

AI-Native Asset Management: A Product Strategy Framework for Intelligent Investment Platforms

Archana Bandhuvula*

Indian Institute of Management, Ahmedabad-India

* Corresponding Author Email: archanabandhuvula@gmail.com - ORCID: 0000-0002-5247-2250

Article Info:

DOI: 10.22399/ijcesn.5520

Received : 20 January 2025

Revised : 27 March 2025

Accepted : 28 March 2025

Keywords

AI-Native Fiduciary Platforms;
ANFP framework;
platform governance;
intelligent investment
architecture;
public-domain platform
assessment

Abstract:

Artificial intelligence (AI) native asset management is a new generation of investment platforms that integrate prediction, portfolio construction, execution, personalisation and governance into a single product architecture. This review brings together existing peer-reviewed research in financial economics, empirical asset pricing, machine learning, Generative AI/Gen AI, and reinforcement-learning trading, financial textual analysis, delegated portfolio management, robo-advisory systems and platform strategy to analyze the foundations and evolution of AI-based investment decision-making. The review highlights the roles of supervised learning, deep learning, reinforcement learning, volatility forecasting, robust optimization and platform-governance theories in asset-management applications. Although existing techniques have improved return prediction, alternative-data analytics, and financial-text interpretation, the use of alternative data in prediction, and in the understanding of financial text, practical deployment remains constrained by estimation error, false discovery risk, transaction costs, market non-stationarity, model opacity, and fiduciary obligations. This study proposes the AI-Native Fiduciary Platforms (ANFP) framework, a strategic model developed by reviewing the literature, which categorises the intelligent investment platforms based on six important design principles: economically constrained prediction, modular intelligence, governed learning, decision-relevant personalisation, explainability and traceable explanation, and platform-level performance measurement. The current analysis shows that the research is still very focused on the prediction and trading aspects and very little on integrated management of client outcomes, execution feedback loops and governance mechanisms. The review highlights a number of research avenues with potential for further progress: counterfactual evaluation of platforms, interpretable fiduciary intelligence, fairness-aware suitability assessment, robust optimization under regime uncertainty, execution-aware learning architectures, and safe protocols for experimentation. The study finds that the widespread adoption of AI in asset management is not about a single predictive model, but a controlled fiduciary decision-making process of predicting, optimizing, executing, personalizing and overseeing, within an intelligent fiduciary system.

1. Introduction

Portfolio management has long involved tension among normative theory, empirical market regularities, and institutional constraints. The problem of designing a portfolio was expressed as a problem of optimization over the expected return, variance and covariance in the Markowitz mean-variance framework [1]. Capital asset pricing theory links expected returns to market beta and equilibrium risk compensation [2], while efficient-market theory suggests that persistent abnormal

return predictability is difficult to sustain [3]. In response to the appearance of intertemporal consumption and investment requirements, and hedging requirements, portfolio management shifted from a static allocation problem to a sequential decision problem under uncertainty [4, 5] and later, Bayesian allocation methods, equilibrium priors and views-based construction of portfolios are integrated into the practical architecture of institutional investment platforms, such as the Black-Litterman model [6]. These theoretical structures remain at the core, but the

principles have been expanded to cover high dimensional data, non-linear interactions, disaggregated liquidity, behavioural diversity and continuous digital intermediation.

AI-native asset management is typically an asset management platform that utilizes machine learning, Generative AI/Gen AI, automated decision infrastructure, data engineering, human-machine interaction, and governance controls as product primitives, not as analytical tools. This difference is key. A traditional asset management platform can incorporate a predictive model into an ongoing research, portfolio construction or advisory workflow. An AI-native platform, on the other hand, must be built to ensure that data inflows, feature governance, model selection, portfolio optimization, client personalization, execution, risk monitoring, compliance review, and feedback-derived benefits are connected in a singular and coherent product system. The investment product is no longer a fund, a mandate, a model portfolio or an advisory account. The investment product becomes a decision-making environment that adapts, whose success depends on the interplay of predictive accuracy, economic robustness, user trust, operational resilience, and regulatory suitability.

The topic has acquired a new level of importance from a scientific perspective because machine learning and Generative AI/Gen AI have demonstrated applicability to empirical asset pricing and investment decision support. Conventional parametric models, however, have been proven to be unable to represent some nonlinear state spaces and complex interactions, which can be captured by deep learning and related architectures [7]. A number of empirical asset pricing studies have concluded that the prediction of returns is economically weak and fragile, while techniques such as tree-based methods and neural networks are superior to linear benchmarks in high dimensional spaces, when used at scale [8]. When predictive models are converted to strategies that can be traded, other key points of backtest design, leakage control, feature selection, cross-validation of dependence and economic validation become highlighted [9]. Beyond prediction and allocation, algorithmic platforms have been found to affect investors' behaviour and diversification in retail and hybrid channels [10].

The asset management industry is global, and so are modern capital markets, which is one reason for the global relevance of AI-based asset management. Institutional portfolios are exposed to multi-asset, macroeconomic, climate-transition, liquidity-regime, geopolitical, and liability-specific risks. Rapidly emerging ways for retail investors to be

given advice include digital advice and model portfolios and automated rebalancing engines. Meanwhile, market structure has also shifted. Electronic trading, a fragmented trading landscape and high-frequency liquidity provision have changed market structure. From the adaptive markets perspective, financial markets are not in a steady state, but rather their efficiency is contingent upon the context in which they are formed, in which they are created as a result of competition, adaptation and institutional ecology [11]. However, the development of AI-native platforms takes place in a setting where predictive connections may cease to exist, investors' tastes may shift and even potentially hostile strategic dynamics can present themselves.

AI influences more than incremental forecasting. In AI-native asset management, product differentiation takes a new form. The traditional way of differentiating was by investment philosophy, star managers, past performance, distribution channels or fees. Furthermore, the ability to integrate proprietary data pipelines, control model governance, depth of personalisation, execution intelligence, tax-aware optimisation, explainability, and continuous product learning are critical capabilities for AI-native differentiation. Second, investment management platforms can create economies of scope across research, portfolio construction, advisory services, and operations. A model that can learn the cash flow patterns of clients can aid the rebalancing process; execution information can be incorporated into a liquidity-oriented portfolio construction model; and compliance exceptions can be further used to adjust the product guardrails. Thirdly, the use of AI changes the nature of fiduciary responsibility because it is not a single individual but a complex system with several components, ranging from data vendors to oversight committees, model developers, portfolio engineers, user-interface designers, and oversight committees.

Several problems remain unresolved. Achieving good predictive performance in past data does not necessarily mean you can achieve the same performance after transaction costs, the capacity limits, the tax effects, the model decay, and the behavioral response. The multidimensional nature of modern financial data is most likely going to present false discoveries unless the research process is rigorously controlled by solid validation. Model uncertainty can amplify small errors in expected-return and covariance estimates of expected returns and covariance in portfolio optimization, which becomes concentration risk or turnover risk in the portfolio. Robo-advisory systems could also include some simpler methods of eliciting

preferences, risk questionnaires, and the underlying assumptions of the models. The technical blueprint for connecting machine-learning capabilities with product-strategy design remains underdeveloped. The financial machine-learning literature has instead tended to be prediction-oriented, with little attention paid to product design, platform scalability, and user uptake and governance.

The scope of the research is therefore interdisciplinary. From a product-strategy perspective, the central decision concerns where intelligence should be embedded within the platform stack, and while there has been a lot of research into asset prices that has become more and more sophisticated, there is not yet a similar level of knowledge about the product strategy. Optimization techniques have been developed by the operations research world, while the intelligent investment platform must have stable controls of different users and regulatory environments. Though financial economics has developed a thorough understanding of delegation and agency issues, AI-powered platforms add delegation chains of algorithms, data providers, and automated interfaces. A review article is therefore needed to integrate these research streams between the empirical machine learning, delegated portfolio management and platform strategy and product governance at an expert level.

The aim of this review is to provide a product strategy framework to design an AI-native asset management platform, which involves critically analyzing peer-reviewed evidence on the use of machine learning and Generative AI/Gen AI for investment prediction, portfolio construction, automated asset management advice, delegated investments, platform design, and model governance. The review assesses the intelligence of investment platforms in a situation where prediction, optimization, personalization, execution and oversight are not analytical modules but are instead product capabilities that are interdependent. The novel aspect of this review is the aggregation of the various pieces of financial-machine-learning, portfolio-construction, robo-advisory, execution and platform-governance literature into the AI-Native Fiduciary Platforms (ANFP) framework. Studies have previously concentrated on the accuracy of the predictions, model switching, portfolio optimization, delegated management or platform governance as separate fields of study. What is missing is a product strategy account of how to harness these features to a fiduciary investment platform. The goal of the ANFP framework is to ensure that the design of an AI-native asset management service is grounded in six principles that arose from the reviews, which are

economically constrained prediction, modular intelligence, governed learning, decision-relevant personalization, traceable explanation and platform-level performance measurement. The framework is not meant to be an empirically validated model, but rather a conceptual synthesis to inform future research, platform design and comparative evaluation.

2. Literature Review

The field can be divided into five technically connected categories of asset management, which are central to AI-native asset management: return prediction, trading and execution intelligence, portfolio construction, delegated advisory design and platform architecture. The evidence base is uneven across these categories. The strongest quantitative evidence comes from return-prediction studies, especially in the cross-sectional equity context. Several methods of trading have been examined, such as reinforcement learning, recurrent networks, or market microstructure representation, but most of them have limitations of nonstationarity, omission of transaction costs, and evaluation instability. While some ideas and theory have been presented in previous research on robo-advisory and delegated-management, these have not been well correlated with the underlying machinery in machine-learning. Platform strategy research has contributed to understanding modularity, network effects, governance and the design of ecosystems, but not so much research has been conducted on how these apply to investment platforms.

The early research in the trading area, based on reinforcement learning, treated market participation as a problem of maximizing the rewards in a sequence. Aiming to more directly optimize a trading performance criterion as opposed to forecasting, Moody and Saffell proposed a direct approach using recurrent reinforcement learning [13] that is more closely related to actual investment outcomes. It is applicable to the AI-native product strategy as it shows that one can present a clear and explicit connection between the goal of the model and economic actions. However, when an imperfect reward function is derived from a limited number of past trajectories the direct trading target may lead to overfitting. Reinforcement learning can be highly sensitive to the representation of the state, the design of the reward, the constraints on the actions that are allowed, and the simulated market impact. These represent vulnerabilities that become product-level concerns when it comes to deployment of reinforcement learning in an investment platform,

where statistical prediction, policy learning, turnover, risk, and execution effects all require validation, but the mapping from the learning of the policy to turnover and risk as well as execution need to be validated.

Deng et al. [14] investigated this in the contexts of representation learning and decision control, directly learning financial signals and trading using deep direct reinforcement learning. This demonstrates that the truly AI-native platforms are not merely a linear application of a model that first makes a prediction and then has a capital allocation engine. Representations, decisions about transactions, and performance goals can become tightly interwoven. This relationship does pose governance issues, though. Where an economic interpretation of the internal representations is not possible, post-trade attribution is more difficult and may not be available for product teams to give their reasons for changes in exposures. This has made reinforcement-learning methods attractive for adaptively executing and tactically allocating, and imposes more strict control requirements than do traditional signal models.

Supervised learning evidence can offer better performance in the context of equity prediction. In their experimental setup, Krauss, Do, and Huck tested the ability of deep neural networks, gradient-boosted trees, and random forests to predict the returns of a group of S&P 500 companies in the short term (one day ahead) and found that the nonlinear models could model these patterns [15]. Later, Fisher & Krauss showed how LSTMs can be used to predict the stock market, although they pointed out that this could be done in a more flexible manner in a recurrent network than in a feed-forward network [16]. These studies are important for financial forecasting and the effect of the research on the financial asset-management products design was minimal. Predicting equity returns in the short horizon needs to be translated to practical trading processes with price-aware execution, borrowing requirements, turnover limits, liquidity screens, capacity estimates, etc. A model that performs well in a high-turnover, highly liquid setting may not be suitable for a large-scale investment platform.

A complementary perspective comes from limit-order-book and market-microstructure studies. Sirignano and Cont demonstrated that deep learning algorithms can be trained to capture general features of price formation from the vast amount of order book data, a form of market microstructure, and that some of these features may be common to a set of instruments [17]. Also, it is crucial to understand that quality of execution is not merely an operational detail, but a product capability and is

applicable to intelligent investment platforms. Risks that can impact a portfolio's alpha include implementation shortfall, latency, spread costs and adverse selection. Execution models are therefore a subproduct of research models and optimization engines. An AI native platform should consider the significance of the results of an execution and not just provide an after-the-fact tool once a portfolio is created.

Studies of textual and alternative data have increased the number of features that go into investment models. To analyze the financial constraints, textual analysis was used and it has been found that the economically significant firm conditions can be extracted from the textual information [18]. Deriving state variables relevant to the market from unstructured information, such as news headlines, is possible, as demonstrated by Manela and Moreira [19] who derived a news-implied volatility from the text of newspapers related to disaster concerns. This information can be leveraged to drive natural-language processing capabilities for smart investing platforms, like risk monitoring, event detection, thematic research, and client communications. Productization is challenging, however, due to publication bias, source coverage, language drift and shifts in media attention that impact product features derived from text. Document provenance, versioning, entity resolution and temporal alignment are some of the features that a good platform should be able to support. In order to avoid look-ahead bias or inconsistent semantics in textual signals, these controls are necessary.

Research on asset pricing has supported the use of machine-learning approaches for investment prediction well. At the same time, Kelly, Pruitt and Su presented a model which treats characteristics as covariances and connects the relationship between firm characteristics and the risk and return of a firm in a comprehensive model [20]. To investigate the possibility of nonlinearities and interactions, Freyberger, Neuhierl and Weber used nonparametric methods for firm characteristics, and indeed they discovered that such interactions and nonlinearities can have economic importance [21]. In the field of asset pricing, Chen, Pelger and Zhu demonstrated that the method of deep learning can be used to estimate the complex stochastic discount factor structure [22]. The significance of these studies lies in their ability to connect the adaptability of AI and machine learning with the economic principles of expected returns, which is very significant in the context of AI-native asset management. The product-level implication is that the intelligent platforms are not just model-free prediction platforms. They should include

economic regularization, factor exposure constraints, risk attribution, scenario diagnostics, and machine intelligence aligned with portfolio goals and investment beliefs.

Research has also examined online portfolio selection with repeated investment decisions. Online portfolio selection techniques were summarized by Li and Hoi, illustrating that there are considerable variations in the assumptions made by the algorithms with respect to mean reversion, trend following, transaction costs and regret minimization [23]. This literature is relevant for AI-native product strategy as it provides the context of asset management as a process of continuous learning and repeated feedback. However, online portfolio algorithms do not work as well when the stationarity assumptions of the past are falsified, or if there are realistic transaction fees. Therefore, care should be taken in the product design to not make online adaptation look better than offline adaptation. The adaptive allocation constraints are risk budgets, turnover limits, capital-gain considerations, client suitability constraints and stress testing.

Personalization is evolving through automated advice. Das et al. demonstrated how portfolio optimization could be extended to take mental accounts into account, acknowledging the possibility that investors might divide their wealth into subportfolios for specific objectives instead of having a global mean-variance account [24]. This insight is important for strategic use in an intelligent investment platform as personalisation should not be limited to risk tolerance scoring. Portfolio design is influenced by a variety of factors including goal segmentation, cash flow needs, tax planning, time horizons and behavioral intentions. Having a platform that simply correlates the results of a questionnaire with a model portfolio doesn't make it very AI-centric. Investor objectives could be included as constraints, state variables, and feedback signals throughout the product stack to facilitate a more powerful design.

Delegated portfolio management research is an explanation of agency problems. Ou-Yang studied optimal contracts and the influence of incentive contracts on the actions of the delegated portfolio manager in the context of continuous-time delegated portfolio management [25]. In their work, optimal decentralized investment management (ODEIM), Van Binsbergen, Brandt, and Koijen discussed the potential effects of delegating investment and the design of benchmarks on investment decisions [26]. A new layer in delegation appears in a system where human investment committees, machine-learning models, optimization engines and automated interfaces are

all responsible for the delegation. These issues create product-strategy requirements for model accountability, override protocols, performance attribution, and incentive-compatible governance. Otherwise, the problems may be confused with errors due to human judgment or algorithmic output without a control structure being evident.

Platform theory is relevant to investment products that incorporate artificial intelligence. Baldwin and Clark showed that modular architectures can be used to support design evolution by separating the pieces and enforcing the discipline of the interfaces [27]. Two-sided network effects can affect the design of information products, as in the case of Parker and Van Alstyne [28] with the participation of investors, advisors, model providers, custodians, data vendors, and execution venues. They also emphasized platform governance and ecosystem innovation, showing that platform value is shaped by governance rules as much as by technical capacity [29]. Tiwana suggested that what platform ecosystems require is an alignment of the platform's architecture, governance and strategy [30]. In the world of asset management, these ideas translate to a smart platform that must be designed as a modular system that is controlled, monitored and audited from data down to models, portfolios, client interfaces and oversight modules.

Table 1 is a literature comparison that shows a pattern that is vital to product strategy. Research that does consider deployment of the product often neglects to explore it in great detail, and research on delegation or platform design rarely considers financial validation and financial machine learning in depth. The coded studies were published across the years shown in Figure 1 and the assignment of model families to the tasks relevant to the platform in this study is summarized in Figure 2. The distribution indicates a transition from "reinforcement learning", "online allocation" to "deep learning", "textual data", and "platform governance", but there are still gaps in integration across these domains.

In this respect, it is noted that the literature which is available is technically correct but on the strategic level, it is not sufficiently concrete. Though substantial progress has been made in research on return prediction, textual analysis, reinforcement learning and volatility modelling in financial machine learning, rarely have these capabilities of machines been demonstrated in the form of durable products, such as financial services. Research on portfolio-optimisation can be used to explain allocations in uncertainty, but sometimes there is a secondary focus on explainability, governance, execution and experience. Two themes of robo-advisory research emerge from the data: access to

digital platforms and automated allocation, but they are generally not well specified with regard to governance of the fiduciary model and the learning taking place on the platform. Platform-strategy research is useful in explaining modularity and ecosystem control; however, there is currently a lack of link to the financial validation, transaction costs, and suitability requirements. The ANFP framework therefore adds by providing linkage among these streams and a single product architecture that is based on reviews.

A second pattern has to do with metrics for evaluation. Typical metrics used in finance machine-learning research are classification accuracy, directional accuracy, Sharpe ratios, return spreads etc., or out-of-sample fit. Typical studies in delegated management are incentives, welfare or contract efficiency. Modularity, governance and ecosystem control are the areas of platform research. There is a need for multi-aspect evaluation in the world of AI-native asset management, where predictive quality must be accompanied by an economic value, user outcomes, compliance robustness, operational resilience and scalability. The required product strategy framework is provided in the Discussion section.

3. Methodology

The review aimed to be a technical and strategy-based analysis of scholarly journal articles and books about applications of artificial intelligence in asset management. The topics in the keyword logic were machine learning and Generative AI/Gen AI asset pricing, deep learning portfolio management, reinforcement learning trading, online portfolio selection, robo-advising, delegated portfolio management, platform ecosystems, financial text analysis, model validation, portfolio optimization, transaction costs, and financial machine learning. The time frame was limited from 1952 to 2024, as the subject necessitates a connection with the fundamental theories of investment, recent machine learning developments, and recent product design. Only peer-reviewed articles and academic books were used as evidence sources. Conference papers and preprint articles were excluded from the cited evidence, as were industry white papers, regulatory websites, and commercial reports.

The reviewed studies employed a variety of methods which can be categorized by the analytical roles. The structure of portfolio choice and asset-pricing restrictions was built on fundamental financial-economics research that employed equilibrium modelling, dynamic programming, the use of continuous-time optimization, and statistical inference. Supervised learning, deep learning,

reinforcement learning, nonparametric estimation, textual analysis and high-dimensional regularization were used for machine-learning studies. Theories applied for product studies and platform studies were modularity theory, ecosystem governance, and information-product economics. AI-native asset management sits at the intersection of financial economics, machine learning, and product strategy. The review thus checked the evidence that is available for each methodological tradition for the areas of product architecture, investment validity and governance feasibility.

Bias-variance trade-off, regularization, cross-validation, and nonlinear representation were key statistical-learning techniques that governed the investment studies reviewed. The classical supervised learning literature highlights the importance of flexibility to minimize approximation error, but also the risk of increased estimation error unless this flexibility is mitigated by regularization, validation and sample-aware design [32]. Neural-network approaches depend strongly on architecture, optimization, training stability, and the structure of the generating data [33]. Temporal dependence, nonstationarity, low signal-to-noise ratios, and adaptive competition exacerbate these problems in financial environments. Thus, an AI-native platform should not simply be a model because it does a good job on a cross-sectional machine-learning benchmark. Versioned data, purged validation, feature lineage, monitoring, and post-deployment drift detection are mandatory components of the product system.

In the literature surveyed, the reinforcement-learning approach to trading was represented as a sequence of actions in a scenario of uncertain reward. Theoretically, reinforcement learning is suitable for trading, allocation, and execution as it has been shown that actions affect future state, transaction costs are path dependent, and reward functions can be risk-adjusted objectives [34]. However, financial reinforcement learning is different from games and simulated control environments because of the following reasons: Market data are historical, non-repeatable, partially observed, and influenced by strategic agents. The methodological limitation is not limited to statistics. It also has a product component. A well-designed trading policy might be successful under perfect conditions but struggle in real-world settings with liquidity constraints, market impact, or limited borrowing, tax effects, or compliance concerns. Product strategy, therefore, requires reinforcement learning to be a governed part of a wider set of decisions—and not an independent product manager.

Volatility and covariance modelling continued to be the focus of much of the literature as portfolio construction is very much dependent on the way risk is estimated. In fact, the autoregressive conditional heteroskedasticity (ARCH) model and the generalized autoregressive conditional heteroskedasticity (GARCH) model are two widely-used benchmarks that incorporate time-varying volatility, and have played a significant role in risk prediction [35, 36]. Comparative analysis of volatility models has demonstrated that none of the specifications is superior for all assets, horizons and loss functions [37]. In an AI-native platform, it means risk engines need to be ensemble capable, have a long-term view, and be stress-tested. Any prediction layer to estimate what returns to expect must be integrated with a risk layer to estimate covariance, tail loss, liquidity risk, and turnover consequences. While analytical convenience with methodological separation of alpha modeling from risk modeling is possible, the strategic risk is high when the product output is handled in the form of automatic portfolios.

A frequent failing in financial machine-learning applications is weak backtesting methodology. Many trading rules can be tested with the same historical sample and yield apparently significant but data-snooped trading rules [38]. To solve the multiple-testing issue in forecasts, superior predictive ability tests have been developed [39]. Overfitting has also been formalized as a material risk in quantitative finance, particularly for strategy selection across many model variants [40]. These are key considerations for intelligent investment platforms as product teams can optimize both the model families, hyperparameters, asset universes, rebalancing frequencies and risk overlays all at once. For these reasons, any apparent AI advantage may just be a sign of search intensity and not necessarily the result of an economic shift.

The data types analyzed in the reviewed studies range from firm characteristics and past prices, to returns, accounting metrics, macroeconomic indicators, order book information, investor account information, and unstructured text. Each data type supports different product capabilities. Cross-sectional characteristics aid in the selection of stocks and factor tilting. Time-series prices are useful for tactical allocation, recognizing trends, modeling volatility, and analytics in execution. Short horizon execution and estimation of liquidity are supported by order book data. Data on each account can be used for personalization and behaviour nudging. Textual data are used for event detection, thematic analysis and risk interpretation. The methodological difficulty is temporal alignment. A platform should guarantee features

were available at the time of decision-making, proper revision treatment, and target-variable construction without the use of future information. Examples of model types used in the reviewed studies include linear regression, penalized regression, random forests, gradient-boosting machines, feed-forward neural networks, recurrent neural networks, reinforcement-learning policies, nonparametric characteristic model, and stochastic optimization model. Penalized regression is useful in the situation where interpretability and sparse signal selection are desired – tree ensembles are attractive when nonlinear interactions and mixed data types are desired. In situations where representation learning from unstructured or high-dimensional data is needed, deep networks might be most relevant. Under estimation error, predictions must be translated into decisions and stochastic optimization and robust optimization methods are required. Some of the popular implicit technical tools are rolling window validation, walk-forward testing, nested cross validation, transaction cost modeling, covariance shrinkage, scenario analysis, and stress testing.

Product strategy was thus conceptualized as a design issue based on evidence in this review. Not all studies were evaluated solely based on the improvement of a benchmark metric. Their approaches were evaluated based on a deployable intelligent platform with the constraints of scale, governance, explainability, suitability, operational control and economic robustness. This orientation is fitting, since AI-native asset management is not just a research program in predictive modelling but a product-system problem, where statistical evidence needs to withstand translation to decisions, interfaces, and fiduciary processes.

A limited public-domain platform assessment is included as an illustrative application of the ANFP framework. This assessment uses only publicly available descriptions of selected platforms and does not claim access to proprietary data, internal model logic, client-level outcomes, or confidential governance procedures. The purpose is not to provide an empirical ranking of platforms, but to demonstrate how the ANFP dimensions can be used to compare visible platform capabilities and identify areas for future research.

4. Discussion

From the literature reviewed, it is clear that AI-native asset management cannot be considered as a set of prediction models, but rather a layered product system. The findings create a set of four product domains: Intelligence infrastructure, Investment decisioning, Client and advisor

experience, and Governance control. Intelligence infrastructure includes data lineage, feature stores, model registries, systems for data validation, and monitoring. Investment decisioning includes return prediction, risk estimation, optimization, tax management, rebalancing, and execution. Client and advisor experience encompasses the process of discovering goals, explaining options, providing behavioral support, and personalization. Governance control encompasses all aspects of fiduciary oversight, model-risk management, auditability, escalation and compliance. The Product Strategy Framework of Table 3 shows how these domains relate to measurable capabilities and technical risks.

Figure 3 presents the ANFP architecture as a layered product system. The sequence begins with governed data foundations and research intelligence, proceeds through portfolio decisioning and execution intelligence, and extends to the client and advisor layer under governance oversight. The arrows should not be interpreted as a one-way pipeline only; in an AI-native fiduciary platform, monitoring, exceptions, client responses, execution outcomes, and governance reviews feed back into the research and decisioning layers. The purpose of this figure is to show that AI-native asset management requires a controlled learning architecture rather than isolated prediction tools.

The reviewed study set was quantitatively coded and three observations were made. First, the years of publication of the 15 core studies in the literature comparison table ranged from 2001 to 2024 with a mean publication year of about 2015.5 and a median of 2017. The interquartile range is dominated by the late 2010s, reflecting the emergence of deep learning and large-scale empirical asset-pricing literature. Second, some model types were underrepresented: supervised learning and deep learning were largely the focus of return-prediction research, reinforcement learning was primarily used in areas like trading and execution, and platform-governance research was mostly limited to the theoretical and organizational side. Third, there were few studies that simultaneously dealt with prediction, portfolio construction, execution, user outcomes and governance. The lack of joint treatment is the reason why AI-based product strategy is not yet developed, despite the fast development of individual methods, and why it has not been adopted more widely. The distribution of platform capabilities is shown in Figure 4, and Figure 5 summarizes product-risk severity rankings based on the coding procedure used in this review.

First, the main finding is that predictive intelligence is needed, but not enough. Empirical asset-pricing

evidence suggests that both nonlinear and high-dimensional approaches can enhance the return-prediction model, but estimation errors are very significant for portfolio choice. The 1/N diversification puzzle has shown that optimized portfolios do not necessarily outperform simple equal-weighted portfolio benchmarks out of sample when estimation error is high [44]. The outcome is significant because AI-native platforms might become enticed to make the weak predictive signals into aggressive allocation differences. Estimates of expected return are generally noisy, and optimization can induce noise to create unstable weights. Thus, smart platforms should consider the predicted returns as uncertain inputs and not deterministic commands. The decision engine should incorporate the following: Bayesian shrinkage, robust optimization, position constraints, turnover penalties, and scenario testing.

The second is that the capability of measuring covariance and risk must be placed into the first-class product capability. In high dimensional situations, shrinkage estimators have been demonstrated to improve the conditioning of the covariance matrix [45]. When the number of securities is large compared to the effective sample size in a cross sectional equity platform, then high dimensional covariance methods are required [46]. The risk engine should not be a passive reporting module in product terms. It should work to explicitly limit allocations, give explanations, raise rebalance review, and detect regime change. If return models are being created separately from risk models, this can lead to conflicting advice. A high-alpha signal may load on an unintended factor exposure or liquidity risk. Exposure reconciliation should thus be part of the architecture of an AI-native system, as it should assess each recommendation by considering the following factors: expected return, factor risk, idiosyncratic risk, tail risk, liquidity, turnover, and client constraints.

The third one is about false discovery. There are numerous proposed anomalies in the financial literature and multiple-testing concerns are serious. However, Harvey, Liu and Zhu made the case that the significance level of asset pricing tests should be increased when there are numerous factors investigated [47]. Hou, Xue, and Zhang demonstrated that numerous anomalies break down during replication and under more rigorous testing [48]. Green, Hand, and Zhang found that incremental information in firm characteristics should not be considered in isolation, and characteristic signals should be evaluated alongside other characteristics [49]. Chen and Zimmermann emphasized the importance of open and replicable

cross-sectional asset-pricing research. For AI-native platforms, this proof suggests a series of model developments should be managed in the way of research production. All signals must have a historical validity, economic justification, decay monitoring, and retirement criteria. The product roadmap should contain a “signal lifecycle” area to control experimentation, promotion, deployment, monitoring, and decommissioning.

The fourth finding is that investment-platform intelligence has to be task specific. A one-size-fits-all AI solution is not necessarily the right one for all asset-management tasks. Various data-generating structures are used for machine learning in the prediction of equity returns, bond risk premia, volatility forecasting, execution, and personalization. While the application of machine learning to bonds is less well-known, Bianchi, Büchner and Tamoni have applied it to the task of bond risk-premium estimation, thus establishing the relevance of machine learning for bonds, and indeed to any particular domain in which the features are designed and validated [51]. The difficulty of out-of-sample forecast gains and the improvement possible with restrictions or forecast combination has been revealed in the equity-premium prediction literature [52] and [53]. The evidence shows that AI-native platforms should be based on a multi-model architecture. A single monolithic model would be difficult to validate, explain, update, and govern. The modular intelligence allows different techniques to be mapped onto different product tasks like gradient boosting for tabular cross-sectional signals, recurrent networks for extracting temporal state, probabilistic models for uncertainty representation, convex optimization for constrained portfolio construction, and rule-based controls for compliance-critical decisions.

The fifth finding is that the performance of products should be assessed using a multi-objective scoring system. The traditional investment metrics that are used to evaluate investments include drawdown, return, volatility, risk-adjusted return, tracking error, benchmark-relative performance, turnover, and risk. There are also further metrics needed on AI-native platforms: model drift, feature freshness, data-quality exceptions, prediction stability, explanation consistency, client-goal progress, suitability overrides, tax efficiency, execution shortfall, and governance latency. The product strategy scorecard is proposed in Table 3. The table is split into investment alpha and product alpha. Investment alpha refers to the creation of economic value through improved forecasting and allocation. Product alpha is the value created by personalization, automation, reducing friction,

improved behavioral adherence, and improved oversight. This distinction is particularly important because a platform could even enhance the success of its clients even in the face of relatively low pure market alphas, particularly when diversification, tax management and disciplined rebalancing are enhanced.

The sixth result is related to execution and market impact. In multi-period trading models, the transaction costs, constraints, and rebalancing paths should be incorporated into portfolio decisions, and should not be regarded as a static one-period optimization [54]. Under trading frictions, trading policies have been found to be tractable under dynamic portfolio choice with linear rules of rebalancing [55]. Under distributional uncertainty and estimation error, methods can be found to make decisions using data-driven robust optimization [56]. The results support the product principle that allocation, rebalancing, and execution should be co-designed. Recommendations made by an AI native platform without market impact awareness can turn predictive information into trading expenses. Execution intelligence should be used to inform the development of portfolios by penalizing illiquid changes, estimating slippage, and separating urgent risk reduction from discretionary alpha capture.

The seventh observation is that market microstructure risk should be a part of the platform model. Execution results may be adversely impacted when the environment is high frequency or large order, due to flow toxicity and liquidity asymmetry [57]. The lack of simple normality and the presence of heavy tails and volatility clustering has been observed in asset returns in diverse markets [58]. The quality of the volatility proxy and the evaluation loss function are also important sensitivity factors for the volatility forecast comparison [59]. The results suggest that it is not advisable to simplify production risk engines using Gaussian approximation. Platform governance should include stress scenarios, tail metrics, liquidity adjusted risk and drawdown controls. The model outputs should be suppressed or escalated if the market states are outside the training distribution.

The eighth observation concerns the client layer. The impact of AI on asset management is more than just about portfolio math. The platform acts as an intermediary between investors and their actions. If interfaces aren't created properly, they can lead to excessive trading, risk misinterpretation, or relying too heavily on automated suggestions. On the other hand, goal-aware interfaces can help to increase compliance with long-term plans. Literature on mental accounts shows that investors have rational and/or behavioral separations of goals, and that

platforms can depict these separations without giving up on optimization [24]. The goal based architecture should be technically encoded as constraints and state variables and not just marketing categories. For instance, emergency reserves, retirement income, education savings, and philanthropy contributions may have varying liquidity needs, tax considerations, drawdown limitations, and time horizons. This design is aimed at making personalization a system relevant to optimization.

The ANFP framework can be distilled down to six principles for success in building a fiduciary platform using AI. The principles take the reviewed literature and translate it into a reusable product-strategy structure, which can be used in research, platform design and public-domain comparative analysis. From this, a product strategy framework for asset management in an AI-native world follows. The platform should be based on six principles. Prediction should be economically constrained, first. Portfolios should only be transformed after considering uncertainty, transaction costs, capacity and risk exposures, with consideration to machine-learning scores. Secondly, intelligence should be modular. Prediction, risk, execution, personalization, and monitoring should be separated in different models, with interfaces and audit trails defined. Thirdly, learning is to be controlled. Signals are required to go through validation and deployment gates, monitoring and retirement gates. Fourth, personalization should be decision-relevant. The client's information must change constraints, objectives, tax treatment, cash-flow planning and explanations. Fifth, explanations need to be aligned to decision type. Compliance, advisor, and client explanations require different levels of technical detail, but all should remain traceable to model inputs and portfolio effects. Sixth, the level of the platform should be used to judge the performance. Evaluation of return metrics, user outcomes, operational quality and governance metrics should be combined.

4.1 Illustrative Public-Domain Platform Assessment Using the ANFP Framework

Existing market examples have a partial rather than full level of ANFP maturity as illustrated in the platform mapping. BlackRock Aladdin is most effective when used in the role of an institutional infrastructure for portfolio, risk and workflow integration. Wealthfront's best attribute is that it's an automated investing and tax-aware portfolio-management service that is direct-to-client. Robinhood may be treated as a boundary case: a retail digital brokerage with AI-enabled and automated features rather than a conventional

managed investment platform; and that highlights the growing importance of suitability, overtrading protection and governance. It must be understood as a public domain, framework application, and not as a formal ranking by the empirical method.

The strategic lesson is that AI-native asset management should not be based on model complexity alone. Data rights, validation discipline, optimization design, execution feedback, personalization depth, and trust architecture are more likely to be coordinated to result in durable advantage. The evidence reviewed here suggests that intelligent investment platforms should be realized as a controlled adaptive system. Without control, adaptation is overfitting and fiduciary risk. Without adaptation, there are no dynamic products to take advantage of new information. The balance of learning, constraint, personalization and accountability is thus the product challenge.

5. Future Directions

Going forward, AI-native asset management research needs to shift away from evaluation focused on models and toward platform science. The first is the one that is most obvious: causal and counterfactual evaluation. The majority of investment machine-learning research focuses on prediction or portfolio performance on historical replays. But intelligent platforms change behavior when they are deployed: investors can change contribution patterns, advisors can change recommendations, and trades can impact execution costs. Future research should thus create 'what if' scenarios that predict what investors would have achieved had they not received platform recommendations or had received them from a different platform, for instance. This demands user-level longitudinal data, variation in product exposure, and careful separation of market performance from user behavior.

The second research pathway is related to interpretable intelligence in fiduciary settings. When performance is comparable, Rudin presented an argument for preferring a model that can be explained to a human over a black box, when decisions are high-stakes [61]. Financial advice and delegated portfolio management are high stakes fields because the outcomes of the models impact wealth, retirement security, taxation, and risk exposure. Miller demonstrated in his studies that explanation is not just technical, but also social and contrastive, meaning that explanations should include the answer to the question of why a recommendation was made rather than another alternative that is relevant [62]. Future research on explanations in intelligent investment platforms

should link model interpretation with portfolio implications: factor exposure shifts, tax implications, drawdown implications, liquidity trade-offs, and probabilities of client goals. Expert supervision and client confidence will require more than generic feature-importance charts.

A third path relates to fairness, suitability and heterogeneous welfare. Studies on fairness in machine learning systems have revealed that, despite high formal accuracy of predictive systems, bias and inequities can still be present in the systems' predictions [63]. Asset management adds suitability considerations based on risk capacity, liabilities, time horizon, income stability, tax status, and behavioral resilience. Further research should investigate if there are systematic differences in the risk, fees, tax impacts, and drawdowns of AI-native advisory systems for demographic and economic groups. This inquiry should not be made on the basis of "demographic parity" since investment goals could be drawn from a variety of perspectives. The key question is whether similarly placed investors are provided with similarly placed recommendations and whether differences are economically justified.

The fourth is about robust optimization in the presence of model uncertainty. AI models might not be effective during shifts in financial market regimes. Risk measures, martingale pricing, and uncertainty in markets that are incomplete have been important concepts in stochastic finance long before the recent growth of AI-native asset management. Conditional VaR optimization is an easily computable method of managing downside risk and tail exposure [65]. Tools for sequential decisions under uncertainty, particularly when future information and constraints evolve, are provided by multistage stochastic optimization [66]. These approaches should be incorporated into AI-native platforms, where their predictions should be tested within uncertainty sets, stress regimes and distributional shifts in the production decision engine.

The 5th direction is about execution-aware learning. Market-making and limit-order-book models reveal that execution is an optimization problem with inventory risk, order arrival and adverse selection [67]. The potential of applying machine learning to hedging, in the presence of market frictions, has been established in a deep hedging study [68]. Future smart investment platforms should tie execution algorithms into

portfolio construction with common state variables. The portfolio optimizer, for instance, must be aware when a signal is only profitable in advance of the expected impact of the market; and the execution engine must receive information relating to whether a trade will reduce risk, be tax-efficient, or pursue alpha. This integration would bridge the strategic disconnect between model portfolios and real portfolios.

There is also a sixth direction: product experimentation within the boundaries of governance. While digital platforms can be optimized by experimentation, it is impossible to do the same with investment products. Portfolio changes, risk changes and/or advice changes can have financial implications. Safe experimentation approaches should be developed for AI-native asset management, including sandbox experiments, shadow portfolios, delayed deployment gates, suitability-aware treatment assignment, and downside-risk monitoring. This type of research would enable the product teams to learn and do so within the bounds of fiduciary discipline.

A seventh direction is related to platform architecture for multi-agent collaboration. The human component is likely to be present in AI-native investment platforms, including portfolio managers, quantitative researchers, advisors, compliance, model risk committees, and clients. Tool development should therefore be supported by traceable human-machine collaboration. Future systems should document acceptance, modification and rejection of recommendations; impact of overrides on outcomes; and consistency between model uncertainty and model confidence. This would put in place a feedback cycle of both human governance and machine improvement.

The big issue left to answer is not if machine learning can enhance parts of the investment process. Already there are signs that it can. The question that remains to be answered is what would need to be done to make AI a fiduciary product system. The advancement of progress will depend on the finance, machine learning, operations research, human-computer interaction, and law becoming mutually dependent disciplines. To realize the full potential of AI-native asset management, predictive models, optimization techniques, user interfaces, and governance strategies will need to be tested as an integrated decision ecology.

Table 1. Literature comparison matrix for AI-native asset management studies and platform implications

Reference	Data or evidence base	Method family	Investment task	Product-strategy implication	Main limitation for platform deployment
-----------	-----------------------	---------------	-----------------	------------------------------	---

Moody & Saffell [13]	Historical price series	Direct reinforcement learning	Trading policy learning	Reward functions can be aligned with trading outcomes	Sensitive to overfitting and simulated-market assumptions
Deng et al. [14]	Financial signal and price data	Deep reinforcement learning	Representation and trading	Feature learning and action learning can be combined	Governance and interpretability are difficult
Krauss et al. [15]	S&P 500 constituents	Neural networks, boosting, random forests	Statistical arbitrage	Nonlinear models can support stock-selection products	Frictionless assumptions may overstate deployable value
Fischer & Krauss [16]	Equity return sequences	LSTM networks	Market prediction	Temporal representation can support tactical modules	Performance may decay under changing regimes
Sirignano & Cont [17]	Limit-order-book data	Deep learning	Price-formation analysis	Execution intelligence should be part of the product stack	Requires large-scale infrastructure and monitoring
Buehlmaier & Whited [18]	Corporate text and financial data	Textual analysis	Constraint and risk inference	Unstructured data can enrich firm-state modeling	Source drift and temporal alignment are major risks
Manela & Moreira [19]	Newspaper text and market data	Text-based volatility measure	Disaster-risk monitoring	Text can support risk-state detection	Media coverage may not be stable across time
Kelly et al. [20]	Firm characteristics and returns	Characteristic-based covariance modeling	Asset pricing	Characteristics should be connected to risk structure	Complex estimation limits simple product explanation
Freyberger et al. [21]	Cross-sectional firm characteristics	Nonparametric estimation	Return-predictor analysis	Nonlinear characteristic effects should be allowed	High flexibility increases validation burden
Chen et al. [22]	Large-scale asset-pricing data	Deep learning	Stochastic discount factor estimation	Deep models can estimate economically structured pricing systems	Model opacity requires strong oversight
Li & Hoi [23]	Online portfolio literature	Online learning algorithms	Sequential allocation	Continuous learning can inform rebalancing products	Transaction costs and nonstationarity remain difficult
Das et al. [24]	Goal-based portfolio theory	Mental-account optimization	Personalized allocation	Goals should be encoded as decision constraints	Preference elicitation remains complex
Ou-Yang [25]	Continuous-time delegation model	Contract theory	Delegated management	Incentives and controls must shape platform governance	Abstract contracts require practical translation
Van Binsbergen et al. [26]	Delegated investment model	Decentralized management theory	Benchmark and delegation design	Human and machine delegation require clear accountability	Limited treatment of AI systems
Parker & Van Alstyne [28]	Information-product theory	Platform economics	Multi-sided product design	Investment platforms can exploit cross-side data and service effects	Network effects may conflict with fiduciary safeguards

Table 2. AI-native investment-platform architecture and strategic design requirements

Platform layer	Core function	Typical AI or	Strategic decision	Governance
----------------	---------------	---------------	--------------------	------------

		quantitative component		requirement
Data foundation	Ingests market, client, textual, accounting, and execution data	Feature stores, entity resolution, temporal joins	Which data create defensible product advantage	Provenance, availability timestamps, lineage
Research intelligence	Identifies predictive and risk-relevant signals	Penalized models, tree ensembles, neural networks, text models	Which signals deserve promotion into production	Validation, leakage tests, decay monitoring
Portfolio decisioning	Converts forecasts into allocations	Robust optimization, covariance shrinkage, constraints	How model uncertainty should affect capital allocation	Risk attribution, concentration controls
Execution intelligence	Converts target portfolios into trades	Cost models, liquidity models, order scheduling	When alpha capture justifies trading costs	Slippage monitoring, market-impact review
Personalization engine	Maps investor states to objectives and constraints	Goal models, suitability models, behavioral segmentation	How client data should alter recommendations	Suitability audit and bias testing
Explanation interface	Communicates decisions to advisors and clients	Local explanations, factor attribution, scenario narratives	Which explanation is needed for each audience	Traceability and escalation
Oversight system	Monitors models, portfolios, and exceptions	Drift detection, exception rules, audit logs	When automation should be paused or overridden	Model-risk committee and documentation




Figure 1. Publication-year distribution of coded AI-native asset-management studies




Figure 2. Heatmap of model families and platform-relevant investment tasks



Figure 3. AI-Native Fiduciary Platforms (ANFP) product architecture




Figure 4. Distribution of product capabilities addressed by the coded literature




Figure 5. Product-risk severity ranking for AI-native asset-management platforms

Table 3. Platform-level scorecard for evaluating AI-native investment products

Evaluation dimension	Indicator	Desired direction	Product interpretation	Example decision trigger
Predictive	Rolling out-of-sample	Higher	Signal quality after	Reduce allocation if

validity	information coefficient		temporal validation	rolling signal turns negative
Economic value	Turnover-adjusted excess return	Higher	Alpha after implementation costs	Penalize signal with high cost-adjusted decay
Risk control	Expected shortfall and drawdown	Lower	Downside protection under stress	Escalate when stress loss exceeds client tolerance
Portfolio quality	Diversification and factor exposure stability	Higher stability	Avoids unintended concentration	Rebalance when factor exposure leaves policy band
Execution quality	Implementation shortfall	Lower	Measures realized cost of platform decisions	Delay discretionary trades in illiquid regimes
Personalization	Goal-funding probability	Higher	Links portfolio to client objective	Adjust contribution or allocation path
Tax efficiency	Realized tax drag	Lower	Captures after-tax product value	Prefer loss harvesting or low-turnover substitute
Governance	Model-drift exception rate	Lower	Measures platform reliability	Move model to review state when drift persists
Explanation	Advisor/client explanation alignment	Higher	Prevents conflicting narratives	Require review when explanation and attribution diverge

Table 4. Illustrative Public-Domain Platform Assessment Using the ANFP Framework

Platform	Visible platform role	Strongest ANFP alignment	Main visible limitation	How to write it safely
BlackRock Aladdin	Institutional investment-management and risk platform	Data foundation, risk analytics, portfolio oversight, institutional workflow integration	Public information does not reveal all model logic, client-level outcomes, or internal governance procedures	Treat as institutional infrastructure example, not as a full client-facing fiduciary platform
Wealthfront	Automated investing and tax-loss-harvesting platform	Automation, tax-aware portfolio management, risk-profile personalization, direct-to-client digital interface	Limited visible human-advisor layer; public information does not validate long-term welfare outcomes	Treat as robo-advisory/product automation example
Robinhood	Retail brokerage/trading platform with AI-enabled tools and managed strategies	Retail access, execution interface, active-trading tools, engagement layer	More brokerage/trading oriented; not all tools are fiduciary advisory tools; risk of overtrading should be discussed	Treat as boundary case showing why governance and client-protection layers matter

6. Conclusions

AI-native asset management is a significant paradigm shift in the development, management and regulation of investment products. The current evidence presented in this article highlights the significant role of machine learning in the areas of return prediction, risk modelling, trading, textual analysis, and personalization, but the performance of individual models is not enough to justify long-term competitive products. Architecture is the central challenge. Future intelligent investment

platforms are going to require the convergence of data infrastructure, predictive models, portfolio optimisation, execution systems, client interfaces and governance controls into a single, integrated decision system.

This review shows that financial prediction has not only been found to be noisy, it also has been shown to be nonstationary and susceptible to false discovery. The process of portfolio optimization can amplify estimation error, and execution costs and market impact can take away the apparent alpha. While automated advice could be beneficial

to access and discipline, personalization should be built into the objectives, constraints, tax rules, liquidity requirements, and behavioral support—not simplified risk categories. When multiple actors, models and interfaces are involved in value creation, platform strategy research also shows three key principles that have to be taken into account: modularity, governance and ecosystem coordination.

Six design imperatives are therefore proposed, which are the design principles that will be further emphasized in the proposed product strategy framework: economically constrained prediction, modular intelligence, regulated learning, decision-relevant personalization, traceable explanation, and platform-level performance measurement. These imperatives are consistent with the larger science of artificial systems where artifacts are assessed in terms of the relation of structure to function with respect to environmental constraints [70]. They also highlight the need for the organization to strike a balance between exploration and exploitation, as investment platforms require learning from new data, yet not embarking on an uncontrollable experiment with their customers' capital [71]. Product advantage in the new asset management world will require ongoing reconfiguration of data assets, models, governance processes, and client experiences [72], which is only possible with dynamic capabilities. Last but not least, the econometric bases of financial-market analysis are essential, since intelligent systems are still subject to evaluation conditions of uncertainty, dependence, and lack of observation [73].

The ANFP framework supports this by outlining and prioritizing the needs of the product strategy for the AI-native asset management landscape. Prediction does not by itself create a strong investment product; it must be supported by execution awareness, personalization, explanation, and fiduciary risk governance. Empirical research is needed in the future to validate the ANFP framework with the data at the platform level, client outcome level, advisor workflow level, and with governance records.

This topic is important because investment management is rapidly becoming an adaptive platform industry rather than only an industry of advisory services and portfolio products. The best platforms will not simply be those that add AI to existing processes. These will be the ones that make intelligence controllable, understandable, individual, action-aware, and cost-effective. The main takeaway from the field is that AI-native asset management is not a forecasting competition; it is a fiduciary decision architecture.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.
- **Use of AI Tools:** The author(s) declare that no generative AI or AI-assisted technologies were used in the writing process of this manuscript.

References

- [1] Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1), 77–91.
- [2] Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3), 425–442.
- [3] Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2), 383–417.
- [4] Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics*, 51(3), 239–246.
- [5] Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4), 373–413.
- [6] Black, F., & Litterman, R. (1992). Global portfolio optimization. *Financial Analysts Journal*, 48(5), 28–43.
- [7] Heaton, J. B., Polson, N. G., & Witte, J. H. (2017). Deep learning for finance: Deep portfolios. *Annual Review of Financial Economics*, 9, 145–181.
- [8] Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *Journal of Financial Economics*, 135(3), 639–655.
- [9] López de Prado, M. (2018). *Advances in financial machine learning*. Wiley.
- [10] D'Acunto, F., Prabhala, N., & Rossi, A. G. (2019). The promises and pitfalls of robo-advising. *Review of Financial Studies*, 32(5), 1983–2020.
- [11] Lo, A. W. (2004). The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management*, 30(5), 15–29.

- [12] Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press.
- [13] Moody, J., & Saffell, M. (2001). Learning to trade via direct reinforcement. *IEEE Transactions on Neural Networks*, 12(4), 875–889.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] Sirignano, J., & Cont, R. (2019). Universal features of price formation in financial markets: Perspectives from deep learning. *Quantitative Finance*, 19(9), 1449–1459.
- [18] Buehlmaier, M. M. M., & Whited, T. M. (2018). Are financial constraints priced? Evidence from textual analysis. *Review of Financial Studies*, 31(9), 3437–3472.
- [19] Manela, A., & Moreira, A. (2017). News implied volatility and disaster concerns. *Journal of Financial Economics*, 123(1), 137–162.
- [20] Kelly, B. T., Pruitt, S., & Su, Y. (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics*, 134(3), 501–524.
- [21] Freyberger, J., Neuhierl, A., & Weber, M. (2020). Dissecting characteristics nonparametrically. *Review of Financial Studies*, 33(5), 2326–2377.
- [22] Chen, L., Pelger, M., & Zhu, J. (2024). Deep learning in asset pricing. *Management Science*, 70(2), 714–750.
- [23] Li, B., & Hoi, S. C. H. (2014). Online portfolio selection: A survey. *ACM Computing Surveys*, 46(3), Article 35.
- [24] Das, S., Markowitz, H., Scheid, J., & Statman, M. (2010). Portfolio optimization with mental accounts. *Journal of Financial and Quantitative Analysis*, 45(2), 311–334.
- [25] Ou-Yang, H. (2003). Optimal contracts in a continuous-time delegated portfolio management problem. *Review of Financial Studies*, 16(1), 173–208.
- [26] Van Binsbergen, J. H., Brandt, M. W., & Koijen, R. S. J. (2008). Optimal decentralized investment management. *Journal of Finance*, 63(4), 1849–1895.
- [27] Baldwin, C. Y., & Clark, K. B. (2000). *Design rules: The power of modularity*. MIT Press.
- [28] Parker, G. G., & Van Alstyne, M. W. (2005). Two-sided network effects: A theory of information product design. *Management Science*, 51(10), 1494–1504.
- [29] Gawer, A., & Cusumano, M. A. (2014). Industry platforms and ecosystem innovation. *Journal of Product Innovation Management*, 31(3), 417–433.
- [30] Tiwana, A. (2014). *Platform ecosystems: Aligning architecture, governance, and strategy*. Morgan Kaufmann.
- [31] Baldwin, C. Y., & Woodard, C. J. (2009). The architecture of platforms: A unified view. In A. Gawer (Ed.), *Platforms, markets and innovation* (pp. 19–44). Edward Elgar.
- [32] Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Springer.
- [33] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- [34] Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.
- [35] Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007.
- [36] Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.
- [37] Hansen, P. R., & Lunde, A. (2005). A forecast comparison of volatility models: Does anything beat a GARCH(1,1)? *Journal of Applied Econometrics*, 20(7), 873–889.
- [38] White, H. (2000). A reality check for data snooping. *Econometrica*, 68(5), 1097–1126.
- [39] Hansen, P. R. (2005). A test for superior predictive ability. *Journal of Business & Economic Statistics*, 23(4), 365–380.
- [40] Bailey, D. H., Borwein, J. M., López de Prado, M., & Zhu, Q. J. (2017). The probability of backtest overfitting. *Journal of Computational Finance*, 20(4), 39–69.
- [41] Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B*, 58(1), 267–288.
- [42] Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
- [43] Koller, D., & Friedman, N. (2009). *Probabilistic graphical models: Principles and techniques*. MIT Press.
- [44] DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies*, 22(5), 1915–1953.
- [45] Ledoit, O., & Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2), 365–411.
- [46] Fan, J., Fan, Y., & Lv, J. (2008). High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics*, 147(1), 186–197.
- [47] Harvey, C. R., Liu, Y., & Zhu, H. (2016). ... and the cross-section of expected returns. *Review of Financial Studies*, 29(1), 5–68.
- [48] Hou, K., Xue, C., & Zhang, L. (2020). Replicating anomalies. *Review of Financial Studies*, 33(5), 2019–2133.

- [49] Green, J., Hand, J. R. M., & Zhang, X. F. (2017). The characteristics that provide independent information about average U.S. monthly stock returns. *Review of Financial Studies*, 30(12), 4389–4436.
- [50] Chen, A. Y., & Zimmermann, T. (2022). Open source cross-sectional asset pricing. *Critical Finance Review*, 11(2), 207–264.
- [51] Bianchi, D., Büchner, M., & Tamoni, A. (2021). Bond risk premiums with machine learning. *Review of Financial Studies*, 34(2), 1046–1089.
- [52] Rapach, D. E., Strauss, J. K., & Zhou, G. (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *Review of Financial Studies*, 23(2), 821–862.
- [53] Campbell, J. Y., & Thompson, S. B. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies*, 21(4), 1509–1531.
- [54] Boyd, S., Busseti, E., Diamond, S., Kahn, R. N., Koh, K., Nystrup, P., & Speth, J. (2017). Multi-period trading via convex optimization. *Foundations and Trends in Optimization*, 3(1), 1–76.
- [55] Moallemi, C. C., & Sağlam, M. (2017). Dynamic portfolio choice with linear rebalancing rules. *Operations Research*, 65(2), 294–314.
- [56] Bertsimas, D., Gupta, V., & Kallus, N. (2018). Data-driven robust optimization. *Mathematical Programming*, 167(1–2), 235–292.
- [57] Easley, D., López de Prado, M., & O'Hara, M. (2012). Flow toxicity and liquidity in a high-frequency world. *Review of Financial Studies*, 25(5), 1457–1493.
- [58] Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2), 223–236.
- [59] Patton, A. J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1), 246–256.
- [60] Giglio, S., Kelly, B., & Xiu, D. (2022). Factor models, machine learning, and asset pricing. *Annual Review of Financial Economics*, 14, 337–368.
- [61] 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.
- [62] Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267, 1–38.
- [63] Barocas, S., Hardt, M., & Narayanan, A. (2023). *Fairness and machine learning: Limitations and opportunities*. MIT Press.
- [64] Föllmer, H., & Schied, A. (2016). *Stochastic finance: An introduction in discrete time* (4th ed.). De Gruyter.
- [65] Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2(3), 21–41.
- [66] Pflug, G. C., & Pichler, A. (2014). *Multistage stochastic optimization*. Springer.
- [67] Avellaneda, M., & Stoikov, S. (2008). High-frequency trading in a limit order book. *Quantitative Finance*, 8(3), 217–224.
- [68] Buehler, H., Gonon, L., Teichmann, J., & Wood, B. (2019). Deep hedging. *Quantitative Finance*, 19(8), 1271–1291.
- [69] Glasserman, P. (2003). *Monte Carlo methods in financial engineering*. Springer.
- [70] Simon, H. A. (1996). *The sciences of the artificial* (3rd ed.). MIT Press.
- [71] March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71–87.
- [72] Teece, D. J. (2018). Business models and dynamic capabilities. *Long Range Planning*, 51(1), 40–49.
- [73] Campbell, J. Y., Lo, A. W., & MacKinlay, A. C. (1997). *The econometrics of financial markets*. Princeton University Press.