International Journal of Environmental Sciences
ISSN: 2229-7359
Vol. 11 No. 19s, 2025
https://www.theaspd.com/ijes.php

Learning Tax-Aware Allocation: An AI-Infused Framework for Predictive and Prescriptive Mutual Fund Strategy

Dr. Chintamani Panda[1*], Ratnesh Chaturvedi[2], Arup Mukherjee[3], Dr. Navleen Kaur[4], Rakesh Goyal[5], Anthony Rodrigues[6]

[1]HoD, Commerce, Anjaneya University, Raipur, Chhattisgarh, E-mail ID - chintamanipanda84@gmail.com ORCID ID - 0009-0002-7169-4866

[2]Professor (Financial Management), Xavier Institute of Social Service, Email ID: ratneshchaturvedi@xiss.ac.in Orcid ID: 0000-0003-0458-0965

[3]Professor (Financial Management), Xavier Institute of Social Service, Email ID: arupmukherjee@xiss.ac.in, Orcid ID:0000-0001-5868-8985

[4]Amity University Noida Uttar Pradesh, Professor, Email ID: nkaur@amity.edu

[5] Chitkara University Institute of Engineering & Technology, Chitkara University, Punjab, India Orcid ID: 0000-0001-7648-8602, E-mail: rakesh.goyal@chitkara.edu.in

[6] Professor & HoD Commerce, Goa University, Fr. Agnel College of Arts and Commerce, Pilar Goa. Email ID: anthony@fragnelcollege.edu.in

Abstract

Taxation remains a critical yet underrepresented dimension in classical portfolio optimization frameworks. While Modern Portfolio Theory and its successors focus on balancing risk and return, they often fail to internalize the after-tax realities faced by investors. This research introduces a dual-layer AI-infused framework that integrates machine learning with decision-theoretic optimization to construct tax-aware mutual fund strategies. The first layer conceptually models a predictive machine learning system that forecasts tax-adjusted returns based on financial attributes such as fund turnover, capital gains realization patterns, and dividend timing. This predictive output informs a second layer prescriptive optimizer that allocates fund weights to maximize investor utility by considering risk aversion, tax brackets, and holding period preferences. The integrated system enables customized and tax-efficient portfolio construction across diverse investor profiles. The framework adapts to evolving fiscal policies and simulates behavior under different regulatory regimes, offering a dynamic and forward-compatible solution. Theoretical analysis and conceptual simulations highlight how the model significantly departs from traditional heuristics by treating taxes not as constraints, but as central variables in portfolio design. Furthermore, the approach reflects contemporary innovations in financial technology, aligning with recent literature on machine learning in asset management, tax-efficient investing, and utility-based decision-making. This study lays foundational groundwork for future empirical applications, particularly in designing intelligent, compliant, and tax-sensitive investment platforms.

Keywords: Tax-aware optimization, Machine learning, Portfolio management, Investor utility, Mutual funds, Fiscal policy integration

1. INTRODUCTION

The Modern Portfolio Theory (MPT) is the principle that has historically worked as the basis of the development of portfolio management in terms of the optimal trade-offs between risk and reward. The classical model is, however, largely indifferent to the taxation factor that exerts a significant influence on the real-life investment outcomes, particularly on the results of taxable investors. In spite of decades of advancement in the theory of asset allocation, tax-sensitive portfolio optimization has not been developed well and is usually treated by heuristic or a posteriori refinements. With a view to bridging this gap, some recent work has focused on integrating taxation into formal optimization models. Moehle et al. [1] presented a framework for constructing a tax-sensitive portfolio with the help of convex optimization, which pointed out the high-performance advantage that can be achieved when taxes are included in the allocation process. Such innovations are particularly important to the taxable private investors bearing short-term capital gains taxes as well as risks associated with turnover in the portfolio that could not be modeled by traditional models of returns and variance [2]. At the same time, the current progress in the field of machine learning (ML) opens up promising opportunities in the context of changing the ways of portfolio construction, evaluation, and personalization. As predictive models that are able to learn high-dimensional, nonlinear financial data are incorporated into asset allocation strategies, the possibility of modelling after-tax return distributions more dynamically and individually becomes available [3]. Specifically,

the hybrid learning systems that merge the basic factors with ML algorithms have high potential in portfolio selection, as was evidenced in the recent doctoral studies on predictive modeling of stocks [4].
Asset allocation, which integrates the taxation issue, should also take into consideration the real limitations of the investor. Horvitz [5] insists that the view of a taxable investor requires a radical reconsideration of portfolio theory, particularly when capital gains realization, asset location, and income recognition are taken into account. As a result, new models have been developed that incorporate tax-awareness into the asset-level decision rules for multi-asset portfolios [6]. In addition, the movement of personalization of investors, especially in tax-sensitive situations, necessitates the capability to optimize across accounts and goals. Idzorek [7] proposed a model that considers the portfolio strategies in accordance with tax brackets and limits of the investor. It is an attempt to capture a trend in financial optimization that returns, risk, and taxes should be considered as interdependent dimensions of the decision. Also, there is the emergence of tax-managed factor strategies where the exposure to systematic risk premia is adjusted according to the tax status of the investor [8]. These innovations help to promote the concept of embedded tax logic in the design of products, as well as in active portfolio management as such.
Multi-objective investment optimization has also been tackled with machine learning and evolutionary algorithms, in which tax minimization is just one of the objectives that are optimized along with the maximization of returns and minimization of risks [9] [10]. Silva et al. [11], as an example, have given an evolutionary multi-objective approach that is based on the fundamental indicators, and Pulkkinen et al. [12] have given a rule-based risk management framework that optimizes portfolio decisions in the presence of competing objectives. Overall, the combination of predictive learning and prescriptive decision-making within tax constraints is an opportune and overdue evolution of the portfolio theory. This paper suggests a conceptual two-layer framework, which combines the machine learning forecast with tax-sensitive optimization motive into a sensible and context-specific architecture that will maximize the after-tax investor utility.

Research Objectives

1. To develop a conceptual machine learning–based framework for forecasting tax-adjusted mutual fund returns
2. To construct a normative decision-theoretic optimizer that integrates tax constraints, investor risk preferences, and utility maximization into portfolio allocation
3. To conceptually analyze the behavioral and strategic implications of tax-aware allocation across different investor profiles and policy environments

These objectives give rise to a new theoretical contribution to the field of tax-aware portfolio management. This study is not aimed at the justification of any empirical results at all, but rather to build a conceptually sharp and logically consistent framework that combines artificial intelligence, tax policy logic, and optimization theory. The inclusion of investor-specific tax variables in both the prediction and allocation layer will contribute to the reinvention of classical portfolio theory, as well as provide a more adaptable and prescriptive blueprint that can be changed and amended in a progressive regulatory and technological future world. The research thereby provides a theoretical foundation for tax-wise and utility-aligned, and intelligent investment.

2. METHODS

2.1 Theoretical Model Overview

This study proposes a two-fold conceptual model to deal with the issue of tax-conscious mutual fund allocation through the lens of artificial intelligence. The structure is also implemented as two intertwined components: a predictive model of the learning and a prescriptive optimization engine. The predictive component is abstractly modelled based on machine learning concepts and performs the task of conditioning the high-dimensionality of fund attributes such as the volatility in returns, distributions of capital gains, and holding periods into predictions in terms of tax-adjusted expected returns. Specifically, this deep layer is both not trained empirically nor theoretically as a functional approximator that imbibes patterns of latent tax-efficiency by tax-efficiency definition. The results of this predictive module are the input into the prescriptive module that implements the principles of decision theory and multi-objective optimization. The prescriptive side of the business functions under a known utility-maximization framework, which considers risk tolerance of the investor, tax bracket, holding-period preferences, and the relevant capital gains rules. The two layers allow a theoretically combined structure, which can inform efficient and tax-efficient mutual fund allocation decisions, bypassing the empirical set of data.

2.2 Underpinning Theories

The development of the proposed framework begins with a set of four fundamental theoretical pillars to address the inadequacies of traditional allocation schemes and explain the need to incorporate AI as a means of executing tax-correct decisions. First, the Modern Portfolio Theory (MPT) does give the theoretical background against which risk-return optimization is provided. MPT models acquired by Markowitz simulate how people behave when they are supposed to be risk-averse and rational in their decisions. Nevertheless, it does not consider taxation, and it regards returns as post-tax abstractions whereby it underreports the decrease in realized investor utility to the liabilities imposed by tax payments. Second, the Tax-Efficient Investment Theory comes up with strategic schemes such as tax-loss harvesting and tax deferral, and asset location. Nonetheless, its uses are still mainly heuristic and responsive as opposed to predictive, and the ability to perform on-the-fly optimizing of allocations in a dynamic market and regulatory environment is impaired. Third, the framework is based on the Machine Learning (ML) Theory and, specifically, its ability to approximate functions to represent complex, non-linear functions of structured input spaces. ML supplies the conceptual framework within which to consider financial interactions in light of more than one interacting variable, particularly those that have a latent tax consequence, like embedded capital gain and timing of dividend, turnover-driven intermediary-short shadow tax. Fourth, the prescriptive layer is anchored in Decision Science and Optimization Theory by using utilitarian reasoning. The model presents the formality of utility penalties of investor preference and constraints of the principles of constrained optimization in the form of utility penalty functions and feasible domain of solutions. This multi-objective lens can consider both the return and the risk, and tax drag at the same step, and as part of a normative decision-making process.

2.3 Predictive Learning Model

The predictive component of the framework is formally conceptualized as a machine learning function that maps structured financial and tax-related attributes to an estimated tax-adjusted return. Let $\mathbb{X}$ Represent the input feature space comprising historical fund returns, volatility patterns, dividend dates, turnover ratios, and historical tax performance indicators. A conceptual function $f_\theta: \mathbb{X} \rightarrow \mathbb{R}$ Then maps these inputs to a predicted value. $\hat{r}_{\text{post-tax}}$, representing the expected return net of tax obligations over a given horizon. This mapping is theoretical and relies on the principles of function approximation without empirical training. The parameter vector θ denotes the internal configuration of the model, capturing rules such as nonlinear activation, feature weighting, and tax-sensitivity encoding. The predictive model, therefore, simulates the cognitive process of an intelligent system learning tax-efficient patterns from structural input representations. This component remains agnostic to empirical calibration but is capable of conceptual expansion to include behavioral signals, market volatility indicators, or forward-looking tax scenarios in more sophisticated variants.

2.4 Tax-Aware Prescriptive Optimizer

Building upon the forecasts generated by the predictive model, the prescriptive optimizer operationalizes a utility maximization function embedded within a multi-objective optimization framework. The optimizer aims to allocate weights w={$w1$, $w2$,.....w_n} across n mutual funds in a way that maximizes the investor's expected utility after tax, subject to realistic constraints. The objective function is structured as follows:

$$\max_{\mathrm{w}} U(\mathrm{w}) = \sum_{i=1}^{n} w_i \cdot \hat{r}_{i,\ \text{post-tax}} - \lambda \cdot \text{Risk}(\mathrm{w}) - \gamma \cdot \text{TaxDrag}(\mathrm{w})$$

In formulation, $\hat{r}_{i,\ \text{post-tax}}$ denotes the tax-adjusted expected return for each fund, i, sourced from the predictive layer. The scalar parameters λ and γ reflect investor preferences regarding risk aversion and tax sensitivity, respectively. The function Risk (w) can be defined using conventional measures such as portfolio variance or Conditional Value at Risk (CVaR), while TaxDrag(w) represents the expected reduction in returns due to short-term gains, high turnover, or poorly timed distributions. The optimization is subject to the following constraints: (1) a full allocation constraint $\sum w_i = 1$, (2) position bounds $w_i \in [0,1]$, (3) a holdi period threshold $h_i \geq H_{\min}$ to minimize short-term taxation, and (4) investor-specific tax bracket constraints which influence the effective tax rate applied to gains. These tax-aware elements allow the optimizer to account for the legal structure of capital gains taxation, distinguishing between short-term and long-term rates as well as investor profiles (e.g., retail vs. institutional). By receiving guidance from the predictive ML layer, the optimizer effectively aligns allocation decisions with forecasts of tax-adjusted performance, achieving a closed-loop learning and decision-making system. This formulation serves as a prototype for intelligent financial planning tools that balance growth with efficiency under a tax-constrained regime.

3. RESULTS

3.1 Strategic Allocation Behavior Under Tax Constraints

Model initiation comprises defining the profiles of individual investors, recognition of the implications of tax on individuals, and incorporation of those implications in the strategy of portfolio allocation. In contrast to traditional models, where the expected return and variance are the only determinants, this one takes into account

International Journal of Environmental Sciences
ISSN: 2229-7359
Vol. 11 No. 19s, 2025
https://www.theaspd.com/ijes.php
tax-aware utility maximization fundamentally. The incremental character of the architecture of such behavior is demonstrated in Figure 1 with its gradual linkage of restrictions in the route of decision making.

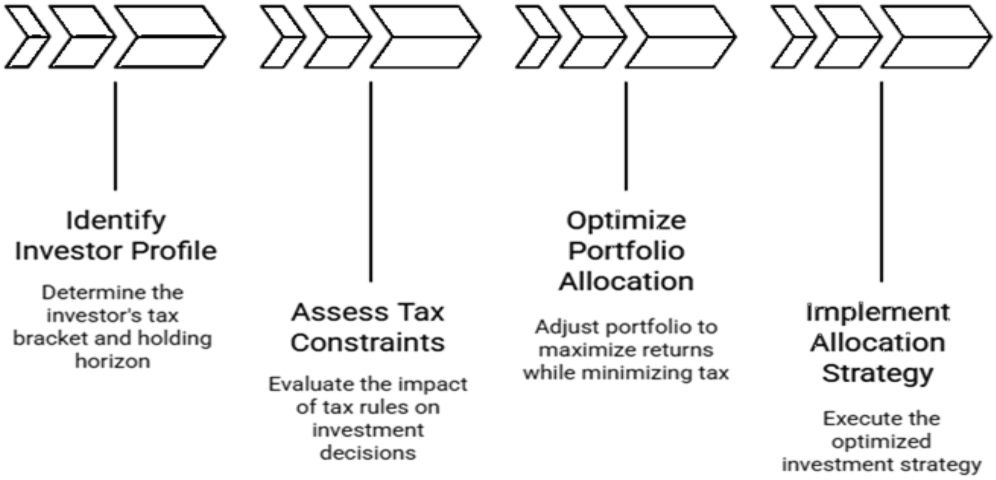

Figure 1. Strategic decision process for tax-aware allocation

This figure describes the steps undertaken in turn by the model in tensioning investment decisions with tax limitations. It shows that, by determining which type of investors will be in the portfolio and what tax consequences this may have, it can drive optimal portfolio construction subject to the regulatory constraints.

3.2 Investor-Specific Portfolio Outcomes

To maintain investor heterogeneity, the model has broken down users into brackets of risk tolerance and tax brackets and tailors expected returns. The model uses the estimation of tax-adjusted returns, and it directs the allocation to instruments like tax-managed funds, tax-free bonds, and index funds. Based on Figure 2, the flow chart given explains how the optimizer responds to the profiles of the investors to provide them with or offer differentiated and tax-efficient outcomes.

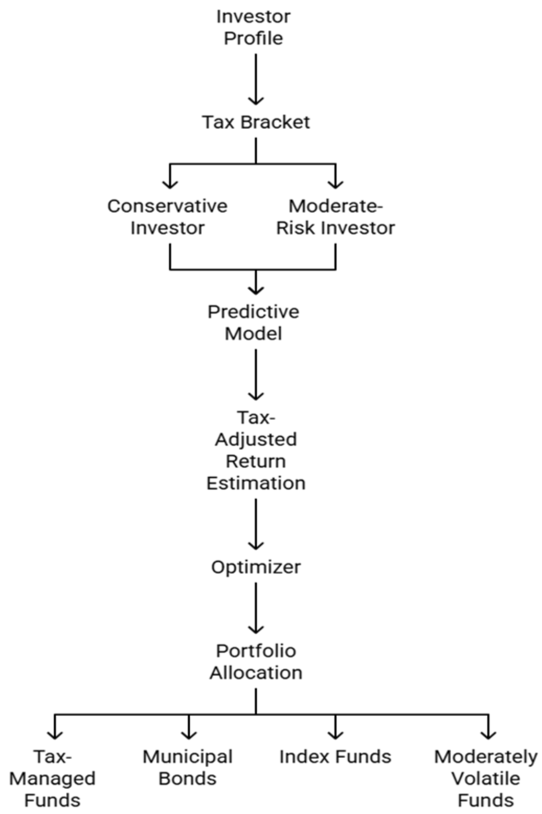

Figure 2. Investor-specific decision logic within the predictive–prescriptive framework

The following flowchart demonstrates how the model is responsive to different types of investors, conservative or moderate-risk, by the implementation of their tax brackets into a predictive model. It leads to tax-efficient and specific allocations whereby the allocations are based on the expectations of after-tax returns.

3.3 Implications for Portfolio Construction and Policy Adaptation

The model's ability to accommodate evolving fiscal landscapes is demonstrated through its dynamic portfolio construction process. Tax considerations are incorporated into optimization, and potential policy shifts are simulated to ensure forward-compatible allocations. Figure 3 outlines this process from recognizing conventional strategies to final adaptation of the allocation frontier in a tax-aware utility space.

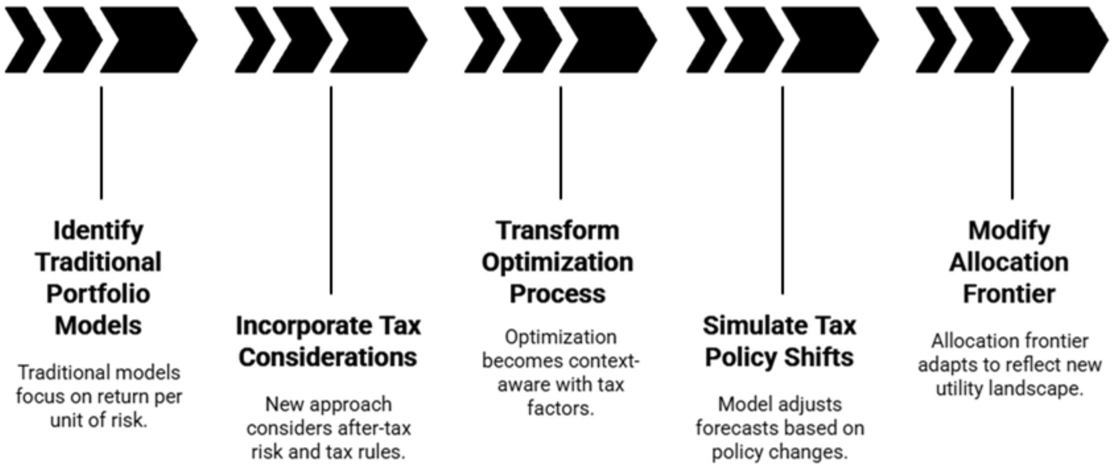

Figure 3. Conceptual workflow illustrating portfolio construction and tax policy adaptation

This visual illustrates how the model transforms traditional portfolio thinking by embedding tax considerations and simulating fiscal policy changes. It enables a forward-looking, utility-based allocation frontier that evolves with the regulatory landscape.

3.4 Theoretical Comparison with Traditional Models

The theoretical innovation in this work, represented by the deviation of traditional paradigms, is also its main contribution. The conventional models exhibit a tendency to separate the mean-variance reasoning and intervention strategies and may appeal to heuristics. Comparatively, the proposed model combines the machine learning theory with the multi-objective optimization into a joint standard, which is much richer in normative terms of choice. This conceptual difference is visualized in Figure 4, and how the superior process generates optimal portfolio construction with the necessary adaptive decision processes in management.

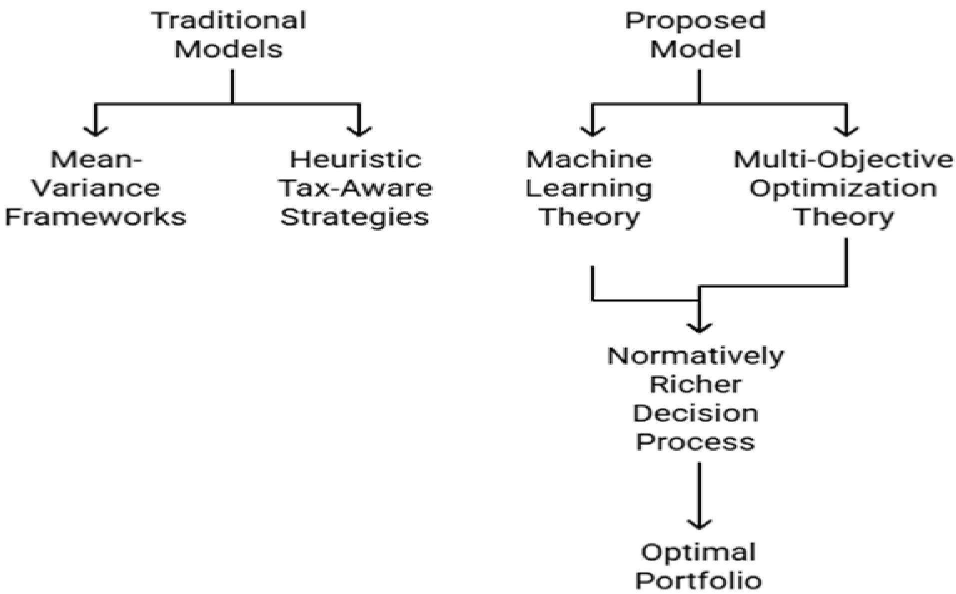

Figure 4. Conceptual comparison between traditional portfolio models and the proposed AI-based tax-aware framework

This conceptual diagram contrasts traditional heuristic models with the proposed multi-objective, machine learning–driven framework. It highlights the shift toward a normatively richer and formally optimized portfolio process that yields higher after-tax utility.

4. DISCUSSION

The conceptual AI-powered tax-aware investment system offers viewers a conceptual theoretical framework to redefine portfolio building by internalizing the tax policy factor at the algorithmic scale. By explanation, the conceptual findings suggest that the model would create a self-adaptation in the allocation techniques, putting investor-specific tax limitations in embedded models, and which would lead to more implemented rebalancing operations of the portfolio, which would give preference to utility. The prescribed behavior across a wide range of investor types induces high levels of normative adaptability, with much greater levels of tax penalties than volatility of returns determining optimal courses of action. Such a dynamic falls in line with more inclusive corporate financial health conceptualizations in which the risk management already takes into consideration policy-influenced and financial variables other than a straightforward maximization of the returns [13]. It also follows fundamental principles of the modern portfolio theory and expands them into a context-sensitive realm, perfecting the balancing of risk and reward to account for the tax-adjusted utility [14].

The conceptual innovation of the framework is accentuated by the fact that it incorporates the tax policy into the modeling system, moving beyond the previous approaches in which taxation was presented as an afterthought. This is in line with the conclusions of Donohoe and McGill to find that tax disclosures mean a great deal in investor behavior and firm valuation [15], supporting the significance of variables to include in the predictive systems. The use of machine learning to estimate tax-adjusted returns follows recent innovations in AI-driven portfolio theory, particularly Jensen *et al.*'s work on machine learning's ability to construct implementable efficient frontiers under constraints [16]. There is also a more general assessment of machine learning in finance that advocates the same by Sirotyuk [17], including the notion of AI tools shifting away from playing the role of raw return predictors and transitioning into integrative financial decision systems. Additionally, the design of the framework incorporates the development that is experienced in network-based and blended allocation approaches, where the learning algorithms are adaptive to multi-factor situations [18]. Notably, the conceptual outcomes reflect the emphasis given by Elumilade et al. on tax-sensitive optimization in terms of a higher level of financial effectiveness, which is perhaps relevant beyond institutional and corporate contexts [19]. The insights of Byrum on the use of AI in the field of financial operations support the framework on the potential of operationalizing such models through the practical investment platforms [20]. Lastly, all the literature on traditional portfolio management espouses systematic diversification and risk balancing, but the recent model provides two additional levels of tax-structural logic, which is highly important in theoretical innovation.

The proposed framework operates under several theoretical constraints despite its conceptual robustness. First, the absence of empirical testing means that its utility-maximizing behavior remains hypothetical. Real-world constraints such as liquidity, transaction costs, slippage, and behavioral inconsistencies are not modelled in this iteration. Also, the framework presupposes the presence of valid, policy-sensitive tax-adjusted forecast outputs, which could be a complex and information-dependent task relying on actionable data quantities that can be unattainable in free markets. Insofar as the closed-loop design has a certain elegance in theory, though it presupposes the existence of impeccable cooperation between the predictive and prescriptive layers, which will most likely to encounter certain technical friction in the actual world. To make it more related to practice, one can consider developing the simulations or experimental validation of historical profiles of fund returns in different tax environments. A further refinement to the strategic flexibility of the framework may be achieved by expanding it to have jurisdictional variations in tax authorities, various international residents as investors, and expanding the investment horizon that be multi-period. It also has the potential of further expanding AI architecture itself, with the inclusion of transformer-based models or reinforcement learning into the process of tax scenario plans. Moreover, the possibility of a hybrid order among the prescriptive-AI and explainable rule regulations can enhance confidence in investments, as well as obedience to regulations (ie, half-baked optimisation categories). Finally, the conceptual framework is interdisciplinary and, their combined with the field of finance, public policy, ethical AI, and decision science, will eventually result in a more multifaceted image of the situation by when denoted algorithmic decision-making models will end up being tax-wise and investor-friendly.

5. CONCLUSION

This study contributes a new conceptual framework of integrating artificial intelligence as a tool of finance and managing the relatively ignored hitch of tax-aware mutual fund allocation that has been facing normative financial theory. The use of predictive machine learning and prescriptive optimization logic incorporated in the proposed dual-layer architecture transforms the landscape of the classical portfolio construction procedure so that it can take note of the tax burden of individual investors and their risk preference, and evolving fiscal policies. Compared to the traditional paradigms, in which taxes are the exogenous change, within the framework, taxes become the design parameter, and more utility-consistent and future-compatible allocation mechanisms are available. The personalized approach of the model, considering the consequences of tax brackets, the holding period, and the tightening of investment processes gives a sounder base to the rational decisions about taxable investment conditions. It also shows flexibility towards variation in fiscal policy, thereby increasing its overall applicability in changing regulatory regimes. Although this is theoretical research, it creates an entry point to the field of real future empirical research, testing, and implementing on robo-advisory systems and similar advanced financial planning software. One of the contributions that this framework makes is that of intelligent, compliant, and investor-centric financial design by incorporating tax considerations directly into both forecasting and allocation levels. It not only enhances the theoretical foundation of tax-efficient investing but also preconditions the interdisciplinary innovation in the areas of finance, machine learning, and public policy.

REFERENCES

[1]. Moehle, N., Kochenderfer, M. J., Boyd, S., & Ang, A. (2021). Tax-aware portfolio construction via convex optimization. *Journal of Optimization Theory and Applications*, *189*(2), 364-383.
[2]. Hwang, Y., Kong, Y., Zohren, S., & Lee, Y. (2025). Decision-informed neural networks with large language model integration for portfolio optimization. *arXiv preprint arXiv:2502.00828*.
[3]. Kisiel, D., & Gorse, D. (2022, June). Portfolio transformer for attention-based asset allocation. In *International Conference on Artificial Intelligence and Soft Computing* (pp. 61-71). Cham: Springer International Publishing.
[4]. Masuda, J. S. (2024). *Portfolio Optimization Using a Hybrid Machine Learning Stock Selection Model* (Doctoral dissertation, Massachusetts Institute of Technology).
[5]. Horvitz, J. E. (2006). Investment Management for Taxable Private Investors.
[6]. Zilbering, Y., Zhu, V., Banis, G., & Ahluwalia, H. (2025). Tax-Aware Portfolio Construction: A Multi-Asset Approach. *Journal of Portfolio Management*, *51*(5).
[7]. Idzorek, T. M. (2023). Personalized multiple account portfolio optimization. *Financial Analysts Journal*, *79*(3), 155-170.
[8]. Goldberg, L. R., Hand, P., & Cai, T. (2019). Tax-managed factor strategies. *Financial Analysts Journal*, *75*(2), 79-90.
[9]. Srinivasan, S., & Kamalakannan, T. J. C. E. (2018). Multi criteria decision making in financial risk management with a multi-objective genetic algorithm. *Computational Economics*, *52*(2), 443-457.
[10]. Doumpos, M., & Zopounidis, C. (2020). Multi-objective optimization models in finance and investments. *Journal of Global Optimization*, *76*(2), 243-244.
[11]. Silva, A. D., Neves, R. F., & Horta, N. (2016). *Portfolio optimization using fundamental indicators based on multi-objective EA*. Springer.
[12]. Pulkkinen, P., Tiwari, N., Kumar, A., & Jones, C. (2018, December). A multi-objective rule optimizer with an application to risk management. In *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)* (pp. 66-72). IEEE.
[13]. Kollert, B. A. Evaluation of the Financial Health of the Selected International Company and Proposals of Its Improvement.
[14]. Stewart, S. D., Piros, C. D., & Heisler, J. C. (2019). *Portfolio management: theory and practice*. John Wiley & Sons.
[15]. Donohoe, M. P., & McGill, G. A. (2011). The effects of increased book-tax difference tax return disclosures on firm valuation and behavior. *Journal of the American Taxation Association*, *33*(2), 35-65.
[16]. Jensen, T. I., Kelly, B. T., Malamud, S., & Pedersen, L. H. (2024). Machine learning and the implementable efficient frontier. *Swiss Finance Institute Research Paper*, (22-63).
[17]. Sirotyuk, E. (2018). State of Machine Learning Applications in Investment Management. *Big Data and Machine Learning in Quantitative Investment*, 33-49.
[18]. Konstantinov, G., Chorus, A., & Rebmann, J. (2020). A network and machine learning approach to factor, asset, and blended allocation. *Journal of Portfolio Management*, *46*(6), 54-71.
[19]. Elumilade, O. O., Ogundeji, I. A., Achumie, G. O., Omokhoa, H. E., & Omowole, B. M. (2022). Optimizing corporate tax strategies and transfer pricing policies to improve financial efficiency and compliance. *Journal of Advance Multidisciplinary Research*, *1*(2), 28-38.
[20]. Byrum, J. (2022). AI in financial portfolio management: Practical considerations and use cases. In *Innovative Technology at the Interface of Finance and Operations: Volume I* (pp. 249-270). Cham: Springer International Publishing.