Universidade de Brasília

Faculdade de Economia, Administração, Contabilidade e Gestão de Políticas Públicas

Programa de Pós-Graduação em Economia

From Modern Portfolio Theory to Artificial Intelligence: A Review and Empirical Analysis

Autor: Rafael de Medeiros Roarelli

Orientador: PhD. Daniel O. Cajueiro

Brasília, DF

2025

Rafael de Medeiros Roarelli

From Modern Portfolio Theory to Artificial Intelligence: A Review and Empirical Analysis

Dissertação apresentada como requisito parcial para a obtenção do título de Mestre em Economia na linha de pesquisa de Economia Aplicada. Título em português: Da Teoria Moderna de Portfolio à Inteligência Artificial: Uma Revisão e Análise Empírica de Alocação de Ativo

Universidade de Brasília

Supervisor: PhD. Daniel O. Cajueiro

Brasília, DF

2025

Rafael de Medeiros Roarelli
From Modern Portfolio Theory to Artificial Intelligence: A Review and Empirical Analysis/ Rafael de Medeiros Roarelli. – Brasília, DF, 2025-
49 p. : il. (algumas color.) ; 30 cm.

Supervisor: PhD. Daniel O. Cajueiro

Dissertação – Universidade de Brasília , 2025.

1. Portfolio. 2. Machine Learning. I. PhD. Daniel O. Cajueiro. II. Universidade de Brasília. III. Programa de Pós-Graduação em Administração. IV. From Modern Portfolio Theory to Artificial Intelligence: A Review and Empirical Analysis

Classificação JEL: C22 B123

Rafael de Medeiros Roarelli

From Modern Portfolio Theory to Artificial Intelligence: A Review and Empirical Analysis

Dissertação apresentada como requisito parcial para a obtenção do título de Mestre em Economia na linha de pesquisa de Economia Aplicada. Título em português: Da Teoria Moderna de Portfolio à Inteligência Artificial: Uma Revisão e Análise Empírica de Alocação de Ativo

Trabalho aprovado. Brasília, DF, 15 de Outubro de 2025:

PhD. Daniel O. Cajueiro
Orientador

PhD. Hebert Kimura
Convidado 1

PhD. Douglas Sad Silveira
Convidado 2

Brasília, DF
2025

.

This dissertation is dedicated to my family, mentors, and friends who have been my steadfast support throughout this academic journey.

I would like to express my sincere gratitude to my advisor, Professor Daniel Cajueiro, for his invaluable guidance and encouragement. I also extend my thanks to the members of the examination committee for their insightful feedback.

To my parents, Arnaldo and Maria Liz, and my brothers, Gustavo and Otávio, I offer my deepest appreciation for your unwavering love and belief in me.

Finally, I dedicate this work to my friends, who have made this journey memorable. To my childhood friends, Henrique Melo, Leonardo, Vilela, Henrique Serra, Di Paula, Thiago, Carlos Eduardo, Fernando, and Gabriel, and to the friends I made during my master's program, thank you for the constant support and encouragement.

Acknowledgements

I would like to express my sincere gratitude to the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES) for the financial support that made this research possible.

I am also deeply grateful to the University of Brasília (UnB) for providing the academic environment and logistical support essential for the development of this dissertation

.

Abstract

Esta dissertação apresenta uma revisão abrangente da teoria de portfólios, traçando sua evolução desde os princípios fundamentais da Teoria Moderna do Portfólio (MPT) até a integração avançada de modelos de machine learning (ML). O estudo inicia-se com a análise do modelo de média-variância de Markowitz e sua posterior extensão para a Teoria Pós-Moderna do Portfólio (PMPT), que introduziu um conceito mais sofisticado de risco de queda (downside risk). O cerne deste trabalho aborda as limitações inerentes a esses paradigmas clássicos — como a sensibilidade aos parâmetros e as premissas estáticas — por meio de uma avaliação sistemática de soluções contemporâneas baseadas em ML. São revisadas técnicas como regularização para esparsidade de portfólio, deep learning para reconhecimento de padrões e reinforcement learning para otimização de políticas dinâmicas. Uma análise empírica é conduzida para comparar o desempenho de um portfólio tradicional de Markowitz com um modelo aprimorado que utiliza um algoritmo Random Forest para prever os retornos dos ativos.

Palavras-chaves: Teoria de Portfólios. Aprendizado de Máquina. Modelo de Markowitz. Alocação de Ativo.

Abstract

This dissertation provides a comprehensive review of portfolio theory, charting its evolution from the foundational principles of Modern Portfolio Theory (MPT) to the advanced integration of machine learning (ML) models. The study begins by examining Markowitz's mean-variance framework and its subsequent extension into Post-Modern Portfolio Theory (PMPT), which introduced a more nuanced concept of downside risk. The core of this work addresses the inherent limitations of these classical paradigms—such as parameter sensitivity and static assumptions—by systematically evaluating contemporary ML-driven solutions. Techniques including regularization for portfolio sparsity, deep learning for pattern recognition, and reinforcement learning for dynamic policy optimization are reviewed. An empirical analysis is conducted to compare the performance of a traditional Markowitz portfolio against an enhanced model that uses a Random Forest algorithm to forecast asset returns

Key-words: Portfolio Theory. Machine Learning. Markowitz Model. Asset Allocation

List of Tables

Contents

Introduction

During human history, the question of how to allocate resources has always been paramount in the lives of families, businesses, and governments. This enduring challenge is epitomized by the popular saying "Do not place all eggs in one basket." Far beyond its colloquial simplicity, this concept underpins modern financial theory, most notably through Markowitz's Modern Portfolio Theory (MPT). MPT formalized diversification as a mechanism to optimize risk-return trade-offs and paved the way for a dedicated branch of finance addressing the problem of efficient asset allocation.

The field of portfolio optimization continues to evolve dynamically, propelled by advances in computational power, econometric estimation, and artificial intelligence (AI). In light of these developments, this work conducts a comprehensive review of portfolio theory, charting its progression from foundational frameworks, such as Modern Portfolio Theory, to contemporary innovations that integrate machine learning (ML) algorithms. Central to this review is a systematic evaluation of ML-driven models, including reinforcement learning (RL) and deep neural networks (DNNs), which aim to address persistent limitations in traditional approaches, such as sensitivity to parameter estimation and non-linear risk modeling.

This study contributes by benchmarking the theoretical and empirical advantages of ML-driven models against classical paradigms. Specifically, it evaluates computational trade-offs, practical applicability, and the extent to which algorithmic advancements reconcile theoretical rigor with real-world financial complexities.

The existing body of literature on portfolio theory has attracted substantial attention, characterized by a multiplicity of literature reviews. For example, Kalayci, Ertenlice e Akbay (2019) conducted a comprehensive review of methods and algorithms for solving the optimization problem. Flint, Chikurunhe e Seymour (2015) focused on diversification measures, proposing a taxonomy that includes cardinality-based, weight-based, return-based, risk-based, and higher-moment measures. Lhabitant (2017) delved into technical aspects of portfolio diversification, addressing issues such as the household diversification puzzle and the diversification versus concentration dilemma.

On a theoretical level, Koumou (2020) explored diversification principles like the law of large numbers, correlation, the capital asset pricing model (CAPM), and risk contribution or risk parity, evaluating their utility for risk-averse investors. Meanwhile, Marhfor (2016) reviewed traditional and modern performance measures, proposing improvements and identifying new research avenues. More recently, Bartram, Branke e Motahari (2020) and Sutiene et al. (2024) turned their attention to AI-based models, examining their applications in portfolio management, trading, risk management, execution, and regulation.

Given the numerous ways to approach a literature review, this work will focus mainly on bringing to the attention of the reader the latest advances in machine learning driven models and compare it with the traditional Markowitz Model. This work is structured to systematically explore and compare two primary groups of portfolio optimization methodologies:, Modern Portfolio Theory (MPT) and Machine Learning (ML)-Driven Models.

The structure of this work is as follows: In Chapter 1, I will provide a brief introduction to portfolio theory. In Chapter 2, I will conduct an in-depth exploration of Modern Portfolio Theory (MPT), highlighting its core principles and inherent limitations. In Chapter 3, I will review Post-Modern Portfolio Theory, emphasizing its contributions to addressing the shortcomings of MPT. In Chapter 4, I will examine Machine Learning-driven models, presenting their theoretical foundations and recent innovations. In Chapter 5, I will perform an empirical analysis, comparing the traditional Markowitz model with a Random Forest model. Finally, in Chapter 6, I will conclude the work by summarizing the key findings, discussing the practical implications of the empirical results, and suggesting potential avenues for future research in algorithmic asset management.

1 Portfolio Theory: Brief Overview

The evolution of portfolio theory reflects the persistent quest of asset managers and researchers to systematize investment decision-making, moving from subjective, intuition-driven approaches to mathematically rigorous frameworks. Traditional portfolio management, rooted in individual stock evaluation, relied on non-systemic and analytically simplistic methods, often constrained by subjective biases and fragmented risk assessments. This paradigm shifted decisively with Markowitz's seminal work in 1952, which introduced Modern Portfolio Theory (MPT)—a groundbreaking framework that redefined diversification as a quantitative science. By analyzing portfolios holistically through Mean-variance framework, Markowitz demonstrated that investors could minimize variance for a given level of return, formalizing the risk-return trade-off, the cornerstone of modern finance.

However, MPT's own limitations, including estimation problems and simplifying assumptions such as variance as a complete risk measure, proved overly complicated and simplistic, respectively, in real-world markets. These limitations catalyzed the emergence of Post-Modern Portfolio Theory (PMPT), which reframed risk as asymmetric and investor-specific. By introducing the Minimum Acceptable Return (MAR) as a personalized benchmark, PMPT redefined risk as deviations below this threshold, distinguishing between "downside risk" (losses relative to MAR) and "upside volatility" (excess returns). This heterogenization of investor preferences and focus on downside protection marked a critical advancement, aligning portfolio optimization with behavioral realities (TODONI, 2015).

Nevertheless, neither MPT nor PMPT represents the end of portfolio theory's evolution. The advent of behavioral finance, computational power, and machine learning (ML) has further expanded the discipline's horizons both in theoretical and estimations aspects. Behavioral critiques have challenged rational investor assumptions, while ML-driven models now address non-linear risk dynamics, high-dimensional datasets, and adaptive market regimes—challenges that classical frameworks struggle to resolve. The following sections will provide an in depth discussion about both the theoretical aspects of each group and its limitations.

2 Modern Portfolio Theory

Its widely acknowledged that Markowitz work was instrumental in paving the way of thinking from the Modern Portfolio Theory. From creating new ways of pricing assets to extending the original, Markowitz model was able to influence many researchers that followed.

Markowitz model, as explained briefly in chapter 1, involved examining a group of stocks and its correlation in order to build an optimal portfolio. The idea is simple: stocks with low correlation are ideal to form an optimal portfolio.

The model is static. Markowitz (1952). Markowitz explains that, instead of using the time series of returns from the i_{th} security, he uses the flow of returns r_i to build his model. The expected return of a portfolio R is defined as the weighted sum of the expected returns of the assets R_i. Both the return of the asset and, consequently, the return of the portfolio are random variables. As such, we have the following equation:

$$R = \sum_{i=1}^{N} \alpha_i R_i$$

The variance of a portfolio is defined by the following equation:

$$V = \sum_{i=1}^{N} \sum_{j=1}^{N} \alpha_i \alpha_j \sigma_{ij}$$

The sum of weights are equal to 1 and all weights are greater or equal than zero. Markowitz original work did not allow the use of short sales an there is no risk-free asset, although later on, as we shall see, those restrictions have been relaxed.

The goal of the model is to find the weights of a portfolio that reach an efficient point (V, R) such that there won't be a point $X = (V_x; R_x)$ so that $R_x > R$ and/or $V_x < V$. In other words, we have an optimization problem where the goal is to minimize the variance of the portfolio subject to a given return. Thus, we have the following optimization problem:

$$\begin{array}{lrl} minimize & V = \sum_{i=1}^{N} \sum_{j=1}^{N} \alpha_i \alpha_j \sigma_{ij} & \\ subject & R & \geq \bar{R} \\ & \sum_{i=1}^{N} \alpha_i & = 1 \end{array}$$

The set of efficient points (or portfolios) is what is called efficient frontier, the famous "Markowitz Bullet". Later on, Tobin (1958) developed the separation theorem, adding the risk free asset to the set of investable assets and developping the Capital

Market Line, which is the Markowitz eficient frontier with the addition of the risk-free asset.

Markowitz's work exhibited certain limitations that have since motivated extensive research within the Mean-Variance framework to address these shortcomings. Section 2.1 focuses on the Capital Asset Pricing Model (CAPM), whereas Section 2.2 provides a brief overview of the multi-period approach and its subsequent extensions. Section 2.3 explores the Bayesian approach for portfolio models.

2.1 The CAPM model

The Capital Asset Pricing Model (CAPM), introduced independently by Sharpe (1964) and Lintner (1965)), and later extended by Black (1972), represents the foundational framework of modern portfolio theory. It establishes a linear relation between expected asset returns and their market risk (beta) and it is based on the on mean–variance optimization framework presented by Markowitz (1952).

While the portfolio model provides an algebraic set of results in asset weights in a risk-return framework, the CAPM model turns this set of results in a testable prediction about the relation between risk and expected return by identifying a portfolio that must be efficient if asset prices are to clear the market of all assets (FAMA; FRENCH, 2004). Sharpe (1964) and Lintner (1965) separately provided it by developing a general equilibrium approach for the asset prices.

Sharpe (1964) contribution to the Capital Asset Pricing Model (CAPM) was to formalize a general equilibrium framework in which asset prices adjust such that investors, through diversification, can achieve optimal portfolios lying along a capital market line. Assuming homogeneous expectations and a risk-free rate accessible to all investors for both borrowing and lending, Sharpe showed that the efficient frontier becomes a single linear frontier, tangent to the optimal portfolio of risky assets (SHARPE, 1964). This result allowed a two-step investment process: first, identify the market portfolio of risky assets, and second, adjust exposure to it via the risk-free asset. The risk of an individual asset, in this setting, is not assessed in isolation but rather in terms of its contribution to the variance of the market portfolio, encapsulated by its covariance with market returns.

Building on this, Lintner (1965) extended the CAPM by relaxing assumptions on investor homogeneity and allowing for differing subjective expectationa. He derived equilibrium conditions under which the expected return on any asset is linearly related to its covariance with the market portfolio, rather than its standalone variance. This nuanced distinction emphasized that within a diversified portfolio, idiosyncratic risk is negligible, and only systematic risk is priced. Lintner's model also introduced the notion of a "market price of risk," defined as the ratio of aggregate expected excess return to the variance of the market portfolio. He rigorously demonstrated that asset prices adjust such that the weighted sum of covariances governs their equilibrium valuations, making the CAPM an empirical statement about market-clearing prices rather than just a normative guideline.

Both Sharpe and Lintner's formulations converge on a critical insight: the expected return on a risky asset is a linear function of its beta—the asset's sensitivity to market returns. Their combined work shifted the focus of asset pricing from total risk to systematic risk and provided a manageable model for empirical validation and practical implementation in portfolio management and asset allocation.

2.2 The Continuous and Discrete Time Model

The static nature of Mean-Variance optimization fundamentally confines its applicability to single-period investment decisions. While Black (1972) acknowledged its validity primarily over infinitesimal intervals, subsequent work by Fama (1970) and Elton e Gruber (1974) explored the stringent conditions under which multi-period investment decisions could be reduced to a single-period framework, highlighting the inherent restrictions of such assumptions.

In response to these limitations, the relaxation of the single-period assumption primarily followed two distinct lines of work (CONSTANTINIDES; MALLIARIS, 1995). One approach developed discrete-time multi-period models through dynamic programming and efficient-market assumptions, while the other extended the framework into continuous time using stochastic calculus and equilibrium theory.

In the context of discrete-time models, Samuelson (1969) pioneered the intertemporal consumption–portfolio problem within a dynamic-programming framework. He demonstrated that optimal consumption and investment rules can be recursively characterized to maximize expected utility over a finite horizon. Building on Samuelson's insights, Hakanson (1970) extended this approach by deriving closed-form optimal consumption, investment, and borrowing–lending strategies for a broad class of utility functions. Hakanson notably showed that for certain preferences, the optimal consumption rule aligns with the permanent-income hypothesis.

Further contributions in discrete time include Rubinstein (1976) development of a valuation framework for uncertain discrete-time income streams, which led to the derivation of option-pricing formulas and illustrated how the Black–Scholes result emerges when trading occurs at discrete dates. Similarly, Jr (1974) examined stock prices, inflation, and the term structure of interest rates within a discrete-time intertemporal setting. His work illustrated finite-horizon effects on optimal portfolio policies and foreshadowed "turnpike" theorems, where optimal constant policies emerge as the investment horizon grows.

In this context, the development of continuous-time models using stochastic calculus provided more flexible and analytically tractable solutions. Merton (1969), Merton (1973) was instrumental in this area, formulating the consumption–portfolio problem via stochastic control. His work yielded closed-form solutions, particularly for Constant Relative Risk Aversion (CRRA) utility functions, demonstrating that the optimal risky-asset fraction remains constant over time and wealth. Merton's 1973 intertemporal CAPM (MERTON, 1973) further extended this, linking state-variable risk exposures to equilibrium asset betas in a continuous-time setting, thereby accounting for changes in investment opportunities.

Breeden (1979) further advanced the continuous-time framework with his Intertemporal Capital Asset Pricing Model (ICAPM), which related asset returns to instantaneous aggregate consumption risks, providing a continuous-time pricing kernel for general preferences. In his 1986 work Breeden (1986), Breeden synthesized consumption, production, inflation, and interest-rate dynamics, effectively bridging discrete- and continuous-time state-preference models. The comprehensive work of Jonathan E. Ingersoll (JR, 1987) in his "Theory of Financial Decision Making" unified stochastic-control and equilibrium pricing approaches across both discrete and continuous time, offering a thorough mathematical treatment of intertemporal consumption–investment problems and term-structure models.

2.3 Portfolio Model with Transaction Costs

Classical portfolio theory, from Markowitz's seminal mean-variance analysis to Merton's continuous-time model, is built on the assumption of frictionless markets. This assumption ignores the explicit costs associated with trading assets. Consequently, the presence of these costs fundamentally alters the investor's problem, inducing a reduction in transaction frequency.

Among the earliest attempts to integrate market frictions into portfolio theory, Pogue (1970) extended the Markowitz model to include transaction costs. His model accounted for two distinct cost components: proportional brokerage fees and the adverse price effects, or market impact, associated with executing large-volume trades.

Constantinides e Malliaris (1995) addresses the transaction cost issue in a continuous-time framework. The authors showed that the presence of transaction costs makes it suboptimal to continuously rebalance a portfolio. Instead, it is better to allow the portfolio's asset allocation to drift within a certain range around the target. This was a groundbreaking insight that laid the foundation for much of the subsequent research in this area.

Davis e Norman (1990) gave a more rigorous mathematical approach to the Constantinides e Malliaris (1995) framework. They modeled the problem as a two-dimensional process (with one risky asset and one risk-free asset), considered fixed transaction costs and showed that the optimal policy is characterized by a "no-transaction" region in the form of a wedge. When the portfolio is inside this wedge, the investor does nothing. When the portfolio hits one of the boundaries of the wedge, the investor should trade just enough to bring the portfolio back to the boundary. This minimal intervention is a key feature of the optimal strategy.

Shreve e Soner (1994) further generalized the model by considering both proportional and fixed transaction costs. They used viscosity solutions to solve the Hamilton-Jacobi-Bellman equations that describe the optimal control problem. Their work provided a more complete and rigorous solution to the problem of optimal consumption and investment in the presence of transaction costs.

2.4 The Bayesian Approach

The Markowitz model supposes that the return and correlations parameters are given. The consequence of this is that the model requires a point estimation for those parameters, resulting in estimation errors and unreasonable results, with the model often prescribing corner solutions with zero weight as well as large positions in assets of markets with small capitalizations (BLACK; LITTERMAN, 1992).

In this context, the Bayesian approach offers a solution by treating the parameters as something uncertain. Specifically, there are three main building blocks underlying a Bayesian framework (AVRAMOV; ZHOU, 2010). First is the formation of prior beliefs, which are typically represented by a probability density function on the stochastic parameters underlying the stock return evolution. Second, we have the formulation of the law of motion governing the evolution of asset returns. Lastly, using both the prior belief and the law of motion, we have the predictive distribution of future asset returns. The predictive distribution, which integrates out the parameter space, characterizes the entire uncertainty about future asset returns. The Bayesian optimal portfolio rule is obtained by maximizing the expected utility with respect to the predictive distribution (AVRAMOV; ZHOU, 2010).

The choice of priors is critical in this approach, and it can be either uninformative or informative. Intuitively, Uninformative priors are priors distributions that offer little information about the parameters. Informative priors are priors distribution that offers prior knowledge (information) about the parameters. In this context, Barberis (2000) examines the role of the parameter uncertainty when returns are predictable. Barberis shows that parameter uncertainty can lead to substantial differences in the optimal allocation to stocks in a long-horizon portfolio choice problem.

As an example. consider the following presented by Brandt (2010). For a single-period portfolio choice with i.i.d returns, we have the following problem:

$$\max_{x_t} \int u\left(x_t' r_{t+1} + R^f\right) p\left(r_{t+1} \mid \theta\right) \mathrm{d}r_{t+1}$$

The investor knows the parametric distribution but not the true parameter values. Given the data Y_T and a prior belief about the parameter $p_0(\theta)$, the posterior distribution can be given by the Bayes Theorem as:

$$p\left(\theta \mid Y_T\right) = \frac{p\left(Y_T \mid \theta\right) p_0(\theta)}{p\left(Y_T\right)} \propto p\left(Y_T \mid \theta\right) p_0(\theta)$$

The posterior distribution can then be used to obtain the investor's subjective return distribution:

$$p\left(r_{t+1} \mid Y_T\right) = \int p\left(r_{t+1} \mid \theta\right) p\left(\theta \mid Y_T\right) \mathrm{d}\theta$$

This then can be used in the investor problem. Formally, the investor solves the following problem:

$$\max_{x_t} \int u\left(x_t' r_{t+1} + R^f\right) p\left(r_{t+1} \mid Y_T\right) \mathrm{d}r_{t+1}$$

Which can be rewritten, giving the following problem:

$$\max_{x_t} \int \left[\int u\left(x_t' r_{t+1} + R^f\right) p\left(r_{t+1} \mid \theta\right) \mathrm{d}r_{t+1}\right] p\left(\theta \mid Y_T\right) \mathrm{d}\theta$$

The example shows that the investor, rather than solving the problem for a given single set of parameters, solves the problem by choosing a prior information about the parameters and then solving an average problem over all possible set of parameter values, where the expected utility of any given set of parameter values, is weighted by the investor's subjective probability of these parameter values corresponding to the truth.

Black e Litterman (1992) developed the model that can be considered the gold standard of the Bayesian approach in portfolio theory, dealing with the issues of extreme long and short positions mentioned earlier by assuming that the investor holds initial views of the market, then updates those views via the Bayesian rule (AVRAMOV; ZHOU, 2010). In particular, if the investor have views identical to the market, the chosen portfolio will be the market portfolio. The farther those views are from the market, the farther away his optimal portfolio will be from the market portfolio.

3 Post Modern Portfolio Theory

Post Modern Portfolio Theory is an effort to overcome some of those called simplifications in Markowitz analysis. As Leković (2021) points out, the main preoccupation of PMPT is related to risk and its measures. While Markowitz proposes the standard deviation as a tool to measure it, PMPT authors argue that upside risk is actually good considering that investor do not view a excessive return in a year in a negative way. Almost all see it in a positive way. In response to the shortcomings of classical risk metrics, proponents of Post-Modern Portfolio Theory (PMPT) developed new risk measures and reconceptualized portfolio optimization frameworks to account for asymmetric investor preferences.

Taking that into account, Downside Risk is defined as the risk that a return of an asset is below the return target. Sortino e Price (1994) also makes the distinction between uncertainty and downside risk: uncertainty is defined as the range of returns that an asset can obtain (which can be above or below someone's target return), while downside risk is a subset of uncertainty (that is, those returns that are below the target).

Some critics, as Sortino points out, argue that if one measures downside risk from the mean of each asset and the distributions of returns are symmetric, downside risk is simply semi-variance, and the results would be the same as the standard deviation. PMPT authors argue that returns actually do not necessarily have a symmetric distribution.

The following model presented by Huelin e Mirza (2011) will be used to present the Mean-Lower Partial Moment framework and explain the introduction of downside risk in portfolio analysis. The idea here is simple: create an alternative measure that captures the concept of downside risk and then incorporate in the problem of portfolio choice. First, Huelin and Mirza consider a downside risk averse investor with portfolio allocation X across k assets with $X' = (X_1, ..., X_k)$ and let F_X denote the probability distribution of returns on the portfolio. The set of feasible portfolios is

$$C = \left\{ \begin{array}{ll} X \mid \sum_i X_i = 1, X_i \geq 0 & without short sales, or \\ X \mid \sum_i X_i = 1 & with short sales \end{array} \right\}$$

As I noted before, the restriction $X_I \geq 0$ can be relaxed in a way that permits short sales. An investor will choose a portfolio from this set that maximizes the adjusted return for a given risk. R will represent the vector of returns with $R' = (R_1, ..., R_k)$.

With this, Huelin and Mirza, using the measure developed by Fishburn (1977), define the nth order lower partial moment of the distribution of returns under allocation X, computed at point τ (target rate), $LPM_n(\tau; X)$, where R_x is the return on the portfolio, as

$$LPM_n(\tau; X) = \int_{-\infty}^{\tau} (\tau - R_X)^n \, dF_X(R_X) = \int_{-\infty}^{\tau} (\tau - X'R)^n \, dF(R)$$

From this equation, we can notice the definition of downside risk expressed mathematically, with any point under the target rate being treated as risk. Bawa e Lindenberg (1977) use the risk free rate as a target rate and the problem becomes minimizing the $LPM_N(r_F, X)$ function. The model in question follows Harlow e Rao (1989) generalization of the LPM and include any possible target rate in order to account for different investor's preference.

Huelin e Mirza (2011) introduced the risk free rate as a proportion X_0 in the portfolio and redefined portfolio X to include risky assets and the risk free asset. The optimization problem then becomes as follows:

$$\min_X LPM_n(\tau; X) = \int_{-\infty}^{0} (\tau - X'R)^n \, dF(R) subjectto \left\{ X_0 R_F + \sum_{i \neq 0} X_i E\left(R_i\right) = \mu \right\} and \left\{ X_0 + \sum_{i \neq 0} X_i \in C \right\}$$

The risk averse investor in this model will select his portfolio by minimizing the downside risk for a specific value of the expected return μ. As in the MPT model, Bawa e Lindenberg (1977) shows that linear combinations of a portfolio X of risky assets and the risk free asset lie along a straight line in the mean-$LPM_n^{\frac{1}{n}}$. Given the convexity of the LPM_n function, we can find the optimal portfolio by drawing a tangent from R_F to $LPM_n(\mu)$, obtaining the market portfolio. Harlow e Rao (1989) obtained a similar result for any given target rate. Interesting to note, because the target rate is explicitly included in the calculation of downside risk, we can obtain an efficient frontier for any given target rate (ROM; FERGUSON, 1994). In other words, we can obtain an efficient frontier designed specifically for an investor desiring to reach a τ target rate.

The risk measure framework presented is not the only one developed. Estrada (2009) developed an alternative called the Mean-Gain Loss Spread Framework (GLS). The GLS takes into account three important index, namely the probability of a loss, the average loss and the gain. The idea is to incorporate into the measure information that investors find valuable when making investment decisions.

Another risk measure is the Conditional-Value at Risk (CVaR). CVaR is an extension of the Conditional-Value at risk and provides information considered complementary, given it measures the expected excess loss above the VaR, if a loss larger that VaR actually happens. In this context, a different approach to find the optimum portfolio was developed by Rockafellar, Uryasev et al. (2000). Basically, Rockafellar, Uryasev et al. (2000) built a loss function to define the CVaR, then they used Linear Programing to solve de optimization problem.

Another framework is the Mean-Drawdown framework (MD), developed by Chekhlov, Uryasev e Zabarankin (2005). The MD framework is based on the idea that not only a return below the target rate is important, but also the size of that loss. In this sense, an investor can consider the maximum consecutive loss he can tolerate. Taking this into account, Chekhlov, Uryasev e Zabarankin (2005) defines drawdown as "the drop in the portfolio value compared to the maximum achieved in the past". Different drawdown measures can be obtained (absolute and relative ones). In this sense, the differente drawdown measures can be generalized in one variable, namely the Conditional drawdown-at-risk (CDaR). The CDaR is based on a similar concept as the CVaR. In this sense, we can use the same approach as with the remaining risk measures and reduce the optimization problem to a linear programming problem (HUELIN; MIRZA, 2011).

4 Machine Learning Driven Models

MPT models are difficult to estimate mainly due to the method of estimating the covariance matrix and the non uniqueness of the solution. With the development of estimation through machine learning, many techniques were applied in the field of portfolio selection to try to overcome this difficulty. To catalog well the ways and strategies to estimate a portfolio through Machine Learning techniques, first we need to define Machine Learning. Rasekhschaffe e Jones (2019) defines Machine Learning as a term for methods and algorithms that allow machines to uncover patterns without explicit instructions. Specifically, in ML, we use data to train a model and then use the trained model to uncover those patterns.

Before delving into the methods, an observation must be made: Machine Learning is a field that has hundreds of algorithms and methods. In this sense, there are still many different methods being incorporated into portfolio selection to this day. As such, I will not cover all the methods and algorithms available, focusing solely on those already incorporated into portfolio theory. Also, many methods not published in journals are already being used as well in the financial industry, resulting in a gap between what is being done in the academy and what is being done in the financial industry (KRAUSS; DO; HUCK, 2017). As such, this work will focus solely on what has been published in journals.

The following sections systematically review key machine learning applications in portfolio management. First, section 4.1 addresses the foundational challenge of high-dimensionality and introduces statistical solutions. Next, section 4.2 explores regularization techniques like LASSO that enforce portfolio sparsity and robustness. Then, section 4.3 covers the use of supervised learning models to forecast asset returns. Finally, section 4.4 presents reinforcement learning, a paradigm for learning dynamic trading policies directly from market data.

4.1 High Dimensional Portfolio

A significant challenge that arises from the financial market is its high dimensional nature. This situation happens when the number of assets is much larger than the number of observations. In this context, traditional Mean-Variance approach has limitations because the covariance matrix becomes ill-conditioned and not invertible, leading to unstable and extreme results (LEDOIT; WOLF, 2004). The following subsections explore solutions to this problem.

4.1.1 Covariance Matrix Estimation: Shrinkage Estimators

Ledoit e Wolf (2004) developed a shrinkage approach. The idea is to "shrink" the unstable sample covariance matrix towards a highly structured and stable target matrix, known as the shrinkage target. This target could be based on a single-index model or even be a simple identity matrix. The result is a weighted average of the sample matrix and the target, which is always well-conditioned and invertible. This method effectively balances the trade-off between the bias introduced by the structured target and the variance reduction achieved by moving away from the noisy sample estimate. Subsequent research has focused on refining this process, for instance, by determining the optimal calibration of these shrinkage estimators for portfolio selection (DEMIGUEL; MARTIN-UTRERA; NOGALES, 2013).

4.1.2 Factors Models

Another method is the use of factors models. The idea here is to impose a factor structure for the covariation among assets to reduce the number of free parameters of the covariance matrix (BRANDT, 2010). By estimating this factor structure, the complexity of the covariance matrix is drastically reduced. The problem shifts from estimating a number of covariances to estimating the factor loadings and the covariance of the factors. Sharpe (1963) first proposed using the covariance matrix implied by a single-factor market model in the mean–variance framework problem (BRANDT, 2010). Subsequent work by Ross (1976), Ross (1977) expanded the single-factor model to a multi-factor one using the Arbitrage Pricing Theory (APT).

In a high-dimensional context, Fan, Fan e Lv (2008) demonstrated how to use a multi-factor model to estimate the covariance matrix, effectively reducing the number of parameters to be estimated from $\frac{2N}{(N+1)}$ to a function of the number of factors, which is much smaller. More recently, Caner, Medeiros e Vasconcelos (2023) have analyzed the performance of factor models in high-dimensional settings by proposing new regression techniques to handle the estimation challenges.

4.2 Regularized Optimization and Sparsity

This approach modifies the optimization objective itself by adding a penalty term to discourage complex or extreme portfolio weights. The goal is often to produce a sparse portfolio, meaning it invests in only a small subset of the available assets.

Brodie et al. (2009) showed that adding an ℓ_1 norm penalty (similar to the LASSO regression) to the Markowitz problem yields sparse and stable portfolios. The optimization problem can be formulated as:

$$\min_{w} w^T \Sigma w - \lambda w^T \mu \quad subjectto \quad \|w\|_1 \leq C$$

where w is the vector of portfolio weights and C is a constraint on the sum of the absolute weights. This regularization technique performs variable selection implicitly, making it well-suited for high-dimensional problems. Fan, Zhang e Yu (2012) extended this framework to handle vast portfolios with gross exposure constraints.

Chen, Dai e Zhang (2020) developed a mean-variance framework with the incorporation of sparse-group lasso (SGLasso) regularization in machine learning. SGLasso is an extension of the lasso designed to deal with predictors organized by group structure. The idea here is to penalize the Euclidean ℓ_2 norm of predefined groups of coefficients (MEIER; GEER; BÜHLMANN, 2008), enforcing a penalty on a group level. In the context of portfolio optmization, the set of assets can be naturally separated into subsets grouped by common characteristics such as economic sector, asset class, geographical location, etc. By applying the Group LASSO penalty to the portfolio weights, with groups defined by these sectors, the model encourages sparsity at the sector level (CHEN; DAI; ZHANG, 2020).

More advanced forms of regularization have also been explored to further enhance portfolio performance. Wu et al. (2024) developed mean-variance model with $\frac{\ell_1}{\ell_2}$ regularization to solve the sparse portfolio problem. Corsaro, Simone e Marino (2021) developed a fused Lasso approach for the multi-period portfolio selection problem in a Markowitz framework, applying two ℓ_1 penalties respectively on the portfolio weights, resulting in sparse solutions, and on the difference of wealth allocated across the assets between rebalancing dates, limiting the number of transactions.

Another advanced form is Graphical Lasso. While the others regularizations techcniques presented in this subsection focused on regularizing the weights, the graphical Lasso main focus is to regularize the inverse covariance matrix by maximizing the quasi-maximum likelihood, with a penalty on the sum of the absolute values of the off-diagonal elements. Goto e Xu (2015) used graphical lasso to construct a sparse precision matrix for portfolio hedging. Awoye (2016) uses the group lasso to estimate a covariance matrix for mean-variance portfolio optimisation. Millington e Niranjan (2017) uses graphic lasso to estimate the covariace matrix of assets returns.

To consolidate the different approaches discussed, Table 1 provides a summary of the key authors and the specific Lasso-based regularization technique they applied to the portfolio optimization problem:

Table 1 – Association of Authors with Lasso Techniques

Author(s)	**Lasso Technique**
(BRODIE et al., 2009), (FAN; ZHANG; YU, 2012)	- Lasso (ℓ_1 norm)
(CHEN; DAI; ZHANG, 2020)	- Sparse-group Lasso (SGLasso)
(MEIER; GEER; BÜHLMANN, 2008), (AWOYE, 2016)	- Group Lasso
(WU et al., 2024)	- $\frac{\ell_1}{\ell_2}$ regularization
(CORSARO; SIMONE; MARINO, 2021)	- Fused Lasso
(GOTO; XU, 2015), (MILLINGTON; NIRANJAN, 2017)	- Graphical Lasso

4.3 Deep Learning

The fundamental of machine learning problem is to find a predictor of an output Y given the input X (HEATON; POLSON; WITTE, 2017). A learning machine is an input-output mapping $Y = F(X)$ where the input space is high dimensional.

As a form of machine learning, deep learning trains a model on data to make predictions, but is distinguished by passing learned features of data through different layers of abstraction (HEATON; POLSON; WITTE, 2017). Specifically, deep learning approach employs hierarchical predictors comprising of a series of L nonlinear transformations applied to X. Each of these transformations are called layers, where the original input is X, the output of the first transformation is the first layer and so on. For a more detailed view of a deep learning structure, see Heaton, Polson e Witte (2017)

These structure provides a unique solution to the allocation problem. As hierarchical, non-linear function approximators, deep neural networks are uniquely capable of learning intricate patterns and high-dimensional interactions directly from large datasets without the need for manual feature engineering or pre-specification of the functional form. This feature is particularly interesting when applying it for a stock selection and asset allocation, given the capacity to uncover unconventional patterns from datasets.

Fischer e Krauss (2018) used Long-Short Term Memory (LSTM) network to predict stock movements in the S&P 500 index, demonstrating that LSTM models outperform traditional methods like logistic regression. Jiang e Liang (2017) presents a deep reinforcement learning approach to the portfolio problem consisting of the Ensemble of Identical Independent Evaluators (EIIE) topology, a Portfolio-Vector Memory (PVM), an Online Stochastic Batch Learning (OSBL) scheme, and a fully exploiting and explicit reward function, allowing continuous learning and adaptation to market changes. They tested its framework with three different neural network architectures to act as the "Identical Independent Evaluators" (EIIE). The networks they implemented and tested were Convolutional Neural Network (CNN), Recurrent Neural Network (RNN) and LSTM network.

Liang et al. (2018) proposed an "Adversarial Training" method which was found to greatly improve training efficiency and enhance average daily return and Sharpe ratio. Jang e Seong (2023) combined MPT with a deep reinforcement learning-based approach by using a tensor decomposition-assisted approach for multimodal data" details a multi-stage process. The core neural network used for feature extraction is a 3D CNN. This CNN processes a tensor of multimodal data. The output tensor from the CNN is then decomposed (using Tucker decomposition), and the resulting "core tensor" is fed into a final deep network (a Feedforward/MLP) to produce the final portfolio weights (JANG; SEONG, 2023).

Bühler et al. (2019) present a deep reinforcement learning approach to hedge a portfolio of derivatives with the presence of market frictions such as transaction costs, market impact, liquidity constraints or risk limits. They used a Feedforward Neural Networks (MLPs), where a separate network can be used at each discrete hedging point, and a Recurrent Neural Networks (RNNs), where a single network processes the time-series data and retains a memory of past states, which is more natural for path-dependent strategies and transaction costs.

Ma, Han e Wang (2021) incorporated return predictions into the portfolio formation using a suite of models, including two machine learning techniques—random forest (RF) and support vector regression (SVR)—and three deep learning architectures: LSTM, CNN and a Deep Multilayer Perceptron (DMLP)

Table 2 provides a consolidated summary of the primary neural networks employed by the key authors driving these innovations, highlighting the specific architectures used in their respective studies.

4.4 Reinforcement Learning

Reinforcement Learning (RL) reframes the portfolio optimization task as a sequential decision-making problem, providing an alternative to static, single-period models. In this approach, an autonomous agent learns to make optimal allocation decisions by interacting with the financial market, which acts as the environment. The agent observes the market's current state—comprising indicators, like price history and volatility, and performs an action, which consists of rebalancing the portfolio weights. After each action, the agent receives a reward, such as the period's return or a change in the Sharpe ratio. The objective of the agent is to learn an optimal policy, $\pi(s_t)$, a mapping from state to action, that maximizes the expected cumulative reward over a long horizon

Neuneier (1995) modeled the portfolio selection problem as a Markov Decision Process (MDP), creating a framework suitable for reinforcement learning. The author then successfully applied the Q-learning algorithm to determine an optimal investment strategy, demonstrating its effectiveness for solving problems with high dimensionality.

Table 2 – Neural Networks Used by Cited Authors

Author(s)	**Neural Network(s) Used**
Fischer & Krauss (2018)	- Long-Short Term Memory (LSTM)
Jiang et al. (2017)	- Convolutional Neural Network (CNN) - Recurrent Neural Network (RNN) - Long-Short Term Memory (LSTM)
Liang et al. (2018)	- Adversarial Training Method
Jang et al. (2023)	- 3D Convolutional Neural Network (CNN) - Feedforward Network (MLP)
Buehler et al. (2019)	- Feedforward Neural Networks (MLPs) - Recurrent Neural Networks (RNNs)
Ma et al. (2021)	- Long-Short Term Memory (LSTM) - Convolutional Neural Network (CNN) - Deep Multilayer Perceptron (DMLP)

Ormoneit e Glynn (2000) also adopted the Markov Decision Process (MDP) framework to develop a kernel-based reinforcement learning method. Their approach introduces an iterative process that updates the value function's estimates by calculating a local average, which is determined by a weighting kernel.

In contrast to these methods that focus on solving the MDP, Moody et al. (1998) and Moody et al. (1998) proposed the use of recurrent reinforcement learning (RRL) algorithm for training trading systems and managing portfolios. Their approach, which they call Direct Reinforcement (DR), treats investment decision-making as a stochastic control problem where strategies are discovered directly. This method bypasses the need to build forecasting models and instead focuses on directly optimizing objective functions that measure performance. These performance criteria include profit or wealth, economic utility, and risk-adjusted returns such as the Sharpe ratio. A key contribution of their work was the introduction of the differential Sharpe ratio, a novel metric designed to facilitate the efficient online optimization of risk-adjusted returns while accounting for transaction costs.

Gold (2003) applied RRL to currency trading. They compared the performance of single layer networks with networks having a hidden layer, and examined the impact of the fixed system parameters on performance. Hens e Wöhrmann (2007) applied RRL to bond and equity markets, assuming a rational investor with a constant relative risk aversion (CRAA) utility function. Maringer e Ramtohul (2010) argued that RRL is not sufficiently equipped to capture the non-linearities and structural breaks present in financial data. In this sense, they proposed what they called the threshold recurrent reinforcement learning (TRRL) model, which is the RRL model with the addition of regime switching. Table 3 summarizes the key studies discussed.

Table 3 – Reinforcement Learning Methods by Author

Author(s)	**Method**
(NEUNEIER, 1995)	Q-learning
(ORMONEIT; GLYNN, 2000)	Kernel-based Reinforcement Learning
(MOODY et al., 1998)	Recurrent Reinforcement Learning (RRL)
(GOLD, 2003)	Recurrent Reinforcement Learning (RRL)
(HENS; WÖHRMANN, 2007)	Recurrent Reinforcement Learning (RRL)
(MARINGER; RAMTOHUL, 2010)	Threshold Recurrent Reinforcement Learning (TRRL)

5 Empirical Analysis

This chapter compares the performance of a machine learning model to that of Modern Portfolio Theory (MPT). The structure is as follows: Subsection 5.1 outlines the dataset, Subsection 5.2 describes the Markowitz model, and Subsection 5.3 explains the Random Forest Forecast incorporated in the MPT model. Subsection 5.4 details the investment strategy. Finally, Subsection 5.5 presents the comparative results.

5.1 Dataset

The data set used in the analysis was taken from Yahoo Finance using the yfinance python library. One hundred assets were chosen at random from the S&P 500 list available online. The assets chosen appears on 5.

The data acquired spans from 2015 until 2023. Specifically, the code uses data from 2015-2021 to train the initial models and then tests the strategies on out-of-sample data from January 2022 to December 2023. The data acquired are the following: benchmark rate, risk free rate from the rebalancing date and the closing price. One observation must be made: if there are no available rate observed for the risk free rate, the last one obtained is the one considered for the day in question.

The forecasting model uses three technical features selected to capture the context, strength, and direction of an asset's price movements. Specifically, the model incorporates the following indicators: Daily Return, Momentum, and the Relative Strength Index (RSI). While this feature set provides a focused analysis of core price dynamics, it defines a specific scope for the forecast. For more comprehensive approaches that incorporate a broader range of variables, see, for example, Khaidem, Saha e Dey (2016).

5.2 Markowitz Model

The theoretical framework of the Markowitz model is established in section 2. This section details the methodology for estimating the market portfolio, a crucial input for the model. The market portfolio was derived by maximizing the Sharpe Ratio, which is geometrically equivalent to finding the tangency portfolio on the efficient frontier. The implementation of this optimization was carried out using the pypfopt library (MARTIN, 2021).

The expected return for each asset was calculated using an exponentially weighted moving average (EWMA). This methodology was selected over a simple historical average to ensure that more recent performance data, considered more relevant for forecasting, is given greater significance in the model. The covariance matrix was estimated using the Sample Covariance Matrix method.

5.3 Random Forest Forecast

The methodology presented here enhances the classic Markowitz portfolio selection model by replacing its static reliance on historical mean returns. Instead, a dynamic forecasting component is introduced using a Random Forest algorithm to predict future asset returns. This forecast is then used as an input in the Markowitz model to determine the optimum weights of the tangency portfolios.

The approach presented is based on the work of Breiman (2001) and Khaidem, Saha e Dey (2016).The first step in the estimation process is to have a function to forecast the expected return for each stock. For this, we have:

$$\hat{r}_{i,t+1} = f_i(\mathbf{x}_{i,t})$$

Where the $\hat{r}_{i,t+1}$ is the predicted return for stock i over the next period and $\mathbf{x}_{i,t}$ is is the feature vector for stock i at time t that includes the three features showed in subsection 5.1.

The functionf_i is the Random Forest Model. A Random Forest is an ensemble of many individual decision trees. Its final prediction is the average of all the individual tree predictions, which makes it accurate and stable. From this function, a return vector μ_{ML} is then generated, which will then be used as an input for the following problem:

$$max \quad S = \frac{R_p - r_f}{\sigma_p} = \frac{\mathbf{w}^T \mu_{ML} - r_f}{\sqrt{\mathbf{w}^T \mathbf{\Sigma} \mathbf{w}}}$$

subject to $\mathrm{w}_i \geq 0$
$\sum_{i=1}^{N} w_i = 1$

In this formulation, $\mathbf{w}$ is the vector of weights, μ_{ML} is the vector of predicted returns from the Random Forest model, Σ is the historical covariance matrix of the stocks and r_f is the risk free rate.

5.4 Investment Strategy

The objective of this analysis is to analyze the impact of Machine Learning forecast on the performance of a portfolio model. In this sense, no transaction costs are considered here.

This strategy employs a quarterly rebalancing schedule of 63 trading days, timed to coincide with the mandatory corporate earnings reporting cycle in the United States. For each rebalancing event, the model is retrained using a one-year rolling window of the most recent 252 trading days of data. This approach ensures the model continuously adapts to current market conditions and incorporates the latest corporate financial disclosures into its predictions.

5.5 Results

The analysis evaluates both models against the benchmark (SPY) using standard financial metric. The performance indicators are the following: Cumulative Return, Sharpe Ratio, Max Drawdown, Alpha and Beta. The results are summarized in table 6:

Table 4 – Performance Comparison: Markowitz vs. ML-Forecast Portfolios

Metric	**Markowitz Portfolio**	**ML-Forecast Portfolio**
Cumulative Return	23.19%	49.92%
Sharpe Ratio	0.41	0.61
Sortino Ratio	0.61	0.95
Max Drawdown	-21.02%	-34.12%
Beta	0.69	1.16
Alpha (annualized)	0.0909	0.2167

As expected, the ML-Forecast portfolio obtained substantially higher cumulative returns. This shows that the machine learning model's predictive inputs led to significantly more effective capital appreciation over the two-year period.

This outperformance is further emphasized by the portfolio's Alpha. The ML-Forecast portfolio generated an annualized alpha of 0.2167, indicating that it produced substantial excess returns beyond what would be expected for its level of market risk. In contrast, the traditional Markowitz portfolio's alpha was modest at 0.0909. This suggests the ML model provided a genuine predictive "edge" that was absent in the historically-driven model.

Regarding its volatility, the results show that the ML portfolio carried more risk related to the traditional model. Specifically, the Beta of 1.16 (compared to 0.69 from the traditional model) indicates that the ML-Portfolio was much more volatile than the market benchmark and the traditional model. This is further emphasized by the Max Drawdown of -34.12% compared to the -21.02% of the traditional model, additionally to the Cumulative Return pointed out. Still, the Sortino Ratio, which penalizes downside volatility, indicates that the ML-Forecast portfolio was particularly effective at converting risk into returns while managing harmful downside deviations.

One observation must be made about the limited scope of the analysis: only three features were considered in the Random Forest Model. The idea is to complement other studies, such as Khaidem, Saha e Dey (2016), and show a forecast with less features. Further studies must be made to analyse the accuracy of the forecast made here and compare the results. In this context, the performance results obtained here must be taken with caution.

6 Conclusion

This work has traced the evolution of portfolio theory, from its foundational principles to its current frontier. The evolution from Modern Portfolio Theory's (MPT) formalization of the risk-return trade-off to Post-Modern Portfolio Theory's (PMPT) more nuanced understanding of downside risk highlights a persistent drive to align financial models with the complex realities of investment decision-making. While these classical paradigms established the core of diversification, they also presented persistent challenges related to parameter estimation, static assumptions, and the modeling of non-linear market dynamics.

The integration of machine learning represents the next logical step in this evolution. As demonstrated throughout this review, ML techniques, ranging from regularization methods that enforce portfolio sparsity to deep learning and reinforcement learning models that uncover complex patterns, are not merely incremental improvements but a significant paradigm shift. These methods directly confront the limitations of earlier theories by offering data-driven solutions to high-dimensionality, non-linear risk, and adaptive market regimes.

The empirical analysis conducted in this study provides concrete evidence of this potential. By incorporating a Random Forest forecasting model into the traditional Markowitz framework, the resulting portfolio demonstrated that it can outperform its classical counterpart. The ML-Forecast portfolio achieved substantially higher cumulative returns (49.92% vs. 23.19%), superior risk-adjusted performance as measured by the Sharpe and Sortino Ratios, and an annualized alpha of 0.2167. These results suggest that machine learning can provide a genuine predictive edge, moving beyond historical averages to generate more effective capital allocation. Nevertheless, those results must be taken with caution, as observed in section 5.5

Ultimately, the findings of this work show that machine learning is an evolving and profoundly promising field within finance. It should not be viewed as a replacement for established financial theory, but rather as a powerful complement. ML offers a sophisticated toolkit capable of enhancing classical frameworks, refining empirical analysis, and developing more adaptive and robust investment strategies. The synergistic relationship between financial theory and machine intelligence will continue to redefine the boundaries of portfolio optimization, paving the way for the next generation of data-informed financial decision-making.

Bibliography

AVRAMOV, D.; ZHOU, G. Bayesian portfolio analysis. *Annu. Rev. Financ. Econ.*, Annual Reviews, v. 2, n. 1, p. 25–47, 2010. Citado 2 vezes nas páginas 23 and 24.

AWOYE, O. A. *Markowitz minimum variance portfolio optimization using new machine learning methods.* Tese (Doutorado) — (UCL) University College London, 2016. Citado 2 vezes nas páginas 30 and 31.

BARBERIS, N. Investing for the long run when returns are predictable. *The Journal of Finance*, Wiley Online Library, v. 55, n. 1, p. 225–264, 2000. Citado na página 23.

BARTRAM, S. M.; BRANKE, J.; MOTAHARI, M. *Artificial intelligence in asset management.* [S.l.]: CFA Institute Research Foundation, 2020. Citado na página 16.

BAWA, V. S.; LINDENBERG, E. B. Capital market equilibrium in a mean-lower partial moment framework. *Journal of financial economics*, Elsevier, v. 5, n. 2, p. 189–200, 1977. Citado na página 26.

BLACK, F. Capital market equilibrium with restricted borrowing. *The Journal of business*, JSTOR, v. 45, n. 3, p. 444–455, 1972. Citado 2 vezes nas páginas 20 and 21.

BLACK, F.; LITTERMAN, R. Global portfolio optimization. *Financial analysts journal*, Taylor & Francis, v. 48, n. 5, p. 28–43, 1992. Citado 2 vezes nas páginas 23 and 24.

BRANDT, M. W. Portfolio choice problems. In: *Handbook of financial econometrics: Tools and techniques.* [S.l.]: Elsevier, 2010. p. 269–336. Citado 2 vezes nas páginas 24 and 29.

BREEDEN, D. T. An intertemporal capital asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics*, v. 7, n. 3, p. 265–296, 1979. Citado na página 22.

BREEDEN, D. T. Consumption, production, inflation and interest rates: A synthesis. *Journal of Financial Economics*, v. 16, n. 1-2, p. 3–45, 1986. Citado na página 22.

BREIMAN, L. Random forests. *Machine learning*, Springer, v. 45, n. 1, p. 5–32, 2001. Citado na página 36.

BRODIE, J. et al. Sparse and stable markowitz portfolios. *Proceedings of the National Academy of Sciences*, National Acad Sciences, v. 106, n. 30, p. 12267–12272, 2009. Citado 2 vezes nas páginas 29 and 31.

BÜHLER, H. et al. Deep hedging. *Quantitative Finance*, Taylor & Francis, v. 19, n. 8, p. 1271–1291, 2019. Citado na página 32.

CANER, M.; MEDEIROS, M. C.; VASCONCELOS, G. F. R. Sharpe Ratio analysis in high dimensions: Residual-based nodewise regression in factor models. *Journal of Econometrics*, Elsevier, v. 235, n. 2, p. 105479, 2023. Citado na página 29.

CHEKHLOV, A.; URYASEV, S.; ZABARANKIN, M. Drawdown measure in portfolio optimization. *International Journal of Theoretical and Applied Finance*, World Scientific, v. 8, n. 01, p. 13–58, 2005. Citado na página 27.

CHEN, J.; DAI, G.; ZHANG, N. An application of sparse-group lasso regularization to equity portfolio optimization and sector selection. *Annals of Operations Research*, Springer, v. 284, n. 1, p. 243–262, 2020. Citado 2 vezes nas páginas 30 and 31.

CONSTANTINIDES, G. M.; MALLIARIS, A. G. Portfolio theory. *Handbooks in operations research and management science*, Elsevier, v. 9, p. 1–30, 1995. Citado 3 vezes nas páginas 21, 22, and 23.

CORSARO, S.; SIMONE, V. D.; MARINO, Z. Fused lasso approach in portfolio selection. *Annals of Operations Research*, Springer, v. 299, n. 1, p. 47–59, 2021. Citado 2 vezes nas páginas 30 and 31.

DAVIS, M. H.; NORMAN, A. R. Portfolio selection with transaction costs. *Mathematics of operations research*, INFORMS, v. 15, n. 4, p. 676–713, 1990. Citado na página 22.

DEMIGUEL, V.; MARTIN-UTRERA, A.; NOGALES, F. J. Size matters: Optimal calibration of shrinkage estimators for portfolio selection. *Journal of Banking & Finance*, Elsevier, v. 37, n. 8, p. 3018–3034, 2013. Citado na página 29.

ELTON, E. J.; GRUBER, M. J. On the optimality of some multiperiod portfolio selection criteria. *The Journal of Business*, JSTOR, v. 47, n. 2, p. 231–243, 1974. Citado na página 21.

ESTRADA, J. The gain-loss spread: A new and intuitive measure of risk. *Journal of Applied Corporate Finance*, Wiley Online Library, v. 21, n. 4, p. 104–114, 2009. Citado na página 26.

FAMA, E. F. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, v. 25, n. 2, p. 383–417, 1970. Citado na página 21.

FAMA, E. F.; FRENCH, K. R. The capital asset pricing model: Theory and evidence. *Journal of economic perspectives*, American Economic Association, v. 18, n. 3, p. 25–46, 2004. Citado na página 20.

FAN, J.; FAN, Y.; LV, J. High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics*, Elsevier, v. 147, n. 1, p. 186–197, 2008. Citado na página 29.

FAN, J.; ZHANG, J.; YU, K. Vast portfolio selection with gross-exposure constraints. *Journal of the American Statistical Association*, Taylor & Francis, v. 107, n. 498, p. 592–606, 2012. Citado 2 vezes nas páginas 30 and 31.

FISCHER, T.; KRAUSS, C. Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, Elsevier, v. 270, n. 2, p. 654–669, 2018. Citado na página 31.

FISHBURN, P. C. Mean-risk analysis with risk associated with below-target returns. *The American Economic Review*, JSTOR, v. 67, n. 2, p. 116–126, 1977. Citado na página 25.

FLINT, E.; CHIKURUNHE, F.; SEYMOUR, A. The cost of a free lunch: Dabbling in diversification. *Peregrine Securities*, 2015. Citado na página 16.

GOLD, C. Fx trading via recurrent reinforcement learning. In: IEEE. *2003 IEEE International Conference on Computational Intelligence for Financial Engineering, 2003. Proceedings.* [S.l.], 2003. p. 363–370. Citado 2 vezes nas páginas 33 and 34.

GOTO, S.; XU, Y. Improving mean variance optimization through sparse hedging restrictions. *Journal of Financial and Quantitative Analysis*, Cambridge University Press, v. 50, n. 6, p. 1415–1441, 2015. Citado 2 vezes nas páginas 30 and 31.

HAKANSON, N. H. Optimal investment and consumption strategies under risk for a class of utility functions. *Econometrica: Journal of the Econometric Society*, v. 38, n. 5, p. 587–606, 1970. Citado na página 21.

HARLOW, W. V.; RAO, R. K. Asset pricing in a generalized mean-lower partial moment framework: Theory and evidence. *Journal of financial and quantitative analysis*, Cambridge University Press, v. 24, n. 3, p. 285–311, 1989. Citado na página 26.

HEATON, J. B.; POLSON, N. G.; WITTE, J. H. Deep learning for finance: deep portfolios. *Applied Stochastic Models in Business and Industry*, Wiley Online Library, v. 33, n. 1, p. 3–12, 2017. Citado na página 31.

HENS, T.; WÖHRMANN, P. Strategic asset allocation and market timing: a reinforcement learning approach. *Computational Economics*, Springer, v. 29, n. 3, p. 369–381, 2007. Citado 2 vezes nas páginas 33 and 34.

HUELIN, L.; MIRZA, K. Portfolio optimisation na downside risk framework. *Germany: Lap Lambert Academic Publishing*, 2011. Citado 3 vezes nas páginas 25, 26, and 27.

JANG, J.; SEONG, N. Deep reinforcement learning for stock portfolio optimization by connecting with modern portfolio theory. *Expert Systems with Applications*, Elsevier, v. 218, p. 119556, 2023. Citado na página 32.

JIANG, Z.; LIANG, J. Deep reinforcement learning for portfolio management. *arXiv preprint arXiv:1706.10059*, 2017. Citado na página 31.

JR, J. B. L. Stock prices, inflation, and the term structure of interest rates. *Journal of Financial Economics*, v. 1, n. 2, p. 131–170, 1974. Citado na página 21.

JR, J. E. I. *Theory of Financial Decision Making.* [S.l.]: Rowman Littlefield, 1987. Citado na página 22.

KALAYCI, C. B.; ERTENLICE, O.; AKBAY, M. A. A comprehensive review of deterministic models and applications for mean-variance portfolio optimization. *Expert Systems with Applications*, Elsevier, v. 125, p. 345–368, 2019. Citado na página 16.

KHAIDEM, L.; SAHA, S.; DEY, S. R. Predicting the direction of stock market prices using random forest. *arXiv preprint arXiv:1605.00003*, 2016. Citado 3 vezes nas páginas 35, 36, and 38.

KOUMOU, G. B. Diversification and portfolio theory: a review. *Financial Markets and Portfolio Management*, Springer, v. 34, n. 3, p. 267–312, 2020. Citado na página 16.

KRAUSS, C.; DO, X. A.; HUCK, N. Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the s&p 500. *European Journal of Operational Research*, Elsevier, v. 259, n. 2, p. 689–702, 2017. Citado na página 28.

LEDOIT, O.; WOLF, M. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of multivariate analysis*, Elsevier, v. 88, n. 2, p. 365–411, 2004. Citado 2 vezes nas páginas 28 and 29.

LEKOVIĆ, M. M. Historical development of portfolio theory. *Tehnika*, v. 76, n. 2, p. 220–227, 2021. Citado na página 25.

LHABITANT, F.-S. *Portfolio diversification*. [S.l.]: Elsevier, 2017. Citado na página 16.

LIANG, Z. et al. Adversarial deep reinforcement learning in portfolio management. *arXiv preprint arXiv:1808.09940*, 2018. Citado na página 31.

LINTNER, J. Security prices, risk, and maximal gains from diversification. *The journal of finance*, JSTOR, v. 20, n. 4, p. 587–615, 1965. Citado na página 20.

MA, Y.; HAN, R.; WANG, W. Portfolio optimization with return prediction using deep learning and machine learning. *Expert Systems with Applications*, Elsevier, v. 165, p. 113973, 2021. Citado na página 32.

MARHFOR, A. Portfolio performance measurement: Review of literature and avenues of future research. *American Journal of Industrial and Business Management*, Scientific Research Publishing, v. 6, n. 4, p. 432–438, 2016. Citado na página 16.

MARINGER, D.; RAMTOHUL, T. Threshold recurrent reinforcement learning model for automated trading. In: SPRINGER. *European conference on the applications of evolutionary computation*. [S.l.], 2010. p. 212–221. Citado 2 vezes nas páginas 33 and 34.

MARKOWITZ, H. Portfolio selection. *The Journal of Finance*, [American Finance Association, Wiley], v. 7, n. 1, p. 77–91, 1952. ISSN 00221082, 15406261. Disponível em: <http://www.jstor.org/stable/2975974>. Citado 2 vezes nas páginas 19 and 20.

MARTIN, R. *PyPortfolioOpt: portfolio optimisation in Python*. [S.l.]: GitHub, 2021. <https://github.com/robertmartin8/PyPortfolioOpt>. Citado na página 35.

MEIER, L.; GEER, S. van de; BÜHLMANN, P. The group lasso for logistic regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, Wiley Online Library, v. 70, n. 1, p. 53–71, 2008. Citado 2 vezes nas páginas 30 and 31.

MERTON, R. C. Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, v. 51, n. 3, p. 247–257, 1969. Citado na página 22.

MERTON, R. C. An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society*, v. 41, n. 5, p. 867–887, 1973. Citado na página 22.

MILLINGTON, T.; NIRANJAN, M. Robust portfolio risk minimization using the graphical lasso. In: SPRINGER. *International Conference on Neural Information Processing*. [S.l.], 2017. p. 863–872. Citado 2 vezes nas páginas 30 and 31.

MOODY, J. et al. Performance functions and reinforcement learning for trading systems and portfolios. *Journal of Forecasting*, Wiley Online Library, v. 17, n. 5-6, p. 441–470, 1998. Citado 2 vezes nas páginas 33 and 34.

NEUNEIER, R. Optimal asset allocation using adaptive dynamic programming. *Advances in neural information processing systems*, v. 8, 1995. Citado 2 vezes nas páginas 32 and 34.

ORMONEIT, D.; GLYNN, P. W. Kernel-based reinforcement learning in average-cost problems: An application to optimal portfolio choice. *Advances in neural information processing systems*, v. 13, 2000. Citado 2 vezes nas páginas 33 and 34.

POGUE, G. A. An extension of the markowitz portfolio selection model to include variable transactions' costs, short sales, leverage policies and taxes. *The Journal of Finance*, JSTOR, v. 25, n. 5, p. 1005–1027, 1970. Citado na página 22.

RASEKHSCHAFFE, K. C.; JONES, R. C. Machine learning for stock selection. *Financial Analysts Journal*, Taylor & Francis, v. 75, n. 3, p. 70–88, 2019. Citado na página 28.

ROCKAFELLAR, R. T.; URYASEV, S. et al. Optimization of conditional value-at-risk. *Journal of risk*, Citeseer, v. 2, p. 21–42, 2000. Citado na página 26.

ROM, B. M.; FERGUSON, K. W. Post-modern portfolio theory comes of age. *Journal of investing*, v. 3, n. 3, p. 11–17, 1994. Citado na página 26.

ROSS, S. A. The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, Elsevier, v. 13, n. 3, p. 341–360, 1976. Citado na página 29.

ROSS, S. A. The capital asset pricing model (capm), short-sale restrictions and related issues. *The Journal of Finance*, JSTOR, v. 32, n. 1, p. 177–183, 1977. Citado na página 29.

RUBINSTEIN, M. E. The valuation of uncertain income streams and the pricing of options. *The Bell Journal of Economics*, v. 7, n. 2, p. 407–425, 1976. Citado na página 21.

SAMUELSON, P. A. Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics*, v. 51, n. 3, p. 239–246, 1969. Citado na página 21.

SHARPE, W. F. A simplified model for portfolio analysis. *Management science*, INFORMS, v. 9, n. 2, p. 277–293, 1963. Citado na página 29.

SHARPE, W. F. Capital asset prices: A theory of market equilibrium under conditions of risk. *The journal of finance*, Wiley Online Library, v. 19, n. 3, p. 425–442, 1964. Citado na página 20.

SHREVE, S. E.; SONER, H. M. Optimal investment and consumption with transaction costs. *The Annals of Applied Probability*, JSTOR, p. 609–692, 1994. Citado na página 23.

SORTINO, F. A.; PRICE, L. N. Performance measurement in a downside risk framework. *the Journal of Investing*, Portfolio Management Research, v. 3, n. 3, p. 59–64, 1994. Citado na página 25.

SUTIENE, K. et al. Enhancing portfolio management using artificial intelligence: literature review. *Frontiers in Artificial Intelligence*, Frontiers Media SA, v. 7, p. 1371502, 2024. Citado na página 16.

TOBIN, J. Liquidity preference as behavior towards risk. *The review of economic studies*, Wiley-Blackwell, v. 25, n. 2, p. 65–86, 1958. Citado na página 19.

TODONI, M.-D. A post-modern portfolio management approach on cee markets. *Procedia Economics and Finance*, Elsevier, v. 32, p. 1362–1376, 2015. Citado na página 18.

WU, Z. et al. Sparse and stable portfolio selection with parameter uncertainty. *Journal of the Operational Research Society*, Taylor & Francis, p. 1–21, 2024. Citado 2 vezes nas páginas 30 and 31.

Annex

ANNEX A – List of Assets

Table 5 – S&P 500 Tickers and Company Names

Ticker	Company Name
MOS	The Mosaic Company
BAC	Bank of America Corp
AKAM	Akamai Technologies, Inc.
PGR	Progressive Corp.
DECK	Deckers Outdoor Corp.
CPAY	Corpay, Inc.
COIN	Coinbase Global, Inc.
BMY	Bristol-Myers Squibb Co.
PFG	Principal Financial Group, Inc.
AVB	AvalonBay Communities, Inc.
NVDA	NVIDIA Corp.
USB	U.S. Bancorp
KR	Kroger Co.
ANET	Arista Networks, Inc.
MAS	Masco Corp.
GPC	Genuine Parts Co.
ALLE	Allegion PLC
ALGN	Align Technology, Inc.
T	AT&T Inc.
CMS	CMS Energy Corp.
ED	Consolidated Edison, Inc.
INVH	Invitation Homes Inc.
MDT	Medtronic PLC
ALB	Albemarle Corp.
LII	Lennox International Inc.
CB	Chubb Limited
PNR	Pentair PLC
NFLX	Netflix, Inc.
PLTR	Palantir Technologies Inc.
GIS	General Mills, Inc.
KO	The Coca-Cola Co.
HPE	Hewlett Packard Enterprise Co.
MLM	Martin Marietta Materials, Inc.
DELL	Dell Technologies Inc.
SWKS	Skyworks Solutions, Inc.
TJX	The TJX Companies, Inc.
ABBV	AbbVie Inc.
DGX	Quest Diagnostics Inc.
SHW	The Sherwin-Williams Co.
CPT	Camden Property Trust
PCAR	PACCAR Inc
EQIX	Equinix, Inc.
CHRW	C.H. Robinson Worldwide, Inc.
CME	CME Group Inc.
WDC	Western Digital Corp.
RJF	Raymond James Financial, Inc.
EQT	EQT Corp.
FIS	Fidelity National Information Services, Inc.
ADSK	Autodesk, Inc.
EXE	Exelon Corp.