasting success? A common mistake is to focus on statistical metrics in isolation. Lower root-mean-squared error or higher directional accuracy does not automatically improve decisions. Financial forecasts are useful only to the extent that they improve ranking, screening, timing, or hedging after costs and constraints. This is why the literature increasingly emphasizes economic evaluation rather than purely statistical evaluation. A model that slightly improves predictive loss but generates unstable weights and high turnover may be less useful than a simpler model with weaker raw accuracy but better decision stability (DeMiguel et al., 2009; Kolm et al., 2014). The same principle applies to corporate finance forecasting: better ESG-score prediction is valuable only if it improves inference about financing conditions, valuation, or sustainability risk.

A further issue is target definition. Many studies use next-day return prediction, but finance offers multiple possible targets: direction, abnormal return, volatility, drawdown, tail loss, distress probability, or ESG-score change. Different targets imply different loss functions and different notions of success. A model optimized for one-step return prediction may be irrelevant for long-horizon strategic allocation. Likewise, a model that predicts aggregate ESG scores may miss material subdimensions, while a financial-constraint model may work for manufacturing firms but fail for digital firms with intangible-heavy balance sheets. Target selection is therefore not a technical detail. It defines the decision problem the model is actually solving.

In summary, the evolution from ARIMA-style models to LSTM, transformers, and hybrids reflects a real broadening of representational capacity, but not a universal solution to financial forecasting. The strongest lesson of the literature is comparative rather than triumphalist. Classical econometric models remain indispensable as transparent baselines. Machine-learning models are useful when nonlinearities and feature interactions matter. Sequence models are useful when temporal structure is central. Transformer architectures are promising when long-range dependency and multimodal integration become important. Hybrid systems are attractive when the forecasting stage must be tightly linked to screening or optimization. The real frontier is therefore not one superior algorithm but the design of forecasting systems matched

to the data, horizon, and decision rule under study (Sezer et al., 2020; Fischer & Krauss, 2018; Zhou et al., 2021).

## From Prediction to Allocation: Portfolio Optimization, Risk-Return Trade-off, and Dynamic Rebalancing

Portfolio optimization is the point at which financial prediction becomes an actionable decision. This stage is theoretically older than AI, but AI has changed both its inputs and its structure. Modern portfolio theory, beginning with Markowitz, formalized asset allocation as a trade-off between expected return and variance (Markowitz, 1952). Sharpe's asset-pricing framework further linked expected returns to systematic risk, helping establish the canonical structure of risk-adjusted choice (Sharpe, 1964). Yet decades of empirical work have shown that the practical implementation of mean-variance optimization is plagued by estimation error. Small changes in expected returns or covariances can generate unstable, concentrated, and economically implausible weights. This is one reason why naive diversification can outperform optimized portfolios out of sample (DeMiguel et al., 2009).

AI enters this problem in two main ways. First, it can improve the estimates that feed an optimizer. Instead of using historical mean returns, one may use machine-learning forecasts. Instead of relying on sample covariances, one may use shrinkage estimators, latent factor models, or learned risk structures. Second, AI can modify the allocation rule itself. Rather than solving a static mean-variance problem, one can use reinforcement learning or online learning to choose sequential portfolio actions directly (Ban et al., 2018; Yang, 2023; Li & Hoi, 2014).

The literature on AI-assisted portfolio formation often begins by inserting forecasts into classical optimizers. Li and Liu represent a straightforward version of this approach: LSTM outputs are used within maximum-Sharpe and minimum-variance frameworks (Li & Liu, 2023). Ma and colleagues provide a broader comparison by combining return prediction from random forests, support vector regression, LSTM, and other models with mean-variance and omega optimization (Ma et al., 2021). Chen and colleagues build a two-stage system in which an XGBoost-based prediction block selects promising assets before a mean-variance allocator determines weights (Chen et al., 2021). Ta and colleagues similarly embed LSTM prediction into a portfolio construction workflow (Ta et al., 2020). These studies share a common assumption: that improved predictive information can be translated into superior portfolios if optimization controls risk.

This assumption is reasonable but incomplete. The central challenge is not only prediction quality but forecast usability. Portfolio construction is highly sensitive to forecast error. A small upward bias in expected returns can induce excessive concentration; a slight ranking error can push the optimizer toward high-turnover positions with weak economic value. For this reason, many practical systems use AI for ranking or preselection rather than exact return estimation. The machine-learning block filters the asset universe or produces ro-

bust relative scores, and the optimizer then imposes diversification constraints. Ban and colleagues formalize this connection by studying the interaction between machine learning and portfolio optimization more generally (Ban et al., 2018).

Risk estimation is equally important. Sample covariance matrices are notoriously unstable in high dimensions, especially when the number of assets is large relative to the length of history. Ledoit and Wolf's shrinkage estimators remain central because they improve conditioning and reduce estimation variance, which often matters more for realized performance than marginal improvements in expected-return forecasts (Ledoit & Wolf, 2004a; Ledoit & Wolf, 2004b). In practical terms, AI-based portfolio models that boast sophisticated return prediction but rely on naïve covariance estimation may still fail out of sample. This is one of the most underappreciated points in the recent literature: portfolio performance is a joint function of alpha estimation, risk estimation, and execution rules, not of predictive architecture alone (Kolm et al., 2014).

An alternative tradition shifts attention from mean-variance optimization to risk budgeting and diversification structures. Risk parity and equal risk contribution portfolios allocate capital by balancing risk contributions rather than maximizing forecast-adjusted utility. Maillard and colleagues analyze the properties of equally weighted risk contribution portfolios, while Roncalli and Weisang extend risk parity toward risk-factor formulations (Maillard et al., 2010; Roncalli & Weisang, 2016). López de Prado introduces hierarchical risk parity, which uses clustering and recursive bisection to avoid some of the instability associated with inverting noisy covariance matrices (López de Prado, 2016). These approaches are highly relevant to AI because they show that richer prediction is not the only route to better portfolios. Sometimes the key gain comes from more robust structural diversification rather than more aggressive expected-return estimation.

Online and reinforcement-learning approaches go further by replacing the static optimizer with a sequential decision process. In online portfolio selection, the algorithm updates weights iteratively as new information arrives. Li and Hoi survey this tradition and show how follow-the-winner, follow-the-loser, pattern-matching, and meta-learning strategies can all be interpreted within a unified sequential decision framework (Li & Hoi, 2014). Li and colleagues' moving average reversion strategy exemplifies the use of online learning ideas in allocation, where the core intuition is that mean reversion can be exploited over multiple periods without explicitly forecasting full return distributions (Li et al., 2015).

Deep reinforcement learning extends this logic by treating portfolio management as a dynamic control problem. The agent learns a policy that maps the current state to portfolio weights while optimizing long-run reward, usually some function of return and risk. Deng and colleagues provide an influential early example of deep direct reinforcement learning for trading signals and allocation (Deng et al., 2017). More recent work by Yang develops a task-context mutual actor-critic framework for portfolio management, and Guan and Liu focus on explaining the behavior of deep reinforcement-learning agents in the portfolio setting (Yang, 2023; Guan & Liu, 2021).

These studies are important because they relax the separation between prediction and optimization. The agent does not first predict returns and then solve a portfolio problem; it learns allocation actions directly from reward feedback.

This direct approach has several attractions. It can incorporate transaction costs, delayed rewards, and path dependence more naturally than one-shot optimization. It also aligns well with dynamic rebalancing, where the value of a trade depends on future states, not merely on current expected return. Yet reinforcement learning also amplifies some of finance's hardest identification problems. Reward functions may be misspecified, exploration is difficult in nonstationary markets, and learned strategies may overfit historical dynamics that are not repeatable. Performance can look impressive in backtests while depending heavily on environment design, feature engineering, or training-window luck. For this reason, reinforcement learning in finance remains promising but methodologically demanding rather than mature.

A critical issue throughout this literature is the difference between statistical forecasting gains and realized utility gains. A model may improve prediction in a regression sense yet fail to improve investor welfare because of turnover, estimation error, concentration, or slippage. Conversely, a model with mediocre point forecasts may still improve portfolio decisions by ranking assets better, stabilizing exposures, or avoiding extreme losses. This is why evaluation should move beyond raw predictive metrics toward realized Sharpe ratios, drawdowns, turnover-adjusted returns, tail risk, and robustness across regimes (Kolm et al., 2014; DeMiguel et al., 2009). The most meaningful contribution of AI to portfolio optimization is therefore not simply better forecasting, but better integration of signal extraction with robust decision rules.

The LSTM-based portfolio paper included in the present review is useful precisely because it illustrates both the promise and the limits of the prediction-to-allocation pipeline (Li & Liu, 2023). The promise lies in combining nonlinear sequence learning with standard portfolio rules. The limit is that the framework still depends on assumptions about the stability of forecast quality, the appropriateness of maximum-Sharpe or minimum-variance objectives, and the transferability of sample-specific gains. Similar comments apply to many recent portfolio studies. The field is progressing, but the key bottleneck remains economic robustness rather than architectural novelty.

The strongest current direction is likely a layered design. In such a design, prediction models generate robust expected-return or ranking signals; risk models provide stabilized covariance or factor estimates; allocation modules impose diversification, turnover, and ESG constraints; and rebalancing policies are learned or calibrated under explicit cost assumptions. This view is more realistic than the search for a single end-to-end algorithm. It also connects naturally to sustainable finance, because ESG characteristics can enter the allocation layer as constraints, penalties, or state variables rather than as a separate normative overlay (Friede et al., 2015; Berg et al., 2022).

## Corporate Finance and Sustainability Analysis

If market forecasting is the first main line of AI in finance, firm-level feature identification is the second. This line includes bankruptcy prediction, credit risk, disclosure analysis, financing constraints, governance assessment, and ESG-score prediction. Its central task is to infer economically meaningful firm characteristics from large, noisy, and heterogeneous data. In recent years, sustainable finance has moved this line to the center of the field because ESG performance is now linked to capital allocation, risk management, and valuation rather than only to ethical screening ([Friede et al., 2015](#); [Amel-Zadeh & Serafeim, 2018](#)).

The empirical case for the financial relevance of ESG is now extensive but not uniform. Meta-evidence indicates that a large share of studies report a nonnegative association between ESG and financial performance, though effect sizes vary by design, period, and outcome measure ([Friede et al., 2015](#)). More specific studies show that corporate social responsibility is associated with improved access to finance, lower costs of capital, and lower loan spreads under some conditions ([Cheng et al., 2014](#); [El Ghouli et al., 2011](#); [Goss & Roberts, 2011](#)). ESG-related behavior also appears linked to lower firm risk and stronger resilience during crisis episodes ([Albuquerque et al., 2019](#); [Lins et al., 2017](#); [Albuquerque et al., 2020](#)). These findings have made ESG information financially consequential even for investors without explicit social mandates.

At the same time, the ESG literature also highlights deep measurement problems. ESG ratings diverge substantially across providers because they differ in scope, indicators, weights, and treatment of controversies ([Berg et al., 2022](#); [Christensen et al., 2022](#)). This creates an immediate opportunity for AI but also a methodological warning. Machine-learning models can predict provider-specific scores with reasonable accuracy, as demonstrated in studies such as Lin and Hsu, but high predictive accuracy does not solve the underlying disagreement about what is being measured ([Lin & Hsu, 2023](#)). In other words, AI can reproduce rating systems more easily than it can resolve their conceptual inconsistency.

This tension is especially important when ESG analysis is linked to corporate finance. The paper by Liu examines whether financing constraints affect firms' ESG ratings in the Chinese stock market ([Liu, 2022](#)). Substantively, this is a plausible question. Firms facing tight financing conditions may underinvest in environmental compliance, social programs, governance improvement, or disclosure quality because they prioritize liquidity preservation and short-run survival. Alternatively, constrained firms may use ESG signaling strategically to reduce financing frictions. The relationship is therefore theoretically open rather than obvious. What matters for the present review is that the paper sits at the intersection of two literatures that are often studied separately: financial-constraint measurement and ESG outcome assessment.

The measurement of financing constraints has long been controversial. Kaplan and Zingales famously argue that investment-cash flow sensitivity is not a valid general proxy for

financing constraints, opening a broader debate about construct validity ([Kaplan & Zingales, 2000](#)). Almeida and colleagues focus on the cash-flow sensitivity of cash, Foley and coauthors show how cash holdings are shaped by tax and internal-liquidity considerations, and Whited and Wu propose a structural index of financing constraints risk ([Almeida et al., 2004](#); [Foley et al., 2007](#); [Whited & Wu, 2006](#)). Hadlock and Pierce develop the SA index to move beyond the limitations of the KZ framework, and Farre-Mensa and Ljungqvist question whether commonly used measures actually capture financing constraints in the intended sense ([Hadlock & Pierce, 2010](#); [Farre-Mensa & Ljungqvist, 2016](#)). This literature matters because AI studies that use financing-constraint labels or indices inherit these construct ambiguities. A high-performing model trained on a weak proxy does not produce a strong economic inference.

Once this point is recognized, the contribution of AI to corporate finance becomes clearer. Its value lies less in replacing theory-driven proxies than in combining multiple sources of information to improve latent-variable inference. Financial constraints, governance quality, and sustainability orientation are not directly observed. They are inferred from accounting variables, market outcomes, disclosures, financing structure, ownership patterns, and sometimes textual cues. AI can integrate these heterogeneous sources more flexibly than traditional linear models. But the inference remains only as credible as the conceptual mapping between the latent construct and the observed proxies.

Firm heterogeneity is another reason why AI has become important. ESG and financing dynamics are unlikely to be homogeneous across industries, ownership structures, regulatory environments, and life-cycle stages. Material ESG issues differ across sectors, which is precisely why materiality-based analyses often perform better than generic aggregate scoring ([Khan et al., 2016](#)). Constrained firms in asset-heavy industries may face different sustainability trade-offs from digital firms with intangible capital. Firms in bank-dominated financial systems may respond differently from firms in market-based systems. Emerging markets add further layers involving disclosure quality, state involvement, and policy intervention. AI methods are attractive in this setting because they can detect interactions and nonlinear heterogeneity that are cumbersome in conventional specification search.

Textual and disclosure analysis intensify this potential. Annual reports, management discussion sections, sustainability reports, and earnings calls contain information about governance quality, risk management, strategic orientation, and perhaps greenwashing. Earlier studies show that textual disclosure helps distress prediction ([Mai et al., 2019](#)), and broader reviews confirm that finance text mining has become a substantial field ([Gupta et al., 2020](#)). In sustainable finance, this means AI can be used not only to forecast a rating but also to identify the textual pathways through which a firm signals sustainability commitment or obscures risk. This is particularly important when rating disagreement is high, because textual analysis can uncover which dimensions drive a specific model's prediction.

The rise of ESG-score prediction studies illustrates both the opportunities and the limits of this approach. Lin and Hsu show that machine-learning models can predict ESG scores for Taiwanese nonfinancial companies using structured firm information (Lin & Hsu, 2023). Such work is useful for three reasons. First, it helps identify the variables most associated with provider-specific ESG assessments. Second, it can extend coverage to firms or markets with sparse ratings. Third, it can be used as a screening or monitoring tool when ratings are missing or delayed. But its limitation is equally clear: a predicted ESG score is still conditional on the training label. If the original label is noisy, divergent, or conceptually narrow, prediction quality does not guarantee substantive validity.

A related issue concerns causality. Much of the literature, including work on CSR and access to finance or ESG and cost of capital, relies on observational designs (Cheng et al., 2014; El Ghoul et al., 2011; Goss & Roberts, 2011). AI can improve prediction within such data, but it does not by itself identify causal effects. This distinction is especially important in corporate finance. If constrained firms have lower ESG ratings, is that because financing pressure reduces sustainability investment, because poor ESG raises financing costs, or because both reflect deeper firm quality? Prediction-focused AI models often blur these alternatives. For research aimed at explanation rather than screening, model design must therefore be tied to identification strategy, institutional context, and temporal ordering.

Despite these cautions, AI has real advantages in ESG-related corporate finance analysis. It can capture multidimensional firm states, integrate text with accounting data, identify heterogeneous effects, and improve early-warning systems. It is particularly valuable when the research question is classification or ranking under large-scale, noisy information. It is less decisive when the question is structural explanation under contested constructs. The most defensible position is that AI should augment, not replace, theory-based corporate finance analysis. In the case of financing constraints and ESG, this means using AI to discover patterns, improve measurement, and test heterogeneity, while retaining careful attention to what the chosen labels and proxies actually mean (Hadlock & Pierce, 2010; Farre-Mensa & Ljungqvist, 2016; Berg et al., 2022).

In this sense, the corporate finance line and the portfolio line are converging. Investors increasingly use firm-level sustainability and financing information in cross-sectional allocation, risk control, and engagement strategies. Corporate finance researchers increasingly use machine learning to model disclosure-based or rating-based firm characteristics. The next step is to connect these streams more explicitly. A firm's financing constraints can influence its ESG investments and disclosure quality; these, in turn, can affect its cost of capital, market risk, and inclusion in investor portfolios. AI is well placed to model such linked systems, but only if it moves beyond narrow single-task prediction.

## Data and Indicators

Any evaluation of AI in financial decision-making must take data architecture as seriously as model architecture. Many published comparisons between algorithms are in fact comparisons between data pipelines, label definitions, and preprocessing choices. Finance is a domain in which small differences in frequency, timing, survivorship treatment, or universe selection can materially alter conclusions. This is true for both market forecasting and firm-level sustainability analysis.

The most conventional input set remains market data: prices, returns, volume, volatility measures, order-flow proxies, and sometimes factor returns. These variables are attractive because they are high frequency, standardized, and directly tradable. They are also limited because they are close to the outcome being predicted, which raises the danger of learning transient autocorrelation or microstructure artifacts rather than stable economic relationships. For this reason, classical baselines remain important, and sequence models must be evaluated under carefully defined rolling or expanding windows rather than randomly shuffled samples (Sezer et al., 2020; Fischer & Krauss, 2018).

Financial statements and accounting ratios form the second major data family. These are central in corporate finance, bankruptcy prediction, ESG estimation, and cross-sectional stock selection. They are slower moving than prices but often more interpretable. Variables related to leverage, profitability, liquidity, asset growth, cash holdings, and investment are widely used in both firm-level and asset-pricing studies. The financing-constraint literature, however, shows that even apparently straightforward accounting variables can encode contested constructs (Kaplan & Zingales, 2000; Hadlock & Pierce, 2010). AI models trained on such data therefore require conceptual discipline: users must know whether the model is predicting a realized outcome, a provider score, or a theoretical latent attribute.

ESG data constitute a third family and bring distinct problems. ESG ratings are multidimensional, provider specific, and often incomplete outside large firms or developed markets. They may be available as aggregate scores, pillar scores, or issue-level indicators. Their main practical attraction is that they operationalize sustainability in investable form. Their main analytical weakness is cross-provider divergence (Berg et al., 2022; Christensen et al., 2022). Researchers using ESG data must therefore specify whether they are interested in predicting a given provider's score, identifying material sustainability exposure, or inferring an underlying latent construct that transcends rating systems. These are not equivalent tasks.

Textual data have become indispensable because many financially relevant signals are disclosed narratively rather than numerically. News articles, earnings-call transcripts, annual reports, sustainability reports, analyst commentary, and social media all supply information about expectations, risk perception, and organizational quality (Tetlock, 2007; Bollen et al., 2011; Gupta et al., 2020). Text is particularly important when structured data are lagged or sparse. For example, ESG controversies may surface in news before they enter provider

ratings. Governance problems may appear in tone, emphasis, or omission before they are visible in standard ratios. Likewise, distress risk can often be inferred from managerial language and disclosure complexity (Mai et al., 2019). AI has comparative advantage here because modern language models and document embeddings can convert such unstructured information into tractable features.

Alternative data broaden the field further. Satellite imagery, web traffic, geolocation data, job postings, patent text, supply-chain records, and search intensity have all been used in parts of the broader finance literature, though not always in the core papers reviewed here. The relevance of alternative data lies not only in novelty but in temporal lead. These sources can capture operational or reputational changes before they appear in statements or prices. In ESG-related analysis, alternative data may help monitor physical risk exposure, environmental incidents, labor conditions, or public scrutiny. But they also raise acute concerns about access inequality, reproducibility, and legal or ethical boundaries.

Data frequency and alignment are another underappreciated issue. Market prediction often uses daily or intraday observations, whereas ESG and corporate finance analysis typically use quarterly or annual data. Integrating these two levels requires careful temporal design. A portfolio model that uses annual ESG scores for daily rebalancing may implicitly assume that sustainability signals change more often than they do. Conversely, a firm-level study that matches financing data to contemporaneous ESG scores may obscure lag structures. The integrated perspective proposed in this review therefore requires multi-frequency modeling. Short-horizon market states and long-horizon firm attributes should not be forced into the same temporal resolution without justification.

Label construction can be even more consequential than the raw inputs. In forecasting, researchers must decide whether the label is raw return, excess return, sign, volatility, abnormal performance, or a portfolio-relative ranking. In corporate finance, labels may be distress events, rating categories, financing-constraint indices, or sustainability scores. These labels are not neutral. They reflect theoretical choices about what matters. The example of financial constraints is especially instructive: the KZ, WW, and SA frameworks are not interchangeable, and studies have shown that common measures can disagree materially (Whited & Wu, 2006; Hadlock & Pierce, 2010; Farre-Mensa & Ljungqvist, 2016). A model's apparent performance can therefore change not because the algorithm is better, but because the label is easier to predict.

Preprocessing is similarly decisive. Normalization, winsorization, lagging rules, handling of missing values, treatment of delisted firms, and prevention of look-ahead bias all matter. In text analysis, tokenization, vocabulary restriction, document truncation, and the distinction between domain-adapted and general language models affect results. In ESG work, whether controversies are treated as level variables or event shocks can alter interpretation. Thus, one of the persistent problems in AI-finance comparisons is that published studies often compare end results without making preprocessing decisions sufficiently explicit.

The policy implication is that data governance should be treated as part of the research design, not as a preliminary technical step. This is especially important when AI systems move from academic experimentation to real decision support. A forecasting system for portfolio management needs clear rules about data vintage, revision handling, universe construction, and cost assumptions. An ESG-prediction system needs clear rules about provider selection, missingness, coverage bias, and update frequency. Without such governance, model performance is difficult to interpret and almost impossible to reproduce.

For integrated AI-based financial decision-making, the most promising data architecture is layered. High-frequency market data should capture near-term state changes; firm-level structured data should capture slower-moving financial conditions; text should capture narrative and disclosure-based signals; ESG data should represent sustainability-related exposures and assessments; and alternative data should be added only when their timing, legality, and economic meaning are clear. This layered architecture is what allows the two main lines of finance AI—market prediction and firm-level identification—to be analyzed jointly rather than separately.

## Method Comparison

A recurring source of confusion in AI-finance research is the failure to distinguish between predictive and explanatory aims. Predictive models are judged by out-of-sample accuracy, ranking power, economic utility, or decision performance. Explanatory models are judged by interpretability, identification, construct validity, and the plausibility of causal or structural claims. These goals overlap, but they are not identical. In finance, the difference is especially consequential because many tasks involve high-stakes decisions under legal, fiduciary, or regulatory constraints.

Predictive models dominate the market-forecasting literature. Their objective is straightforward: improve forecasts that can be translated into trading, hedging, or allocation. In this setting, model complexity is acceptable if it improves stable out-of-sample performance. LSTM, transformer, and reinforcement-learning systems are therefore often justified even when they are opaque, because the immediate benchmark is economic performance rather than interpretive clarity (Fischer & Krauss, 2018; Zhou et al., 2021; Yang, 2023). But the justification weakens when the model is deployed in institutional contexts that require explanation, such as fiduciary asset management, regulated advice, or risk governance.

Explanatory models dominate much of corporate finance and sustainable finance. When the research question concerns whether financing constraints affect ESG outcomes, whether CSR reduces the cost of capital, or why ESG ratings diverge, the aim is not merely prediction. It is to infer a mechanism, a channel, or a structural relation (Cheng et al., 2014; El Ghouli et al., 2011; Berg et al., 2022). In these contexts, black-box predictive success is insufficient. A model may predict ESG scores well while obscuring whether it is relying on size, industry, disclosure volume, or genuinely material sus-

tainability information. It may identify correlations between financing variables and ESG scores without clarifying temporal ordering or omitted variables.

This is why explainability has become central in recent discussions of finance AI. Rudin argues that in high-stakes settings, the field should prefer inherently interpretable models where possible rather than rely on post hoc explanations for black boxes (Rudin, 2019). Arrieta and colleagues provide a broader framework for explainable AI, distinguishing among concepts, audiences, and explanatory techniques (Arrieta et al., 2020). In finance specifically, bibliometric evidence shows rapid growth in explainable AI research, indicating that the field increasingly recognizes opacity as a practical constraint rather than a tolerable side effect (Chen et al., 2023).

The trade-off, however, should not be oversimplified. Interpretable models are not always superior, and black-box models are not always irresponsible. The relevant question is whether the model's complexity is proportionate to the task and whether the decision environment requires local, global, or procedural explanation. For example, a high-frequency portfolio signal used internally within a diversified multi-signal platform may tolerate lower interpretability than a model used to deny credit, assign a sustainability score, or justify fiduciary stewardship actions. In other words, the appropriate level of explanation depends on the financial context.

Portfolio management provides a useful middle case. Investors often care less about the internal semantics of a signal than about stable realized performance and controlled exposures. Yet as reinforcement learning and end-to-end systems become more common, explainability becomes necessary for understanding why the strategy loads on particular market states, factors, or asset clusters. Guan and Liu's work on explainable deep reinforcement learning is important precisely because it tries to connect policy behavior to interpretable reference weights and prediction power (Guan & Liu, 2021). This kind of work suggests that finance does not have to choose between prediction and explanation in absolute terms. Instead, it can design systems in which opaque components are surrounded by interpretable diagnostic layers.

In ESG-related corporate finance, the need for explanation is even stronger. ESG ratings affect investor perception, capital allocation, and in some settings regulatory classification. If a machine-learning model predicts ESG scores, users must know whether it is learning material sustainability characteristics or merely proxying for firm size, international visibility, or disclosure capacity. This is not only a methodological problem; it is also a fairness and accountability problem. Divergence across rating providers already reveals that scoring systems embody implicit judgments (Berg et al., 2022; Christensen et al., 2022). Adding AI to this environment without interpretive discipline risks amplifying rather than reducing opacity.

A second comparison concerns causal versus associational modeling. Predictive models are usually associational. They learn patterns that improve forecasts, regardless of whether those patterns are causal. Explanatory finance research often seeks something stronger. Yet AI methods are frequently im-

ported into explanatory settings without corresponding identification strategies. The result can be technically impressive but substantively ambiguous studies. For example, a model may show that certain accounting and market variables predict financing-constraint categories or ESG scores, but this does not establish the causal role of those variables. Therefore, when the objective is explanation, AI should be paired with research designs that address endogeneity, timing, and heterogeneity rather than treated as a substitute for them.

A third comparison concerns evaluation metrics. Predictive models are often judged by error, AUC, directional accuracy, or portfolio return statistics. Explanatory models require additional criteria: variable stability, sign consistency, theoretical coherence, and robustness across specifications. In some cases, a simpler model with slightly weaker predictive power may be preferable because it yields a more credible substantive interpretation. This is especially true in sustainable finance, where the policy and governance implications of model outputs can be large.

The practical conclusion is that finance researchers should specify model purpose before selecting model class. If the goal is high-frequency ranking under dense data and low explanation requirements, complex predictive models may be appropriate. If the goal is firm-level inference under contested constructs and governance relevance, interpretable or hybrid models are often more defensible. The broad theme of this review is not that one class should dominate the other, but that AI in finance should be evaluated as part of a decision system with explicit epistemic goals.

## Limitations and Risks

The literature reviewed above shows substantial progress, but it also reveals recurring limitations that are too often treated as technical inconveniences rather than central research problems. Four risks are especially important: overfitting, out-of-sample instability, interpretability deficits, and institutional dependence.

Overfitting remains the most obvious challenge. Financial data have low signal-to-noise ratios, regime shifts, structural breaks, and adaptive agents. These characteristics make flexible models vulnerable to learning patterns that do not survive outside the training sample. Deep architectures can fit nonlinear relations extremely well, but finance does not reward in-sample fit. It rewards robust, repeatable economic value after costs and under changing conditions (Sezer et al., 2020; Krauss et al., 2017). The risk is particularly acute when researchers tune architectures extensively on a fixed historical window or report the best-performing specification without fully accounting for search. Apparent innovation may then reflect selection bias rather than genuine predictive improvement.

Out-of-sample failure is closely related but deserves separate emphasis. Even when a model is not overfit in a narrow statistical sense, it may still fail because the environment changes. Financial time series are non-stationary not only in mean and variance but also in institutional structure, liquidity, regulation, and investor behavior. A model trained in a

low-rate environment may break in a tightening cycle. A sentiment model trained on one media ecosystem may weaken after platform changes. An ESG model trained on a stable rating methodology may lose validity after provider revisions. Thus, genuine robustness requires temporal, cross-market, and cross-regime testing rather than a single train-test split (DeMiguel et al., 2009; Kolm et al., 2014).

Interpretability deficits are the third risk. In portfolio applications, opacity can conceal unstable exposures, hidden leverage to macro conditions, or spurious dependence on a narrow set of features. In corporate finance and ESG analysis, opacity is more serious because model outputs may influence financing decisions, stewardship, screening, or public claims about sustainability quality. Post hoc explanation methods can help, but they do not always recover the actual reasoning process of a complex model (Rudin, 2019; Arrieta et al., 2020). This is why interpretable design should be considered at the start of model development rather than appended at the end as a compliance layer.

Institutional dependence is the fourth and perhaps least discussed limitation. Financial models are embedded in specific market structures, accounting systems, disclosure regimes, investor bases, and regulatory frameworks. A result obtained in the U.S. equity market may not generalize to China's stock market, and a financing-constraint proxy validated in one setting may behave differently in another. ESG is especially sensitive to institutional context because reporting standards, enforcement, ownership structures, and policy priorities vary widely. The paper on financing constraints and ESG in the Chinese stock market is therefore not just a finance application; it is an example of how institutional setting shapes both the data and the substantive interpretation (Liu, 2022).

These limitations also interact. A model can be overfit partly because the institutional regime represented in the training data is unusually stable. A model can appear interpretable but still lack validity because its features proxy for country-specific disclosure habits rather than economically meaningful behavior. A model can show out-of-sample success in one market and fail in another because the target variable itself is constructed differently. Therefore, robustness should be treated multidimensionally: across time, across datasets, across institutions, and across target definitions.

A related risk is benchmark misselection. Many finance AI papers benchmark against weak baselines, such as poorly specified linear models or naïve optimizers, making gains look larger than they are. Strong baselines should include transparent econometric models, robust covariance estimators, shrinkage techniques, naïve diversification, and where relevant simple textual models. Without such baselines, the incremental value of AI is difficult to judge. This point is especially relevant in portfolio optimization, where historical experience shows that complicated expected-return estimates often underperform simpler diversification rules (DeMiguel et al., 2009; Ledoit & Wolf, 2004b).

Another limitation concerns reproducibility. Proprietary data, inconsistent preprocessing, and incomplete reporting of hyper parameters remain common. This problem is acute

with alternative data and some ESG datasets. Reproducibility is not only a norm of good science; it is a practical necessity in finance because small implementation details can materially affect outcomes. The field would benefit from more standardized evaluation protocols, clearer reporting of data vintages, and stronger separation between model development, validation, and final testing.

Finally, there is a conceptual risk in treating AI as neutral infrastructure. Models always embed objectives. In forecasting, the objective may prioritize short-horizon excess return. In portfolio management, it may prioritize Sharpe ratio, downside risk, or turnover efficiency. In ESG analysis, it may prioritize prediction of existing scores rather than independent assessment of material sustainability performance. These objective choices shape results. They determine which signals are learned and which trade-offs are ignored. Therefore, one of the most important future tasks is to make objective functions explicit and align them with the intended decision context.

## Conclusion and Future Directions

This review has argued that AI in financial decision-making should be understood through a broader architecture rather than through isolated technical applications. The first core line is market prediction, where models attempt to extract tradable signals from prices, returns, volatility, and text. The second core line is firm-level feature identification, where models infer financing constraints, ESG characteristics, distress risk, and other latent corporate attributes. These lines have evolved separately, but they increasingly interact in real financial decisions.

The literature on forecasting shows that the movement from classical econometric models to LSTM, transformers, and hybrid systems has expanded representational power, but not eliminated the basic problems of weak signals and unstable regimes (Sezer et al., 2020; Fischer & Krauss, 2018; Zhou et al., 2021). The portfolio literature shows that predictive gains matter only when they survive the translation into robust allocation under noisy covariance estimates, transaction costs, and rebalancing constraints (Markowitz, 1952; Kolm et al., 2014; Ban et al., 2018). The sustainable finance literature shows that ESG and financing information now affect access to capital, cost of capital, risk, and resilience, but also suffer from substantial measurement disagreement and institutional heterogeneity (Cheng et al., 2014; Berg et al., 2022; Christensen et al., 2022).

The two focal papers included by the user are therefore best interpreted not as isolated cases, but as representative nodes in this larger architecture. The LSTM-based portfolio paper captures the prediction-to-allocation logic that defines much of quantitative asset management (Li & Liu, 2023). The financing-constraint-and-ESG paper captures the firm-side logic in which financial frictions and sustainability outcomes are jointly analyzed (Liu, 2022). The real research opportunity lies in connecting these domains. Future studies should model how firm-level sustainability and financing signals enter cross-sectional expected returns, downside risk, portfolio

constraints, and dynamic rebalancing. They should also examine how investor allocation mechanisms feed back into firms' financing conditions and ESG incentives.

Methodologically, future work should pursue three directions. First, it should build multi-layer models that combine forecasting, risk estimation, and optimization rather than overemphasizing any single stage. Second, it should incorporate explainability more deeply in corporate finance and sustainable finance settings where model outputs influence high-stakes decisions (Rudin, 2019; Arrieta et al., 2020). Third, it should move toward multi-frequency and multi-modal designs that integrate market data, financial statements, ESG indicators, and text under clear temporal logic.

In short, the next stage of AI in finance is not simply more complex forecasting. It is the construction of integrated, accountable, and decision-relevant systems in which market prediction, portfolio design, and ESG-related corporate finance analysis are treated as interdependent components of the same financial decision process.

## References

- Ahmed, S., Alshater, M. M., El Ammari, A., & Hammami, H. (2022). Artificial intelligence and machine learning in finance: A bibliometric review. *Research in International Business and Finance*, 61, 101646. <https://doi.org/10.1016/j.ribaf.2022.101646>
- Albuquerque, R., Koskinen, Y., & Zhang, C. (2019). Corporate social responsibility and firm risk: Theory and empirical evidence. *Management Science*, 65(10), 4451–4469. <https://doi.org/10.1287/mnsc.2018.3043>
- Albuquerque, R., Koskinen, Y., Yang, S., & Zhang, C. (2020). Resiliency of environmental and social stocks: An analysis of the exogenous COVID-19 market crash. *The Review of Corporate Finance Studies*, 9(3), 593–621. <https://doi.org/10.1093/rcfs/c-faa011>
- Almeida, H., Campello, M., & Weisbach, M. S. (2004). The cash flow sensitivity of cash. *The Journal of Finance*, 59(4), 1777–1804. <https://doi.org/10.1111/j.1540-6261.2004.00679.x>
- Amel-Zadeh, A., & Serafeim, G. (2018). Why and how investors use ESG information: Evidence from a global survey. *Financial Analysts Journal*, 74(3), 87–103. <https://doi.org/10.2469/faj.v74.n3.2>
- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- Ban, G.-Y., El Karoui, N., & Lim, A. E. B. (2018). Machine learning and portfolio optimization. *Management Science*, 64(3), 1136–1154. <https://doi.org/10.1287/mnsc.2016.2644>
- Bao, W., Yue, J., & Rao, Y. (2017). A deep learning framework for financial time series using stacked autoencoders and long-short term memory. *PLOS ONE*, 12(7), e0180944. <https://doi.org/10.1371/journal.pone.0180944>
- Berg, F., Kölbel, J. F., & Rigobon, R. (2022). Aggregate confusion: The divergence of ESG ratings. *Review of Finance*, 26(6), 1315–1344. <https://doi.org/10.1093/rof/rfac033>
- Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1), 1–8. <https://doi.org/10.1016/j.jocs.2010.12.007>
- Chen, W., Zhang, H., Mehrlawat, M. K., & Jia, L. (2021). Mean-variance portfolio optimization using machine learning-based stock price prediction. *Applied Soft Computing*, 100, 106943. <https://doi.org/10.1016/j.asoc.2020.106943>
- Chen, X.-Q., Ma, C.-Q., Ren, Y.-S., Lei, Y.-T., Huynh, N. Q. A., & Narayan, S. (2023). Explainable artificial intelligence in finance: A bibliometric review. *Finance Research Letters*, 56, 104145. <https://doi.org/10.1016/j.frl.2023.104145>
- Cheng, B., Ioannou, I., & Serafeim, G. (2014). Corporate social responsibility and access to finance. *Strategic Management Journal*, 35(1), 1–23. <https://doi.org/10.1002/smi.2131>
- Christensen, D. M., Serafeim, G., & Sikochi, A. (2022). Why is corporate virtue in the eye of the beholder? The case of ESG ratings. *The Accounting Review*, 97(1), 147–175. <https://doi.org/10.2308/TAR-2019-0506>
- DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *The Review of Financial Studies*, 22(5), 1915–1953. <https://doi.org/10.1093/rfs/hhm075>
- Deng, Y., Bao, F., Kong, Y., Ren, Z., & Dai, Q. (2017). Deep direct reinforcement learning for financial signal representation and trading. *IEEE Transactions on Neural Networks and Learning Systems*, 28(3), 653–664. <https://doi.org/10.1109/TNNLS.2016.2522401>
- El Ghouli, S., Guedhami, O., Kwok, C. C. Y., & Mishra, D. R. (2011). Does corporate social responsibility affect the cost of capital? *Journal of Banking & Finance*, 35(9), 2388–2406. <https://doi.org/10.1016/j.jbankfin.2011.02.007>
- Farre-Mensa, J., & Ljungqvist, A. (2016). Do measures of financial constraints measure financial constraints? *The Review of Financial Studies*, 29(2), 271–308. <https://doi.org/10.1093/rfs/hhv052>
- Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2), 654–669. <https://doi.org/10.1016/j.ejor.2017.11.054>
- Foley, C. F., Hartzell, J. C., Titman, S., & Twite, G. (2007). Why do firms hold so much cash? A tax-based explanation. *Journal of Financial Economics*, 86(3), 579–607. <https://doi.org/10.1016/j.jfineco.2006.11.006>
- Friede, G., Busch, T., & Bassen, A. (2015). ESG and financial performance: Aggregated evidence from more than 2000 empirical studies. *Journal of Sustainable Finance & Investment*, 5(4), 210–233. <https://doi.org/10.1080/20430795.2015.1118917>
- Goodell, J. W., Kumar, S., Lim, W. M., & Pattnaik, D. (2021). Artificial intelligence and machine learning in finance: Identifying foundations, themes, and research clusters from bibliometric analysis. *Journal of Behavioral and Experimental Finance*, 32, 100577. <https://doi.org/10.1016/j.jbef.2021.100577>
- Goss, A., & Roberts, G. S. (2011). The impact of corporate social responsibility on the cost of bank loans. *Journal of Banking & Finance*, 35(7), 1794–1810. <https://doi.org/10.1016/j.jbankfin.2010.12.002>
- Guan, M., & Liu, X.-Y. (2021). Explainable deep reinforcement learning for portfolio management: An empirical approach. In *Proceedings of the Second ACM International Conference on AI in Finance* (pp. 50:1–50:9). ACM. <https://doi.org/10.1145/3490354.3494415>
- Gupta, R., Chen, M. Y., Hardle, W. K., Lee, T. M., & Mamaysky, H. (2020). A comprehensive review of text-mining applications in finance. *Financial Innovation*, 6, 49. <https://doi.org/10.1186/s40854-020-00205-1>
- Hadlock, C. J., & Pierce, J. R. (2010). New evidence on measuring financial constraints: Moving beyond the KZ index. *The Review of Financial Studies*, 23(5), 1909–1940. <https://doi.org/10.1093/rfs/hhq009>

27. Kaplan, S. N., & Zingales, L. (2000). Investment-cash flow sensitivities are not valid measures of financing constraints. *The Quarterly Journal of Economics*, 115(2), 707–712. <https://doi.org/10.1162/003355300554782>
28. Kara, Y., Boyacioglu, M. A., & Baykan, Ö. K. (2011). Predicting direction of stock price index movement using artificial neural networks and support vector machines: The sample of the Istanbul Stock Exchange. *Expert Systems with Applications*, 38(5), 5311–5319. <https://doi.org/10.1016/j.eswa.2010.10.027>
29. Khan, M., Serafeim, G., & Yoon, A. (2016). Corporate sustainability: First evidence on materiality. *The Accounting Review*, 91(6), 1697–1724. <https://doi.org/10.2308/accr-51383>
30. Kolm, P. N., Tütüncü, R., & Fabozzi, F. J. (2014). 60 years of portfolio optimization: Practical challenges and current trends. *European Journal of Operational Research*, 234(2), 356–371. <https://doi.org/10.1016/j.ejor.2013.10.060>
31. Krauss, C., Do, X. A., & Huck, N. (2017). Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500. *European Journal of Operational Research*, 259(2), 689–702. <https://doi.org/10.1016/j.ejor.2016.10.031>
32. Ledoit, O., & Wolf, M. (2004a). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2), 365–411. [https://doi.org/10.1016/S0047-259X\(03\)00096-4](https://doi.org/10.1016/S0047-259X(03)00096-4)
33. Ledoit, O., & Wolf, M. (2004b). Honey, I shrunk the sample covariance matrix. *The Journal of Portfolio Management*, 30(4), 110–119. <https://doi.org/10.3905/jpm.2004.110>
34. Li, B., & Hoi, S. C. H. (2014). Online portfolio selection: A survey. *ACM Computing Surveys*, 46(3), Article 35. <https://doi.org/10.1145/2512962>
35. Li, B., Hoi, S. C. H., Sahoo, D., & Liu, Z.-Y. (2015). Moving average reversion strategy for on-line portfolio selection. *Artificial Intelligence*, 222, 104–123. <https://doi.org/10.1016/j.artint.2015.01.006>
36. Li, H., & Liu, T. (2023). Portfolio optimization based on the LSTM forecasting model. In *Proceedings of the 2nd International Conference on Financial Technology and Business Analysis* (Vol. 48, pp. 97–106). *Advances in Economics, Management and Political Sciences*. <https://doi.org/10.54254/2754-1169/48/20230431>
37. Lin, H.-Y., & Hsu, B.-W. (2023). Empirical study of ESG score prediction through machine learning: A case of non-financial companies in Taiwan. *Sustainability*, 15(19), 14106. <https://doi.org/10.3390/su151914106>
38. Lins, K. V., Servaes, H., & Tamayo, A. (2017). Social capital, trust, and firm performance: The value of corporate social responsibility during the financial crisis. *The Journal of Finance*, 72(4), 1785–1824. <https://doi.org/10.1111/jofi.12505>
39. Liu, T. (2022, December). Financial constraint' impact on firms' ESG rating based on Chinese stock market. In *Proceedings of the 2022 4th International Conference on Economic Management and Cultural Industry (ICEMCI 2022)* (pp. 1085–1095). Atlantis Press. [https://doi.org/10.2991/978-94-6463-098-5\\_122](https://doi.org/10.2991/978-94-6463-098-5_122)
40. López de Prado, M. (2016). Building diversified portfolios that outperform out-of-sample. *The Journal of Portfolio Management*, 42(4), 59–69. <https://doi.org/10.3905/jpm.2016.42.4.059>
41. Ma, Y., Han, R., & Wang, W. (2021). Portfolio optimization with return prediction using deep learning and machine learning. *Expert Systems with Applications*, 165, 113973. <https://doi.org/10.1016/j.eswa.2020.113973>
42. Mai, F., Tian, S., Lee, C., & Ma, L. (2019). Deep learning models for bankruptcy prediction using textual disclosures. *European Journal of Operational Research*, 274(2), 743–758. <https://doi.org/10.1016/j.ejor.2018.10.024>
43. Maillard, S., Roncalli, T., & Teïletche, J. (2010). The properties of equally weighted risk contribution portfolios. *The Journal of Portfolio Management*, 36(4), 60–70. <https://doi.org/10.3905/jpm.2010.36.4.060>
44. Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1), 77–91. <https://doi.org/10.1111/j.1540-6261.1952.tb01525.x>
45. Patel, J., Shah, S., Thakkar, P., & Kotecha, K. (2015). Predicting stock and stock price index movement using trend deterministic data preparation and machine learning techniques. *Expert Systems with Applications*, 42(1), 259–268. <https://doi.org/10.1016/j.eswa.2014.07.040>
46. Roncalli, T., & Weisang, G. (2016). Risk parity portfolios with risk factors. *Quantitative Finance*, 16(3), 377–388. <https://doi.org/10.1080/14697688.2015.1046907>
47. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
48. Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review: 2005–2019. *Applied Soft Computing*, 90, 106181. <https://doi.org/10.1016/j.asoc.2020.106181>
49. Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3), 425–442. <https://doi.org/10.1111/j.1540-6261.1964.tb02865.x>
50. Ta, V. D., Liu, C.-M., & Tadesse, D. A. (2020). Portfolio optimization-based stock prediction using long-short term memory network in quantitative trading. *Applied Sciences*, 10(2), 437. <https://doi.org/10.3390/app10020437>
51. Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *The Journal of Finance*, 62(3), 1139–1168. <https://doi.org/10.1111/j.1540-6261.2007.01232.x>
52. Whited, T. M., & Wu, G. (2006). Financial constraints risk. *The Review of Financial Studies*, 19(2), 531–559. <https://doi.org/10.1093/rfs/hhj012>
53. Yang, S. (2023). Deep reinforcement learning for portfolio management. *Knowledge-Based Systems*, 278, 110905. <https://doi.org/10.1016/j.knsys.2023.110905>
54. Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., & Zhang, W. (2021). Informer: Beyond efficient transformer for long sequence time-series forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(12), 11106–11115. <https://doi.org/10.1609/aaai.v35i12.17325>