Simple vs. Sophisticated Portfolio Optimization:

U.S. Retail Level Portfolio Study Case

By

Adiletkhan Nassylkhan

Submitted to

Central European University

Department of Economics and Business

In partial fulfillment of the requirements for the degree of

Master of Arts in Economic Policy in Global Markets

Supervisor: Professor Tomy Lee

Vienna, Austria

2025

Copyright Notice

For bibliographic and reference purposes this thesis should be referred to as: Nassylkhan, Adiletkhan. 2025. Simple vs. Sophisticated Portfolio Optimization: U.S. Retail Level Portfolio Study Case. MA thesis, Department of Economics and Business, Central European University, Vienna.

[1] Icon by Font Awesome.

Author's Declaration

I, Adiletkhan Nassylkhan, the undersigned candidate for the MA degree in Economic Policy in Global Markets, declare herewith that the present thesis titled "Simple vs. Sophisticated Portfolio Optimization: U.S. Retail Level Portfolio Study Case" is exclusively my own work, based on my research and only such external information as properly credited in notes and bibliography.

I declare that no unidentified and illegitimate use was made of the work of others, and no part of the thesis infringes on any person's or institution's copyright.

I also declare that no part of the thesis has been submitted in this form to any other institution of higher education for an academic degree.

Vienna, on 9th June 2025

Full Name: Adiletkhan Nassylkhan

Abstract

Over the past decade, exchange-traded funds (ETFs) have revolutionized retail investing by providing low-cost access to diversified portfolios that typically track an index or sector rather than a single stock. This shift raises the question of how portfolio optimization methods perform when applied solely to ETF-based, multi-asset class portfolios. This research evaluates the out-of-sample performance of five advanced portfolio optimization methods relative to the naïve equal-weighted (1/N) rule from the perspective of a retail investor. The analysis is based on a portfolio comprising 30 U.S.-listed ETFs that represent both traditional and alternative asset classes. A 252-day rolling-window simulation is used to generate daily out-of-sample performance data from March 30, 2009, to May 7, 2025, with daily rebalancing[2]. I test the optimization methods using two approaches. First, I apply optimization methods at the whole portfolio level without any constraints on asset class exposure. Second, I impose asset-class-specific weight constraints and perform optimization within each asset class. The results show that in the first case, complex optimizers often produce portfolios heavily concentrated in short-term Treasuries, allowing the naïve strategy to outperform substantially. In the second case, performance improves as optimizers operate within asset classes that share similar characteristics; however, this improvement still does not substantially outperform the naïve approach. Overall, the equal-weighted strategy consistently matches or outperforms the complex methods in terms of Sharpe ratio, turnover, and computation time. These findings suggest that for retail investors, the theoretical benefits of complex optimization are often outweighed by estimation error, high turnover, and computational cost.

Keywords: Portfolio Optimization, Retail Investors, Multi-Asset Portfolios, Rolling-Window Analysis

[2] As the main analysis relies on daily rebalancing, quarterly rebalancing results - provided in the Appendix 8.3 - show similar performance patterns.

Acknowledgments

Effort, action, and prosperity – three words that define the path to success. But staying focused on that path is never easy without the right people behind you. So, I want to thank everyone who supported me throughout this journey, and especially:

To my parents and family – thank you for always being there, cheering me on, and giving me everything I needed to keep going.

To Central European University – for providing a truly inspiring academic environment and a generous scholarship that made this experience possible.

To my supervisor, Prof. Tomy Lee – your thoughtful feedback and encouragement have helped shape the direction and quality of this thesis more than I can express.

To Prof. Marc Kaufmann, whose valuable advice and guidance greatly contributed to the development of this work.

And finally – and most importantly – thank you to my fiancée. Your unwavering support, belief in me, and gentle encouragement helped me reach the finish line of my master's program.

Table of Contents

List of Figures[3]

List of Tables [4]

[3] All Figures were produced by the author using Python with data from Yahoo Finance.
[4] All Tables were produced by the author using Python and Overleaf with data from Yahoo Finance.

List of Abbreviations

Abbreviation	Full Term
EW	Equal Weight
IV	Inverse Volatility
HRP	Hierarchical Risk Parity
MD	Maximum Diversification
GMV	Global Minimum Variance
MSR	Maximum Sharpe Ratio
CVaR	Conditional Value at Risk
ETF	Exchange-Traded Fund
Sharpe Ratio	Risk-adjusted return measure
Active Assets	Number of non-zero weighted assets in the portfolio
Asset Class	Group of securities with similar characteristics (e.g., equities, bonds)
Constrained Setting	Portfolio optimization with fixed weights assigned to asset classes
Unconstrained Setting	Portfolio optimization without limits on asset class exposure
Optimization	The process of selecting portfolio weights to maximize or minimize specific objectives

1. Introduction

Portfolio optimization originates from the fundamental principle of diversification, often captured by the phrase "not putting all your eggs in one basket." It is a quantitative approach that helps investors identify the most efficient combination of assets within a portfolio. The objective is typically to achieve the highest expected return for a given level of risk or to minimize risk for a given level of return. The field has been extensively studied since 1952 when Harry Markowitz introduced the Modern Portfolio Theory (Markowitz 1952). Since then, the growth of data availability and computational capabilities has led to the development of more advanced optimization techniques.

At the same time, the investment landscape has evolved in response to the rapid growth of exchange-traded funds (ETFs). ETFs are investment vehicles listed on stock exchanges that seek to replicate the performance of a particular index, sector, or asset class. Since 2008, ETFs have experienced considerable growth, with global assets under management (AUM) reaching 11.1 trillion USD by the end of 2023. This corresponds to a compound annual growth rate of 19.8% (Morningstar 2023, cited in State Street Global Advisors 2024). According to a recent survey, nearly one-third of respondents expect global ETF AUM to more than double, reaching 30 trillion USD within the next five years, while 60% anticipate that ETF assets will reach at least 26 trillion USD by June 2029 (PwC 2024).

While a substantial body of academic research has focused on portfolio optimization using individual stocks, such as S&P 500 constituents or bond funds, relatively few studies have analyzed optimization within ETF-only portfolios. As retail investors with limited resources and practical constraints increasingly favor ETFs over individual securities to reduce company-specific risk, it becomes especially relevant to examine how portfolio optimization performs in portfolios composed entirely of ETFs. This raises an important question: Do complex optimization methods significantly outperform the naïve equal-weighted (1/N) rule when applied to retail-level, multi-asset portfolios made up entirely of ETFs?

This study evaluates the out-of-sample performance of five advanced portfolio optimization methods in comparison to the naïve (1/N) strategy, using a dataset of 30 U.S.-listed ETFs representing both traditional and alternative asset classes. A rolling-window approach with a 252-day horizon is applied to simulate daily out-of-sample portfolio performance from March 30, 2009, to May 7, 2025, with daily rebalancing[5]. The analysis is conducted under two frameworks: an unconstrained setting, where optimizers freely allocate across all assets, and a constrained setting, where weights are fixed at the asset class level, and optimization is applied within each class.

The results reveal that in the unconstrained case, complex optimization techniques tend to produce impractical portfolios, often heavily concentrated in ETFs that track short-term Treasuries, leading to underperformance relative to the naïve strategy. Although performance improves under the constrained framework, the naïve (1/N) strategy portfolio still performs comparably or better in terms of Sharpe ratio, turnover, and computational efficiency. These results suggest that for retail investors, the naïve (1/N) strategy can offer more robust and cost-effective outcomes than complex optimization methods, particularly when accounting for estimation error and implementation challenges. These findings provide practical guidance for retail investors, financial advisors, and policymakers, underscoring the importance of simplicity.

The remainder of this thesis is structured as follows. Chapter 2 reviews the relevant literature. Chapter 3 describes the dataset and the ETF selection process. Chapter 4 outlines the optimization setup and performance evaluation metrics. Chapter 5 presents the empirical

[5]As the main analysis relies on daily rebalancing, quarterly rebalancing results - provided in the Appendix 8.3 - show similar performance patterns.

results and their interpretation. Chapter 6 proposes policy recommendations. Chapter 7 concludes with a summary of the main insights and suggestions for future research. Chapter 8 provides an appendix with supplementary results, including full dataset details and rolling Sharpe ratios. It also presents an extended analysis of portfolios composed of 60 ETFs. To assess the robustness of the findings, results for a quarterly rebalanced portfolio of 30 ETFs are also included.

2. Literature Review

The literature on portfolio optimization is well-established, with most research focusing on methods for selecting the most efficient mix of securities within a single asset class, such as stocks, bonds, or commodities. However, the growing use of exchange-traded funds (ETFs) by retail investors – attracted by their low costs, liquidity, and built-in diversification – raises questions about the relevance of traditional optimization approaches when applied to portfolios composed entirely of ETFs. This thesis addresses this gap by evaluating whether complex optimization strategies provide added value over the naïve equal-weighted approach in ETF-based portfolios that are already diversified across sectors or asset classes.

DeMiguel, Garlappi, and Uppal (2009) provide a foundational study for evaluating portfolio optimization methods. They compare thirteen optimization strategies against the naïve equal-weighted (1/N) rule using seven monthly datasets of U.S. equity market returns, including sector, industry, international, and factor portfolios. Their findings reveal that none of the optimized strategies consistently outperform the 1/N rule out-of-sample in terms of Sharpe ratio, certainty-equivalent return, or turnover. This underperformance is primarily attributed to estimation error in return and risk parameter inputs derived from historical data. While the study is critical in challenging the practical utility of theoretically optimal models, it remains limited to equity portfolios and does not address ETF-specific characteristics.

Plyakha, Uppal, and Vilkov (2012) further explore portfolio construction by comparing equal-, value-, and price-weighted strategies using data from 1000 randomly generated portfolios composed of 100 stocks from the S&P 500 index. Their findings show that equal-weighted portfolios, rebalanced monthly, outperform the alternatives on key performance metrics, including total mean return, four-factor alpha, Sharpe ratio, and certainty-equivalent return. This outperformance is attributed mainly to a rebalancing premium and higher systematic risk. Despite these contributions, the study also remains confined to equity portfolios and does not examine how such strategies perform in multi-asset, ETF-based settings.

Jacobs, Müller, and Weber (2014) take a broader view by evaluating diversification strategies not only across equities but also across asset classes. Their study compares eleven optimization models with a set of heuristic allocation rules using a dataset that spans nearly four decades and includes global equity indices (MSCI regional), euro-denominated bonds (iBoxx Euro Overall Index), and commodities (S&P GSCI). Although the study does not use ETF return series directly, it uses indices that are transparent, investable, and commonly tracked by ETFs, thereby making the results applicable to ETF-based portfolio strategies. The authors find that heuristic strategies such as equal weighting and GDP-weighting offer diversification benefits comparable to those of advanced portfolio optimization methods.

Together, these studies make a significant contribution to the understanding of portfolio optimization and highlight the challenges associated with implementing complex models in practice. However, they do not fully capture the realities of ETF-based investing. Unlike individual securities, ETFs already encapsulate a degree of diversification by tracking entire sectors, indices, or asset classes. This raises doubts about the incremental benefits of applying optimization techniques to portfolios exclusively composed of such instruments.

3. Data

This study analyzes a diversified portfolio composed of 30 U.S.-listed ETFs. The selection is based on TradingView's list [6] of the 100 most-traded ETFs, ranked by daily trading volume multiplied by share price. While narrowing the list to 30 required discretionary judgments, the selection prioritized highly liquid and widely recognized ETFs to minimize subjective bias and enhance the validity of the analysis. A full description of the selected ETFs is provided in Appendix 8.1.1: Dataset Details – 30 ETFs.

Daily adjusted closing price data were retrieved via the Yahoo Finance API and converted into return series, covering the period from March 31, 2008, to May 7, 2025 (4,305 observations per ETF). The start date of March 31, 2008 was chosen to ensure a consistent and complete time series across all selected ETFs. ETFs with shorter trading histories, such as XLRE (Real Estate) and XLC (Communication Services), were excluded to maintain consistency. For out-of-sample analysis, the period from March 30, 2009, to May 7, 2025 (4,054 observations) is used. The final portfolio includes the following ETFs, classified as shown in Table 1 below.

Table 1: Classification of 30 ETFs

Asset Class	**Number of ETFs**	**Examples**
U.S. Equity Sectors	9	Financials, Technology, Healthcare (e.g., XLF, XLK, XLV)
U.S. Equity Indices	4	Broad Market, Small/Mid Cap (e.g., SPY, QQQ, IWM)
International Equity	3	Global and Emerging Markets (e.g., ACWI, EEM)
Fixed Income	8	Corporate, Treasury, and Muni Bonds (e.g., TLT, LQD, BND)
Alternative Investments	3	Gold, Silver, Real Estate (e.g., GLD, SLV, VNQ)
Cash Equivalents	3	Short-Term Treasuries (e.g., BIL, SHY)
Total	**30 ETFs**	

Figure 1 displays the annualized risk-return profiles of the 30 selected ETFs over the full sample period from March 31, 2008, to May 7, 2025. Each point represents an ETF,

[6] TradingView, Most Traded ETFs, accessed June 7, 2025, https://www.tradingview.com/markets/etfs/funds-most-traded/.

positioned according to its annualized average return (y-axis) and annualized standard deviation (x-axis). The ETFs are color-coded by asset class, clearly revealing clustering patterns that reflect the distinct characteristics of each asset class. As expected, ETFs such as SHY and BIL (black), which track short-term Treasuries, exhibit the lowest volatility and returns. In contrast, equity ETFs (blue) show higher volatility accompanied by higher returns. Fixed-income ETFs (yellow) fall into a moderate risk-return range, while international equities (red) and alternatives (green) exhibit greater dispersion. This highlights the variation in risk-return profiles across asset classes and supports the case for diversified, class-aware portfolio construction.

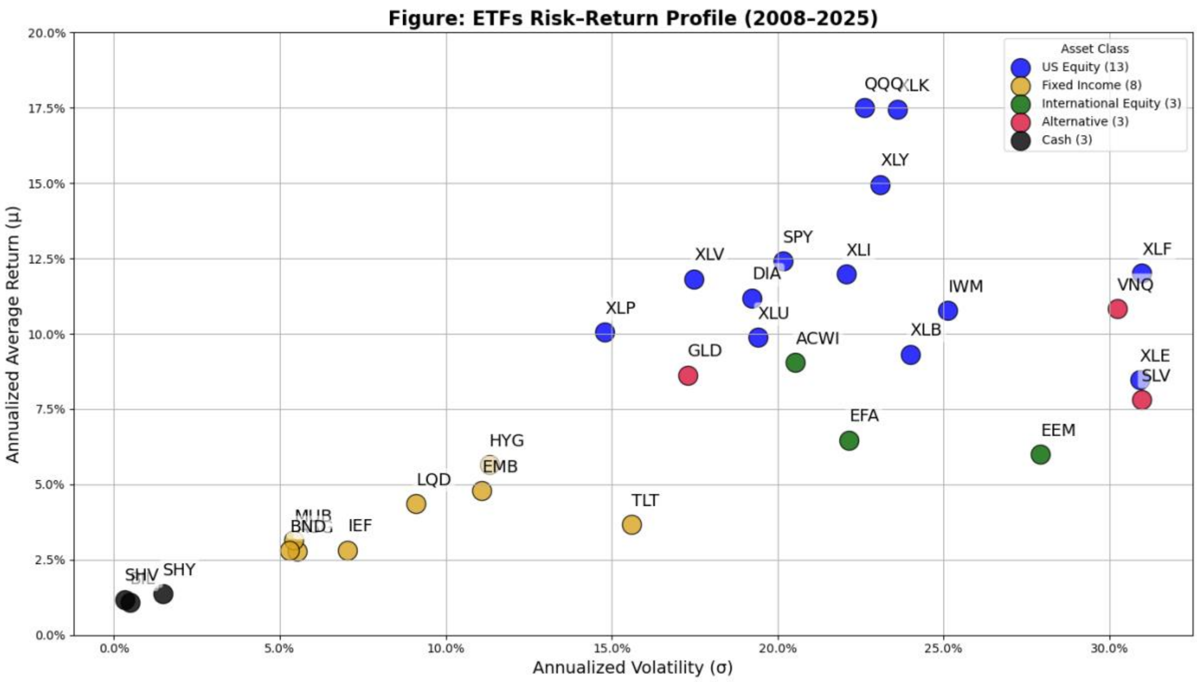

Figure 1: ETFs Risk-Return Profile Full Sample Period

4. Methodology

Let $P_{i,t}$ denote the closing price of asset i at time t. The corresponding simple return in percentage terms is given by: $r_{i,t} = 100 \times \frac{(P_{i,t} - P_{i,t-1})}{P_{i,t-1}}$.

Let $r_t = (r_{1,t}, \dots, r_{N,t})$ represent the vector of asset returns at time t, and let $\omega = (\omega_1, \dots, \omega_N)$ be the vector of portfolio weights. The return of the portfolio at time t, denoted by R_t, is then computed as the weighted sum of individual asset returns:

$$R_t = \omega' r_t = \omega_1 r_{1,t} + \dots + \omega_N r_{N,t}$$

Since portfolio return is the weighted sum of asset returns, the main goal of portfolio optimizers is to determine the set of weights that balances risk and return. To achieve this, optimizers use historical asset return data as input and compute weight allocations based on specific optimization criteria.

4.1 Optimization Setup

To evaluate the performance and robustness of the optimization strategies, two approaches were applied: unconstrained and constrained. In the unconstrained setting, optimization was applied at the whole portfolio level without any restrictions on asset class exposures. This allowed the optimizers to allocate freely across all 30 ETFs based solely on historical risk and return characteristics. Such a setup is valuable for evaluating how each optimizer performs under maximum flexibility without any predefined guidance or human-imposed structure on asset class weights.

Under the constrained approach, the process began by assigning fixed weights to each asset class, as shown in Table 2. Optimization was then applied individually within each asset class. This structure leverages the fact that assets within the same class tend to exhibit similar behavior, thereby making within-class optimization more reliable. Additionally, when optimization is applied at the overall portfolio level, it must select among all 30 ETFs. In contrast, optimizing within asset classes reduces the dimensionality of each problem – for example, choosing among only 16 ETFs in the equity class or 8 in fixed income – thereby lowering estimation error and improving stability.

Table 2: Portfolio Weights by Asset Class

Asset Class	Weight Allocation
Equities	55%
Fixed Income	35%
Alternatives	5%
Short-Term Treasuries	5%

4.2 Walk-Forward Optimization

Given that the optimizer's input is historical data, their outputs are sensitive to the specific sample used. To address this and reduce the risk of overfitting, a rolling window approach – a dynamic form of walk-forward optimization – is used to enhance the robustness and generalizability of the results. This method follows the structure outlined by Jacobs, Müller, and Weber (2014). Still, it is adjusted based on daily frequency and daily rebalancing to generate rolling estimations for evaluating the out-of-sample performance of portfolio strategies. The process is structured as follows:

- Step 1: Training window – The model is trained on a rolling window of the most recent 252 trading days (app. one year) to compute optimal portfolio weights.
- Step 2: Testing window – These weights are then applied to the following trading day, which serves as out-of-sample data to evaluate performance.
- Step 3: Iteration – The window is rolled forward by one day, and the process is repeated. Thus, the portfolio is rebalanced daily. For quarterly rebalancing results, please refer to the Appendix 8.3.

Overall, this approach generates a sequence of daily out-of-sample returns, allowing for a more realistic and reliable evaluation of each strategy's performance over time.

4.3 Portfolio Allocation Strategies

This section provides a brief overview of six optimization methods used in this research to find the optimal portfolio weights, $\mathrm{w} = (\omega_1, \ldots, \omega_N)$.

4.3.1 Equally Weighted (EW)

Also known as the naïve or 1/N strategy, the equally weighted portfolio is one of the simplest forms of portfolio optimization. It assigns an equal share of capital to each asset in the portfolio without relying on expected return or risk estimates. Despite its simplicity, it often serves as a surprisingly difficult benchmark to outperform. For a portfolio with N assets, the weights are:

$$\omega_i = \frac{1}{N}, for\ i = 1, \ldots, N.$$

4.3.2 Inverse-Volatility (IV)

Originally proposed by Carvalho, Xiao, and Moulin (2011), the Inverse-Volatility (IV) strategy assigns asset weights based on the inverse of each asset's historical volatility (standard deviation). Assets with lower risk receive higher weights, thereby contributing to a reduction in overall portfolio risk. It is essential to note that this approach relies on historical volatility and assumes that past volatility will continue to persist in the future. The weights are:

$$\omega_i = \frac{1/\sigma_i}{\Sigma_{j=1}^{N}(1/\sigma_i)}$$

4.3.3 Maximum Diversification (MD)

Introduced by Choueifaty and Coignard (2008) and further developed by Choueifaty et al. (2013), the Maximum Diversification (MD) optimizer aims to construct the most diversified portfolio by maximizing the diversification ratio D. This ratio is defined as the weighted average of individual asset volatilities divided by the portfolio volatility. The strategy allocates higher weights to assets that contribute more to diversification. The optimization problem is formulated as follows:

$$max\ D = \frac{\omega^T \sigma}{\sqrt{\omega^T \Sigma \omega}} \quad subject\ to \quad \sum_{i\,=\,1}^{N} \omega_i = 1, \omega_i \geq 0$$

4.3.4 Maximum Sharpe Ratio (MSR)

The Maximum Sharpe Ratio (MSR) strategy builds on the mean-variance optimization framework introduced by Markowitz (1952) and was further developed by Sharpe (1966), who introduced the Sharpe Ratio. This strategy seeks to identify the portfolio on the efficient frontier that maximizes the Sharpe Ratio by achieving the best trade-off between expected return and volatility. The optimization problem is formulated as follows:

$$max\ \frac{\omega^T \mu\ -\ r_f}{\sqrt{\omega^T \Sigma \omega}}$$

Where μ is the vector of expected returns, is the risk-free rate, and Σ is the covariance matrix of asset returns. To enhance stability and mitigate estimation error, the covariance matrix in this study was estimated using the Ledoit-Wolf shrinkage method (Ledoit and Wolf 2004).

4.3.5 Global Minimum Variance (GMV)

The Global Minimum Variance (GMV) strategy represents a special case of mean-variance optimization, as proposed initially by Markowitz (1952). On the efficient frontier, it represents the portfolio with the lowest possible risk, regardless of the expected return. This approach completely disregards return forecasts and focuses solely on minimizing total portfolio volatility. The optimization problem is defined as:

$$min\ \omega^T \Sigma \omega \qquad subject\ to \quad \sum_{i=1}^{N} \omega_i = 1$$

As with the MSR strategy, the covariance matrix Σ was estimated using the Ledoit–Wolf shrinkage method (Ledoit and Wolf 2004).

4.3.6 Hierarchical Risk Parity (HRP)

Hierarchical Risk Parity (HRP), introduced by López de Prado (2016), constructs portfolios based on a hierarchical clustering algorithm that groups assets according to their similarities. Unlike traditional optimizers that rely on inverting the covariance matrix, HRP allocates weights based on the hierarchical structure of asset correlations, aiming to form risk-balanced clusters. This technique enhances stability and reduces estimation error, particularly in high-dimensional settings.

Although HRP lacks a closed-form optimization expression, its procedure typically involves four key steps: (1) estimating the correlation matrix, (2) constructing a hierarchical clustering dendrogram, (3) reordering the matrix using quasi-diagonalization, and (4) allocating weights recursively based on cluster variances.

4.4 Statistical Measures

To evaluate each portfolio optimizer, this study uses the built-in performance evaluation methods provided by Skfolio, a Python library developed by Delatte and Nicolini (2023) (see Appendix). Among the available metrics, the following key performance indicators are selected, as they are widely used in portfolio optimization research:

Cumulative Returns: Captures the total percentage gain or loss over the investment period, taking into account compounding. It is computed as the product of sequential daily returns minus one:

$$Cumulative\ Returns = (1 + r_{t+1}) \times (1 + r_{t+2}) \times \ldots (1 + r_{t+n}) - 1$$

Sharpe Ratio: Measures risk-adjusted performance by comparing excess return (portfolio return minus the risk-free rate) to volatility (Sharpe 1966). A higher Sharpe ratio indicates a more favorable return per unit of risk. This study sets the risk-free rate to 0%, a common simplification in daily return-based analysis to avoid adding noise from near-zero short-term rates.

$$Sharpe\ Ratio = \frac{R_p - r_f}{\sigma_p}$$

Average Daily Turnover: Measures the frequency of portfolio weight changes, calculated as the average absolute change in weights between consecutive daily rebalancing periods. Higher turnover implies increased trading costs.

Maximum Drawdown: This represents the largest decline from a portfolio's peak to its trough, highlighting the worst-case loss scenario and providing insight into the portfolio's downside risk.

CVaR at 95%: Conditional Value at Risk (CVaR) at 95% measures the expected average loss in the worst 5% of cases. Unlike Value at Risk (VaR), which indicates the minimum loss beyond a confidence threshold, CVaR quantifies the severity of losses that exceed that threshold. This makes CVaR a more comprehensive measure of tail risk and is particularly useful for assessing downside exposure in portfolio optimization, especially when return distributions exhibit fat tails or skewness (Rockafellar and Uryasev 2000).

5. Results and Discussion

This section presents the main results of the portfolio optimization analysis, which is based on a portfolio composed of 30 ETFs. It compares performance under both unconstrained and constrained settings, highlights limitations of the optimizers, and evaluates computational efficiency alongside key portfolio metrics such as turnover and the average number of assets. Additional analyses using a 60-ETF portfolio and quarterly rebalancing – designed to test model robustness on larger samples and different rebalancing frequencies – are provided in Appendix 8.2 (Portfolio of 60 ETFs) and Appendix 8.3 (Results: Quarterly Rebalancing), respectively.

5.1 Results: Unconstrained Optimization

Table 3: Unconstrained: Out-of-Sample Performance Metrics (30 ETFs)

Performance Metrics	EW	IV	MD	MSR	GMV	HRP
Ann Mean	9.74%	2.51%	1.46%	4.13%	1.67%	1.14%
Ann Std Dev	10.72%	1.63%	0.43%	7.22%	1.09%	0.22%
Ann Sharpe Ratio	0.91	1.54	3.38	0.57	1.54	5.15
MAX Drawdown	26.08%	3.65%	0.95%	26.64%	4.51%	0.31%
CVaR at 95%	1.59%	0.24%	0.057%	1.18%	0.14%	0.023%

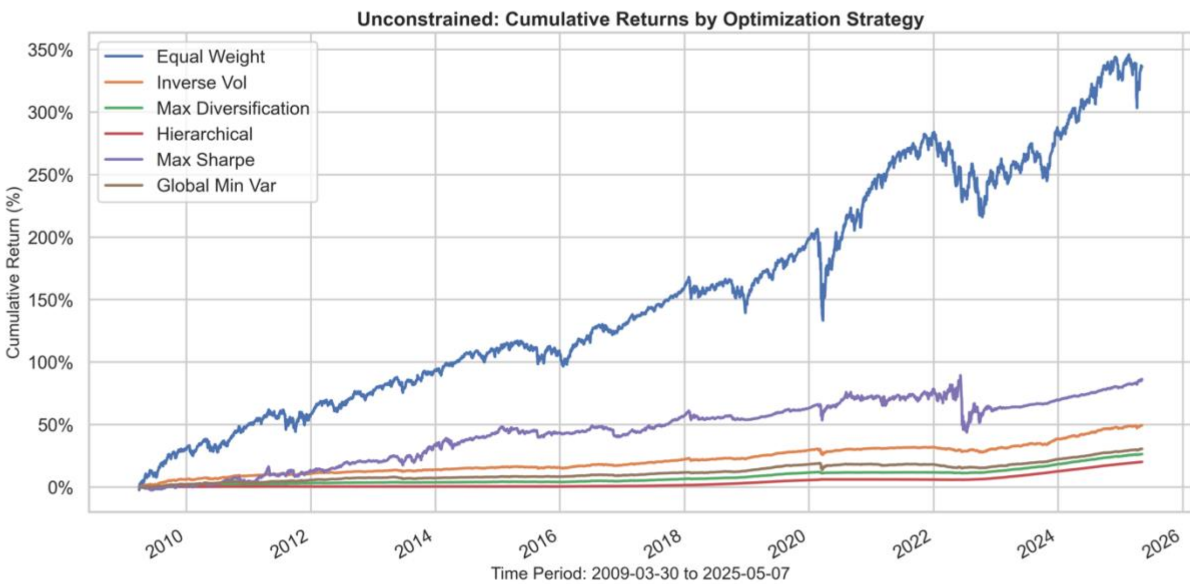

Figure 2: Unconstrained: Out-of-Sample Cumulative Portfolio Returns (30 ETFs)

Table 3 and Figure 2 summarize the out-of-sample performance metrics under the unconstrained setting, where optimizers were applied at the overall portfolio level without any limits on asset class exposure.

The Equal Weight (EW) strategy delivered the highest annualized mean return (9.74%) and achieved over 300% cumulative growth over the 16-year period, outperforming all other methods. However, this performance was accompanied by relatively high volatility (10.72%) and a maximum drawdown of 26.08%, resulting in a Sharpe ratio of 0.91. Despite its simplicity, EW achieved strong risk-adjusted returns. This aligns with evidence that many individual investors tend to favor such straightforward allocation approaches (Benartzi and Thaler 2007). In contrast, the Maximum Sharpe Ratio (MSR) strategy yielded weaker outcomes, with a lower Sharpe ratio (0.57) and a comparable drawdown of 26.64%, indicating limited robustness – likely due to instability in the covariance matrix and resulting estimation errors.

Hierarchical Risk Parity (HRP) and Maximum Diversification (MD) delivered exceptionally high Sharpe ratios – 5.15 and 3.38, respectively – primarily due to their concentrated allocations to short-term Treasury bills. This resulted in extremely low portfolio volatility (0.22% for HRP and 0.43% for MD) but also low annualized mean returns: 1.14% for HRP and 1.46% for MD. Consequently, both strategies achieved low cumulative returns over the full period despite their high Sharpe ratios.

Inverse Volatility (IV) and Global Minimum Variance (GMV) delivered more balanced and consistent performance profiles. IV achieved a moderate return of 2.51% with low volatility (1.63%), resulting in a Sharpe ratio of 1.54. GMV showed a comparable outcome, yielding a return of 1.67% with lower volatility (1.09%) and an identical Sharpe ratio of 1.54.

In summary, EW outperformed in cumulative returns and delivered results more consistent with investor preferences. At the same time, models like HRP and MD excelled on risk-adjusted metrics – primarily due to their heavy exposure to short-term Treasury bills.

These findings highlight the limitations of unconstrained optimizers in multi-asset portfolios and underscore the importance of incorporating practical constraints in model design.

5.2 Optimizer's Limitations

In the unconstrained setting, all optimizers – except Equal Weight (EW) and Maximum Sharpe Ratio (MSR) – produced allocations heavily concentrated in short-term Treasury ETFs. As shown in Figure 3, Hierarchical Risk Parity (HRP) consistently allocated a large portion of the portfolio to Ultra-Short Treasuries (SHV) and 1–3 Month Treasury Bills (BIL), reflecting its strong preference for low-volatility assets throughout the full out-of-sample period.

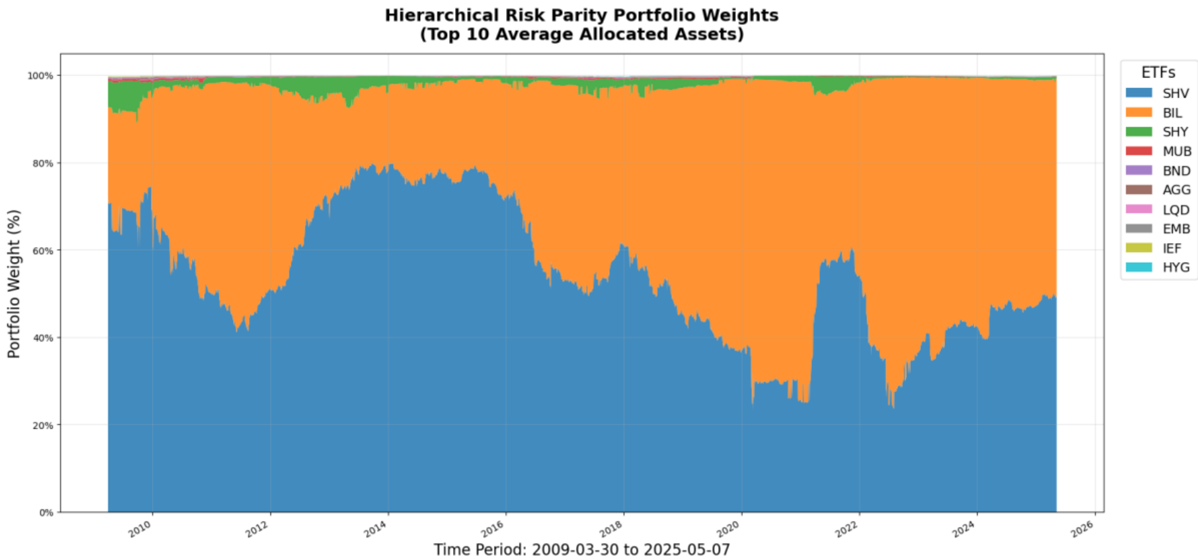

Figure 3: Unconstrained: HRP Portfolio Weights (Top 10 Average Allocated Assets)

While this allocation effectively reduced volatility, it also significantly constrained return potential, resulting in portfolios that were highly stable but lacked meaningful growth. This outcome reflects how certain optimizers rank assets based on return per unit of risk, as illustrated in Figure 4. Under this criterion, short-term Treasury bills with near-zero volatility appear disproportionately attractive, often leading optimizers to favor these instruments and construct overly concentrated portfolios.

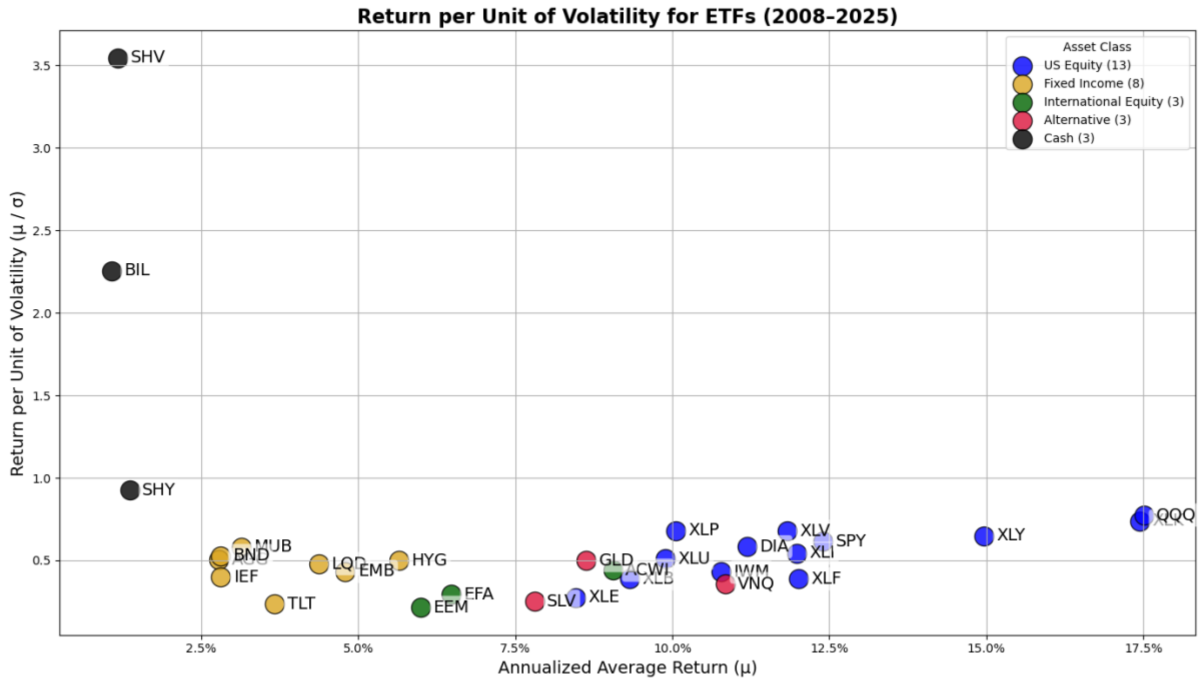

Figure 4: Return per Unit of Volatility for ETFs

The MSR portfolio exhibited unstable and unpredictable allocation patterns, as illustrated in Figure 5, where asset weights shifted abruptly over time. This instability arises from estimation errors in both expected returns and the covariance matrix, two critical components of the MSR optimization process. Because the method relies on inverting the covariance matrix and is highly sensitive to even slight inaccuracies in return estimates, minor errors can lead to substantial and erratic changes in portfolio composition (DeMiguel, Garlappi, and Uppal 2009). When applied to a universe of 30 ETFs within a rolling-window framework, the inversion process became particularly unreliable, resulting in highly volatile and unbalanced allocations. However, when constraints were introduced, and MSR was applied separately within individual asset classes, performance improved. The smaller number of assets within each class enhanced the stability of the optimization and produced more interpretable and consistent portfolio weights, as discussed in the following section.

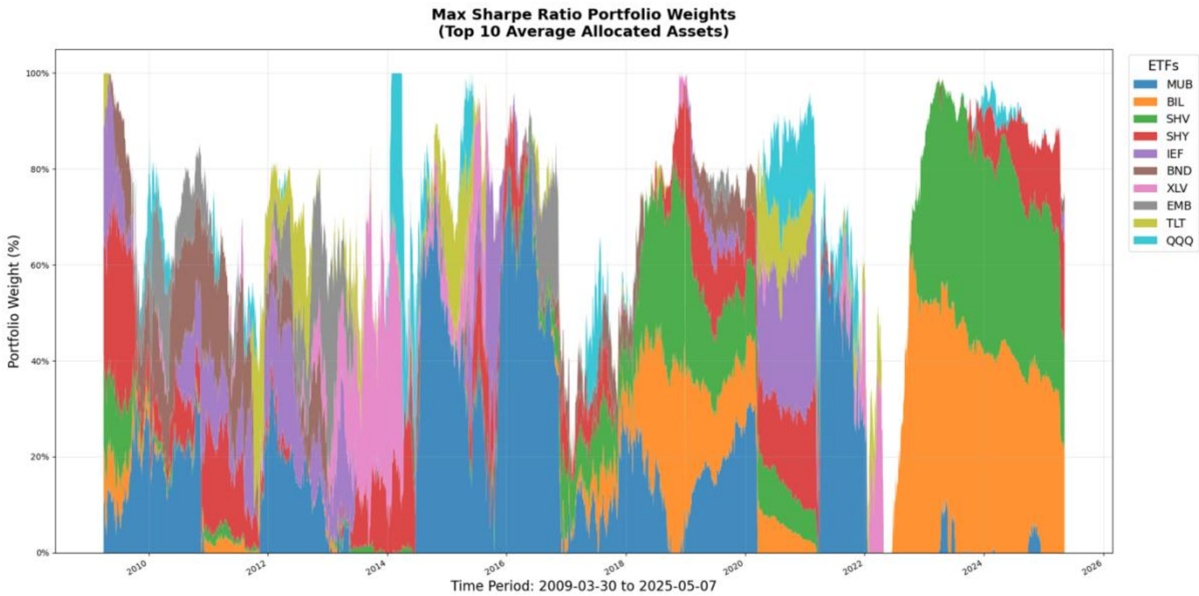

Figure 5: Unconstrained: MSR Portfolio Weights (Top 10 Average Allocated Assets)

5.3 Results: Constrained Optimization

Table 4: Constrained: Out-of-Sample Performance Metrics (30 ETFs)

Performance Metrics	EW	IV	MD	MSR	GMV	HRP
Ann Mean	9.69%	9.34%	8.57%	8.13%	7.79%	9.02%
Ann Std Dev	10.62%	10.13%	10.10%	12.19%	8.14%	9.59%
Ann Sharpe Ratio	0.91	0.92	0.85	0.67	0.96	0.94
MAX Drawdown	25.93%	25.44%	30.41%	26.46%	22.96%	24.60%
CVaR at 95%	1.58%	1.50%	1.48%	1.92%	1.18%	1.42%

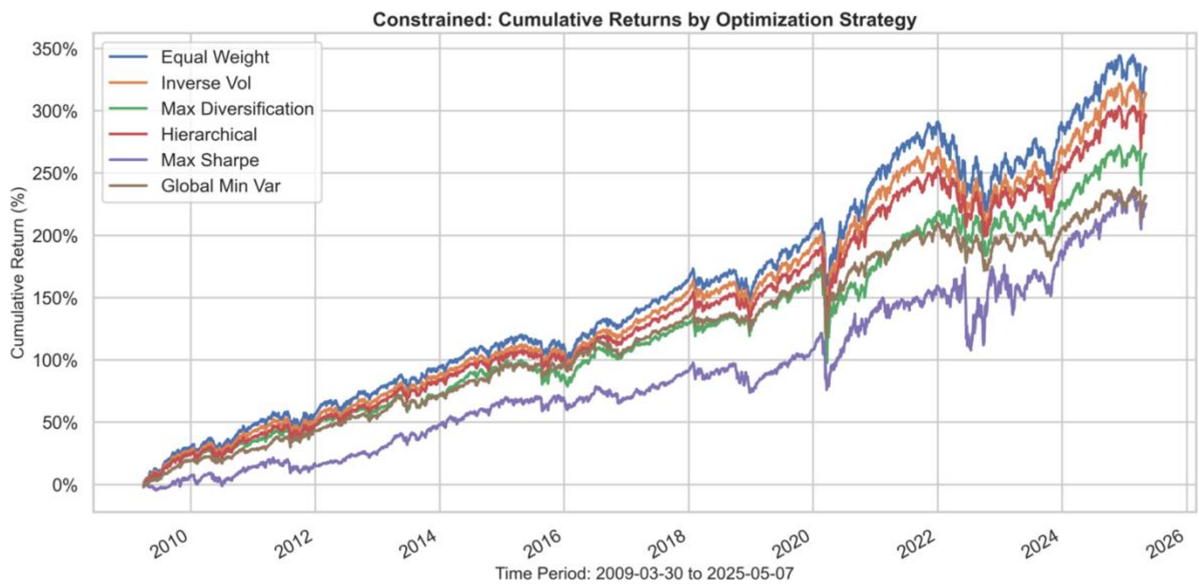

Figure 6: Constrained: Out-of-Sample Cumulative Portfolio Returns (30 ETFs)

Table 4 and Figure 6 summarize the out-of-sample performance metrics under the constrained setting. First, fixed weight allocations were assigned to each asset class, and then optimization was performed within each class individually.

Under this setup, all strategies produced more balanced and interpretable results. Equal Weight (EW) remained one of the top performers, achieving the highest annualized mean return (9.69%) and maintaining a decent Sharpe ratio (0.91). While it exhibited moderate volatility (10.62%) and a relatively high drawdown (25.93%), its consistent growth makes it a reliable benchmark for comparison.

Inverse Volatility (IV) and Hierarchical Risk Parity (HRP) also delivered strong results, with annual returns of 9.34% and 9.02%, respectively. Their Sharpe ratios (0.92 for IV and 0.94 for HRP) were slightly higher than that of EW, reflecting improved risk-adjusted performance. Importantly, their maximum drawdowns and CVaR levels were also marginally lower than those of EW, suggesting better downside protection within a diversified structure.

Maximum Diversification (MD) achieved a return of 8.57% but experienced the largest drawdown (30.41%) among all strategies, along with a relatively lower Sharpe ratio of 0.85. This indicates that while MD benefited from diversification, it remained vulnerable to market corrections even under the constrained setting.

Maximum Sharpe Ratio (MSR) showed noticeable improvement compared to the unconstrained case. Its Sharpe ratio increased to 0.67, and its return (8.13%) was decent, though it still lagged behind other strategies in terms of risk-adjusted efficiency.

Global Minimum Variance (GMV) continued to prioritize portfolio stability, yielding the lowest annualized volatility (8.14%) and the lowest CVaR (1.18%) among all strategies. Although it had the lowest return (7.79%), its Sharpe ratio (0.96) was the highest, indicating a highly efficient balance between risk and return in this setting.

In summary, applying fixed-weight constraints led to more realistic and robust optimization outcomes. Risk-focused strategies such as IV, HRP, and GMV became more competitive, while MSR showed improved stability. These results highlight the value of assigning asset-class weights and applying optimization within each class – particularly in multi-asset portfolio construction, where managing a broad set of asset classes can introduce significant estimation challenges. However, it is worth noting that none of the more complex optimizers significantly outperformed the simple EW strategy across the key performance metrics.

5.4 Cumulative Computation Time

Table 5: Computation Time in Seconds (30 ETFs)

Optimizers	Unconstrained	Constrained
EW	3	5
IV	18	36
MSR	76	350
GMV	76	395
MD	85	449
HRP	320	421

When evaluating portfolio optimizers for retail-level applications, computation time becomes a practical concern. Unlike institutional settings with access to high-performance computing, individual investors typically rely on personal devices with limited processing power. In such cases, optimizers that are too resource-intensive may not be suitable for frequent rebalancing or responsive portfolio adjustments. Therefore, understanding how long each method takes to run can help determine which strategies are both practical and accessible for everyday use.

Table 5 presents the computation times for each optimizer under both unconstrained and constrained conditions. These results were obtained using a MacBook Air 13" (2020) – a commonly used consumer laptop – making the findings directly relevant for retail investors. As expected, the Equal Weight (EW) strategy was the fastest, completing the task in just 3 seconds without constraints and 5 seconds with constraints due to its simple, non-iterative logic. Inverse Volatility (IV) was also efficient, though its computation time doubled from 18 to 36 seconds when constraints were applied – due to the added steps needed to enforce asset class limits.

More complex strategies, such as Maximum Sharpe Ratio (MSR), Global Minimum Variance (GMV), and Maximum Diversification (MD), required substantially longer runtimes. For example, MD took 85 seconds in the unconstrained case and 449 seconds in the constrained setup. These increases reflect the additional computational burden of working with large covariance matrices and solving iterative optimization problems. Hierarchical Risk Parity (HRP) had the highest overall computation time (320 seconds unconstrained, 421 seconds constrained), though it remained relatively stable across settings due to its clustering-based approach.

In summary, while all strategies remained computationally feasible on a standard laptop, the gap between simple and complex methods is noticeable. For retail investors managing their own portfolios, strategies such as EW and IV offer not only reasonable performance but also practical speed and ease of use. For those willing to wait longer or rebalance less frequently, more advanced methods, such as HRP or MD, remain viable.

5.5 Average Active Assets and Turnover

Table 6: Average Active Assets and Turnover (30 ETFs)

Optimizers	Unconstrained		Constrained	
	Avg. Active Assets	Avg. Daily Turnover	Avg. Active Assets	Avg. Daily Turnover
EW	30.00	0.56%	30.00	0.54%
IV	30.00	0.36%	30.00	0.59%
HRP	10.43	0.77%	30.00	1.33%
MD	11.92	1.43%	16.29	1.97%
GMV	9.63	0.55%	12.83	1.76%
MSR	8.45	9.64%	9.14	13.05%

When evaluating portfolio strategies, it is essential to look beyond returns and risk-adjusted metrics to consider indicators that reflect real-world usability – particularly for retail investors. Two such indicators are the average number of active assets and average daily turnover, both of which provide insight into a strategy's diversification and cost efficiency. A higher number of active assets suggests broader diversification, while a lower count indicates concentration. Turnover was calculated as the average sum of absolute changes in portfolio weights between consecutive days, reflecting how frequently and substantially the portfolio is rebalanced. Higher turnover implies more trading activity, which may result in increased transaction costs – a crucial consideration for retail investors managing their own portfolios.

Table 6 compares these two indicators across all portfolio strategies under both unconstrained and constrained settings. As expected, Equal Weight (EW) and Inverse Volatility (IV) maintain complete diversification across all 30 ETFs in both setups, with minimal daily turnover – demonstrating high stability and low trading costs. These characteristics make them particularly well-suited for retail investors seeking simplicity and cost control.

In the unconstrained case, more complex optimizers, such as the Maximum Sharpe Ratio (MSR), Global Minimum Variance (GMV), and Hierarchical Risk Parity (HRP), allocate to a narrower subset of assets, averaging only 8 to 10 ETFs. This reflects a strong tendency towards concentration, likely due to their sensitivity to estimated risk-return trade-offs. These optimizers also show relatively low turnover in the unconstrained setting, as they repeatedly favor a small set of preferred assets.

Under constrained optimization, however, all strategies shift toward broader diversification. HRP reaches full asset inclusion, and both GMV and MSR increase their average number of holdings. Yet, this diversification comes at a cost: higher daily turnover. Notably, MSR's turnover climbs to over 13%, implying more frequent trading and potentially higher transaction costs.

In summary, while constraints enhance diversification and reduce concentration risk, they also tend to increase portfolio activity. For retail investors, this underscores a critical trade-off between stability and realism in portfolio design. Simpler strategies such as EW and IV may offer a more accessible balance of performance, diversification, and operational ease.

6. Policy Implications

The findings of this thesis highlight the value of promoting simple, transparent, and accessible investment strategies for retail investors. The equal-weight (EW) approach consistently delivered strong performance while remaining low in complexity and transaction costs. As such, financial literacy programs and digital investment platforms should prioritize teaching and enabling this method. EW should also be considered a default benchmark against which more complex optimization strategies are evaluated. Educational initiatives – especially those targeting first-time investors – should include practical modules that cover basic ETF investing, risk diversification, and rebalancing techniques using intuitive strategies, such as EW. Moreover, robo-advisors and investment platforms should offer EW and other low-turnover strategies as default or entry-level options to support informed and cost-effective decision-making.

7. Conclusion

With exchange-traded funds (ETFs) now forming a major part of retail-level portfolios, it is essential to assess whether complex portfolio optimization methods offer meaningful advantages over the naïve equal-weighted (1/N) rule. While much of the existing research focuses on optimizing portfolios of individual securities, this study evaluates how complex optimization methods perform when applied to portfolios composed entirely of ETFs, which themselves offer built-in diversification by tracking broad indices, sectors, or asset classes.

Using a dataset of 30 U.S.-listed ETFs and a rolling-window approach with a 252-day window, I evaluated out-of-sample performance from March 30, 2009, to May 7, 2025, comparing five optimization strategies under both unconstrained and asset class-constrained frameworks. The findings show that in unconstrained optimization – where allocation was applied at the whole portfolio level without any restrictions on asset class exposures – many advanced strategies, despite their theoretical appeal, produced highly concentrated portfolios in short-term Treasuries. This resulted in limited long-term cumulative returns and underperformance compared to the naïve strategy.

Although constrained optimization – where fixed-weight exposure to asset classes was imposed and optimizers were used within each asset class – improved portfolio balance and stability, the equal-weight approach still performed comparably or better across key metrics, such as the Sharpe ratio, turnover, and computational efficiency. These results suggest that for retail investors, the 1/N strategy remains a robust, cost-effective, and easy-to-implement solution.

This study provides several key insights for retail investors, financial advisors, and policymakers. First, the equal-weight strategy remains a difficult benchmark to surpass. Second, applying asset class constraints enhances optimization by reducing overfitting and improving portfolio structure. Third, optimizers should be used within the same asset classes; otherwise, differences in asset characteristics can distort allocations and reduce effectiveness.

Future research could implement other optimization strategies, expand the ETF universe to include international or thematic exposures and integrate tax and transaction cost considerations. It would also be worthwhile to explore weekly, monthly, or yearly rebalancing using different rolling windows, such as 100 or 252 days. Assessing how these results generalize to institutional settings or different market regimes may also offer more profound insight into the trade-off between optimization complexity and practical performance.

8. Appendix

Methodological Tools: PyPortfolioOpt and Skfolio

For portfolio optimization, I used two Python libraries: PyPortfolioOpt and Skfolio. PyPortfolioOpt was applied to compute portfolio weights for the HRP, MSR, and GMV strategies (Martin 2021), offering flexible implementations of risk- and return-based allocations. In parallel, Skfolio was used for EW, MD, and IV portfolios (Delatte and Nicolini 2023). Its compatibility with scikit-learn, built-in performance metrics, and pipeline-friendly design made it suitable for comparative analysis. Performance metrics were calculated using Skfolio's Population class (from Skfolio import Population), ensuring reproducibility through a documented, open-source package. These tools enabled the consistent construction and evaluation of diverse portfolio strategies.

8.1 Portfolio of 30 ETFs

8.1.1 Dataset Details: 30 ETFs

Table 7 provides a brief description of the ETFs, including their tickers and asset classes, all of which can be found and downloaded via the Yahoo Finance API.

Table 7: Overview of ETF asset classes, tickers, and descriptions

Asset Class	ETF Ticker	Description
U.S. Equity Sectors (9 ETFs)	XLF	Financials: banks, insurance, capital markets
	XLV	Health Care: pharmaceuticals, biotech, medical devices
	XLK	Technology: software, hardware, IT services
	XLE	Energy: oil, gas, and energy equipment
	XLI	Industrials: aerospace, transportation, construction
	XLY	Consumer Discretionary: retail, autos, entertainment
	XLU	Utilities: electric, gas, and water utilities
	XLP	Consumer Staples: food, beverages, household goods
	XLB	Materials: chemicals, packaging, metals, mining
U.S. Equity Indices (4 ETFs)	SPY	S&P 500 index ETF
	QQQ	Nasdaq-100 index ETF
	DIA	Dow Jones Industrial Average ETF
	IWM	Russell 2000 small-cap ETF
International Equity (3 ETFs)	ACWI	Global equity: All Country World Index
	EFA	Developed Markets ex-US
	EEM	Emerging Markets
Fixed Income (8 ETFs)	LQD	Investment-grade corporate bonds
	HYG	High-yield corporate bonds
	TLT	Long-term U.S. Treasuries
	AGG	Total U.S. bond market
	IEF	7–10 Year U.S. Treasuries
	MUB	Municipal bonds
	EMB	Emerging markets sovereign bonds
	BND	Broad U.S. bond market
Alternative Investments (3 ETFs)	GLD	Gold ETF (SPDR)
	SLV	Silver ETF (iShares)
	VNQ	U.S. Real Estate Investment Trusts
Cash Equivalents (3 ETFs)	BIL	1–3 Month Treasury Bills
	SHV	Ultra-short Treasuries
	SHY	1–3 Year U.S. Treasuries

8.1.2 Rolling Sharpe Ratios

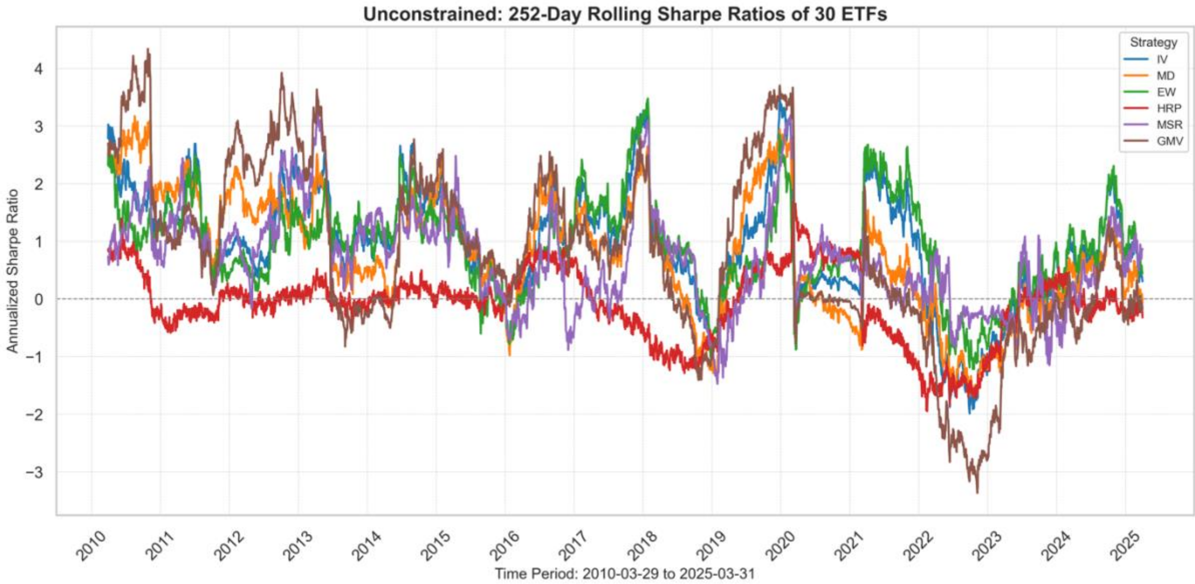

Figure 7: Unconstrained: 252-Day Rolling Sharpe Ratios (30 ETFs)

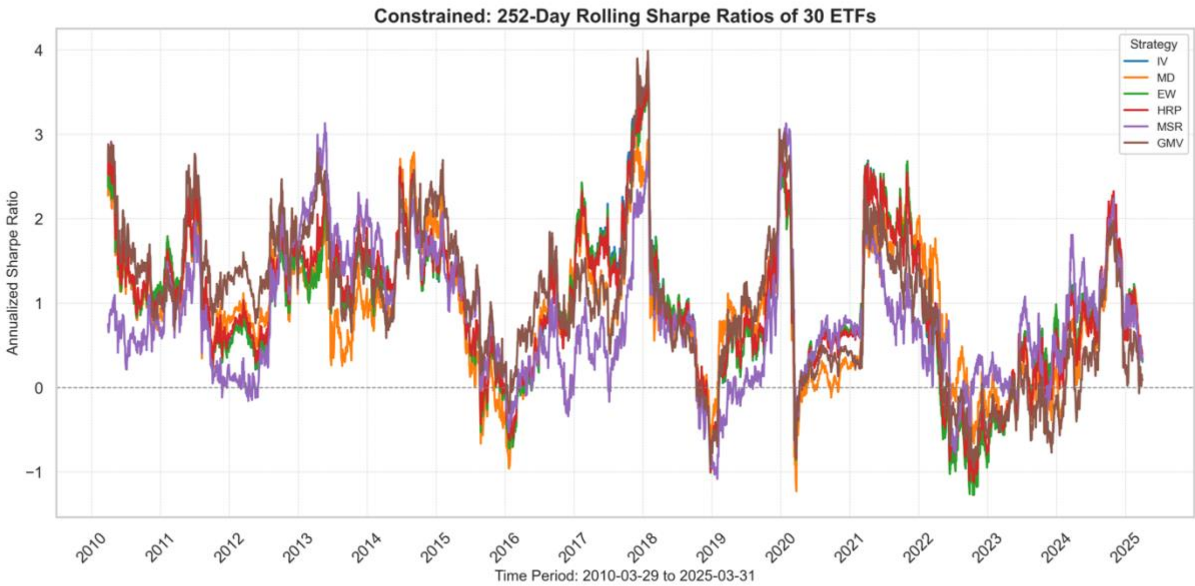

Figure 8: Constrained: 252-Day Rolling Sharpe Ratios (30 ETFs)

8.2 Portfolio of 60 ETFs

The Equal Weight (EW) strategy consistently performs well across 60 ETFs defined in Table 8 below – twice the number analyzed in the original study – based on Sharpe ratio, turnover, and computation time. In the unconstrained setting, EW achieved nearly a 500% return, as shown in Figure 9, while other optimization methods failed to surpass 100%. Although all strategies performed better in the constrained setting, EW still outperformed the rest. Its Sharpe ratio remained within a practical and reasonable range of 0.81 to 0.84. In terms of computation time, EW was completed within 3-4 seconds, and its daily turnover and number of active assets were among the lowest, ranking second only to the IV strategy in the unconstrained setting. These findings suggest that for retail investors, the potential advantages of complex optimization methods are often offset by estimation errors in expected returns and risk parameters.

8.2.1 Dataset Details: 60 ETFs

Table 8 presents the expanded 60 ETF portfolio, offering significantly greater diversification compared to the 30 ETF version. The equity category alone increased by 28 ETFs, primarily due to the inclusion of additional U.S. equity sectors and broad market indices. In contrast, other asset classes saw only modest increases or remained unchanged, primarily due to the limited availability of highly liquid ETFs that specifically track Fixed Income, Alternative Investments, and Cash Equivalents.

Table 8: Classification of 60 ETFs

Asset Class	Number of ETFs	Examples
U.S. Equity Sectors	16	Financials, Tech, Healthcare (e.g., XLF, XLK, XLV)
U.S. Equity Indices	14	S&P 500, Nasdaq-100, Small/Mid-Cap (e.g., SPY, QQQ, IWM)
International Equity	8	Global, Emerging, Europe (e.g., ACWI, EEM, VGK)
Other Equity ETFs	6	Growth, Value, Semiconductors (e.g., VUG, IWD, SOXX)
Fixed Income	8	Long and Short-Term Bonds (e.g., TLT, LQD, BND)
Alternative Investments	5	Commodities, Real Estate (e.g., GLD, VNQ, USO)
Cash Equivalents	3	Treasury Bills and Short-Term Bonds (e.g., BIL, SHY)
Total	**60 ETFs**	

8.2.2 Results: Unconstrained Optimization (60 ETFs)

Table 9: Unconstrained: Out-of-Sample Performance Metrics (60 ETFs)

Performance Metrics	EW	IV	MD	MSR	GMV	HRP
Ann Mean	11.84%	3.76%	1.39%	4.04%	1.68%	1.15%
Ann Std Dev	14.65%	3.37%	0.46%	7.99%	1.27%	0.22%
Ann Sharpe Ratio	0.81	1.12	3.02	0.51	1.32	5.16
MAX Drawdown	33.77%	7.55%	1.07%	23.87%	5.23%	0.31%
CVaR at 95%	2.19%	0.51%	0.060%	1.36%	0.16%	0.023%

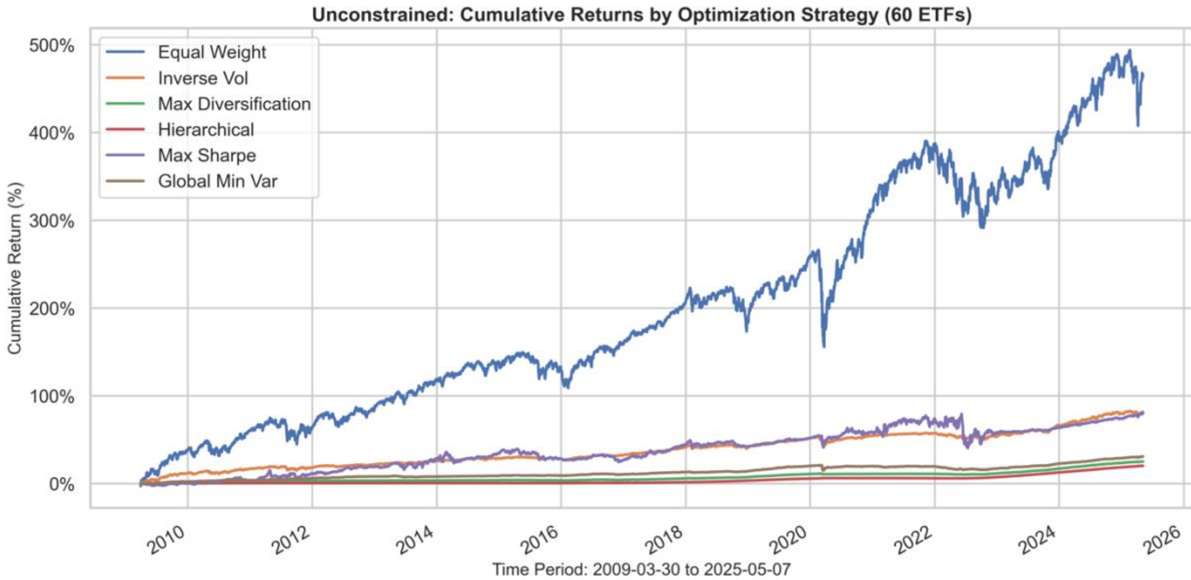

Figure 9: Unconstrained: Out-of-Sample Cumulative Portfolio Returns (60 ETFs)

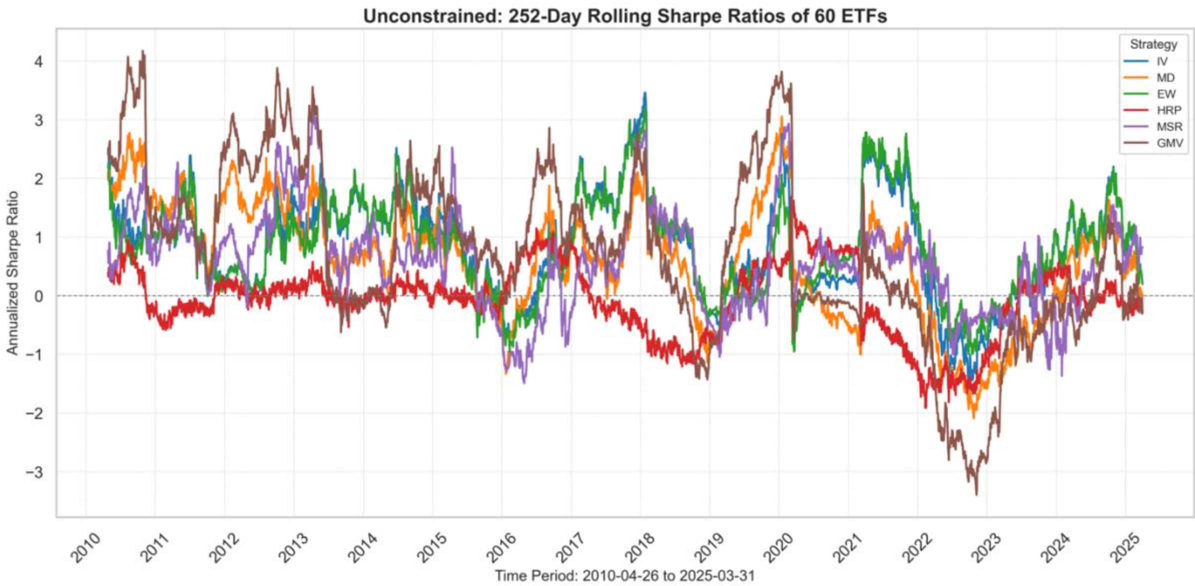

Figure 10: Unconstrained: 252-Day Rolling Sharpe Ratios (60 ETFs)

8.2.3 Results: Constrained Optimization (60 ETFs)

Table 10: Constrained: Out-of-Sample Performance Metrics (60 ETFs)

Performance Metrics	EW	IV	MD	MSR	GMV	HRP
Ann Mean	9.44%	9.20%	8.41%	7.29%	7.54%	8.80%
Ann Std Dev	11.21%	10.66%	10.83%	13.12%	8.06%	10.03%
Ann Sharpe Ratio	0.84	0.86	0.78	0.56	0.94	0.88
MAX Drawdown	26.83%	25.99%	28.68%	26.02%	22.01%	25.21%
CVaR at 95%	1.66%	1.59%	1.59%	2.07%	1.17%	1.49%

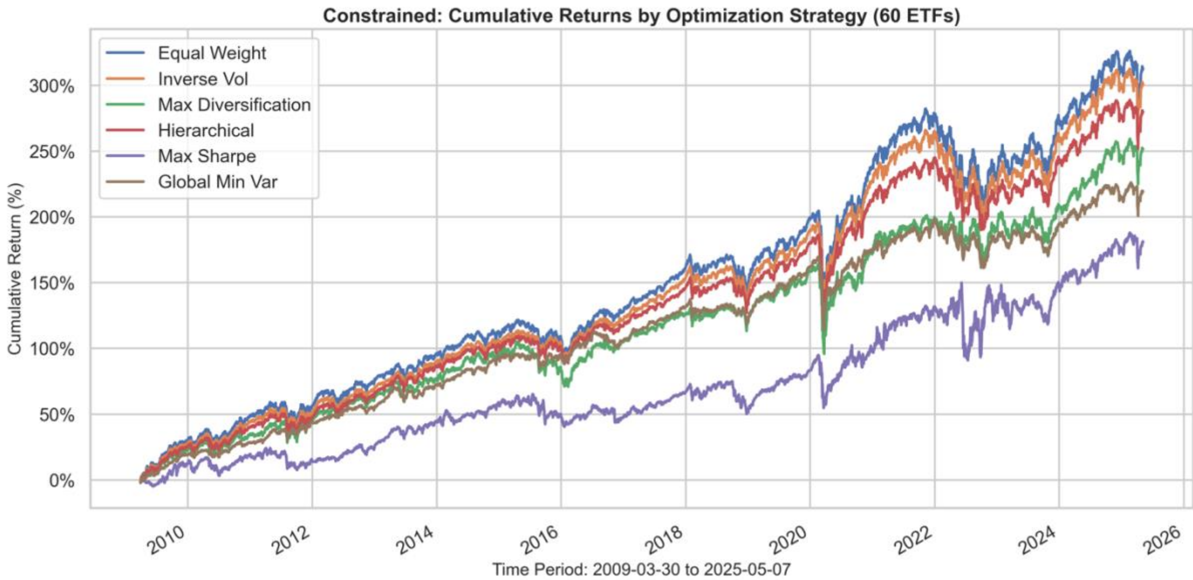

Figure 11: Constrained: Out-of-Sample Cumulative Portfolio Returns (60 ETFs)

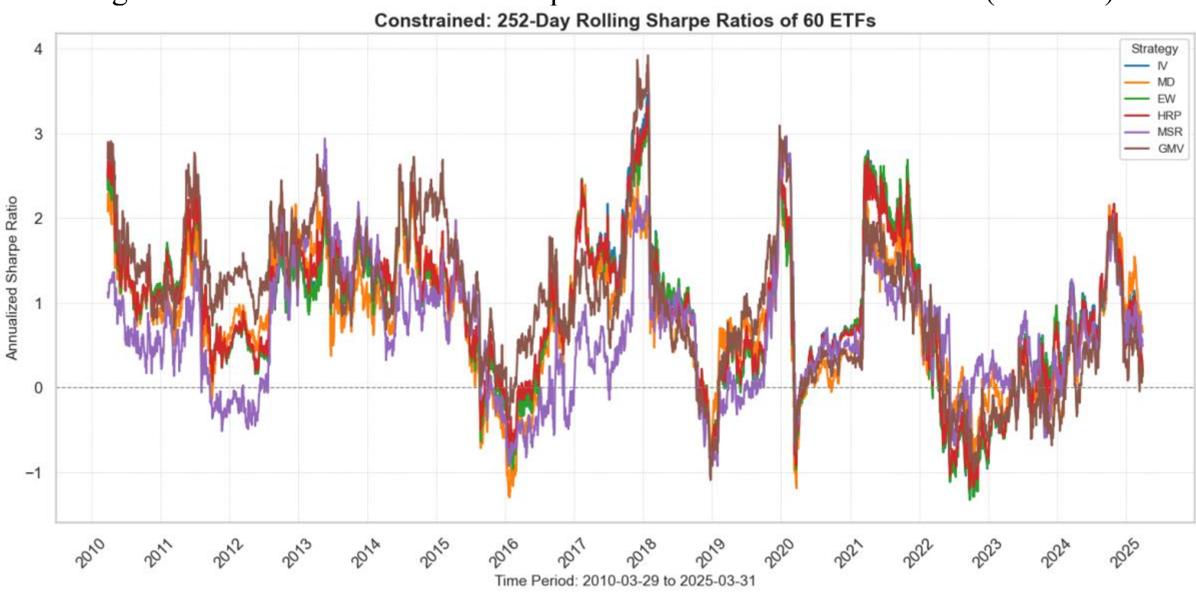

Figure 12: Unconstrained: 252-Day Rolling Sharpe Ratios (60 ETFs)

8.2.4 Cumulative Computation Time (60 ETFs)

Table 11: Computation Time in Seconds (60 ETFs)

Optimizers	Unconstrained	Constrained
EW	4	3
IV	22	39
MSR	137	295
GMV	79	241
MD	109	311
HRP	632	555

8.2.5 Average Active Assets and Turnover (60 ETFs)

Table 12: Average Active Assets and Turnover (60 ETFs)

Optimizers	Unconstrained		Constrained	
	Avg. Active Assets	Avg. Daily Turnover	Avg. Active Assets	Avg. Daily Turnover
EW	60.00	0.59%	54.00	0.59%
IV	60.00	0.45%	54.00	0.62%
HRP	9.90	0.73%	54.00	3.39%
MD	14.10	1.69%	18.51	2.22%
GMV	11.68	0.62%	15.45	1.88%
MSR	9.91	10.94%	10.66	15.49%

8.3 Results: Quarterly Rebalanced (30 ETFs)

8.3.1 Results: Quarterly Rebalanced, Unconstrained Optimization (30 ETFs)

Table 13: Unconstrained: Quarterly Reb, Out-of-Sample Performance Metrics (30 ETFs)

Performance Metrics	EW	IV	MD	MSR	GMV	HRP
Ann Mean	9.55%	2.46%	1.39%	4.08%	1.69%	1.14%
Ann Std Dev	10.64%	1.76%	0.53%	7.15%	0.98%	0.22%
Ann Sharpe Ratio	0.90	1.40	2.62	0.57	1.73	5.13
MAX Drawdown	24.70%	4.29%	1.88%	13.03%	3.28%	0.35%
CVaR at 95%	1.60%	0.26%	0.66%	1.19%	0.13%	0.023%

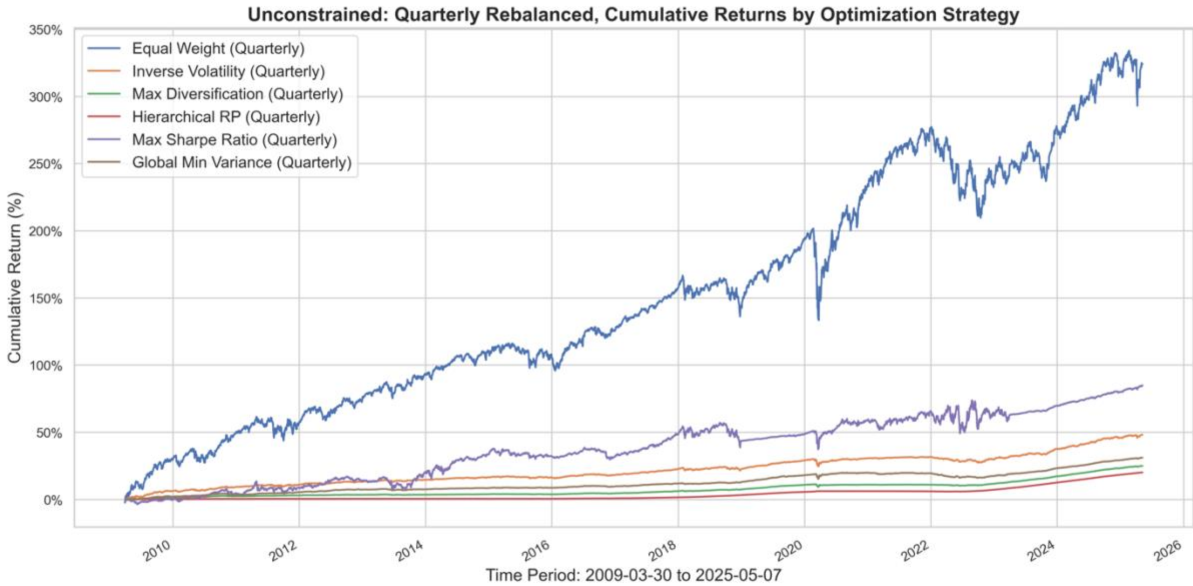

Figure 13: Constrained: Quarterly Reb, Out-of-Sample Cumulative Portfolio Returns (30 ETFs)

8.3.2 Results: Quarterly Rebalanced, Constrained Optimization (30 ETFs)

Table 14: Constrained: Quarterly Reb, Out-of-Sample Performance Metrics (30 ETFs)

Performance Metrics	EW	IV	MD	MSR	GMV	HRP
Ann Mean	9.53%	9.24%	8.29%	9.20%	7.73%	8.93%
Ann Std Dev	10.52%	10.12%	10.25%	12.70%	8.56%	9.70%
Ann Sharpe Ratio	0.91	0.91	0.81	0.72	0.90	0.92
MAX Drawdown	24.51%	24.74%	30.07%	26.52%	26.56%	25.46%
CVaR at 95%	1.58%	1.52%	1.52%	1.99%	1.24%	1.45%

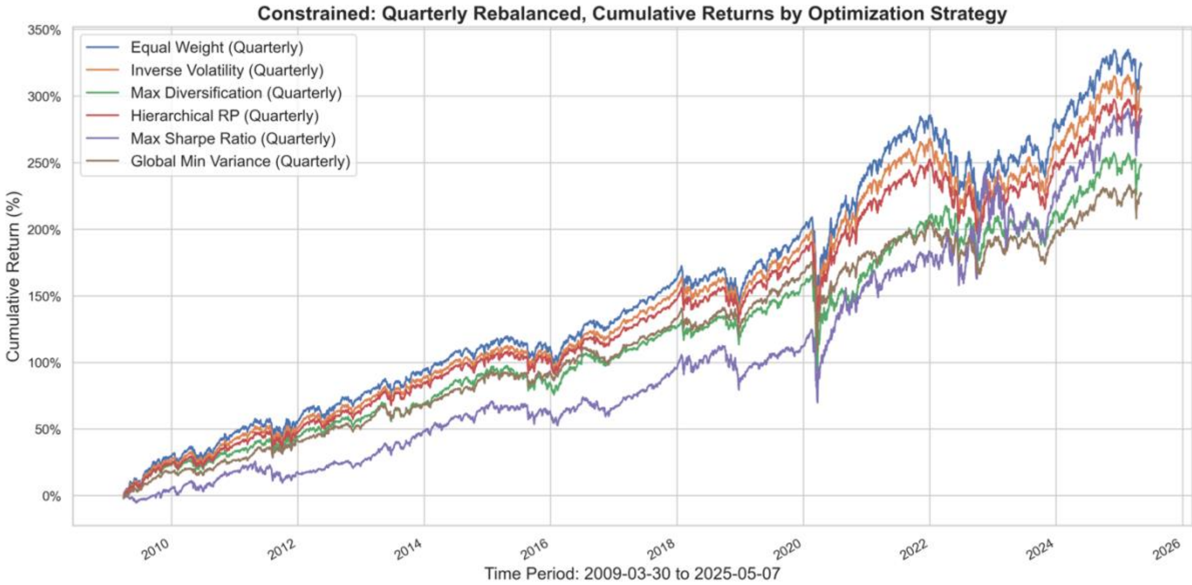

Figure 14: Constrained: Quarterly Reb, Out-of-Sample Cumulative Portfolio Returns (30 ETFs)

8.3.3 Quarterly Reb: Cumulative Computation Time (30 ETFs)

Table 15: Quarterly Reb, Computation Time in Seconds (30 ETFs)

Optimizers	Unconstrained	Constrained
EW	4	4
IV	6	5
GMV	6	12
MSR	7	13
MD	8	16
HRP	12	9

8.3.4 Quarterly Reb: Average Active Assets and Turnover (30 ETFs)

Table 16: Quarterly Reb, Average Active Assets and Turnover (30 ETFs)

Optimizers	**Unconstrained**		**Constrained (Quarterly)**	
	Avg. Active Assets	**Avg. Daily Turnover**	**Avg. Active Assets**	**Avg. Daily Turnover**
EW	30.00	0.06%	30.00	0.06%
IV	30.00	0.08%	30.00	0.08%
HRP	10.85	0.13%	30.00	0.18%
MD	12.15	0.19%	16.42	0.31%
GMV	9.41	0.12%	12.48	0.34%
MSR	8.29	0.96%	8.98	1.04%

Reference List

Aroussi, Ran. 2025. yFinance Documentation. Accessed June 7, 2025. https://ranaroussi.github.io/yfinance/index.html

Benartzi, Shlomo, and Richard Thaler. "Heuristics and Biases in Retirement Savings Behavior." Journal of Economic Perspectives 21, no. 3 (2007): 81–104. https://doi.org/10.1257/jep.21.3.81

Carvalho, Raul Leote de, Lu Xiao, and Pierre Moulin. 2012. "Demystifying Equity Risk-Based Strategies: A Simple Alpha Plus Beta Description." The Journal of Portfolio Management 38 (3): 56–70. https://ssrn.com/abstract=1949003

Choueifaty, Yves, Tristan Froidure, and Julien Reynier. 2013. "Properties of the Most Diversified Portfolio." Journal of Investment Strategies 2 (2): 49–70. https://doi.org/10.2139/ssrn.1895459

Choueifaty, Yves. 2008. "Towards Maximum Diversification." SSRN Electronic Journal. https://doi.org/10.2139/ssrn.4063676

De Miguel, Victor, Lorenzo Garlappi, and Raman Uppal. 2009. "Optimal Versus Naïve Diversification: How Inefficient is the 1/N Portfolio Strategy?" The Review of Financial Studies 22 (5): 1915–1953. https://doi.org/10.1093/rfs/hhm075

Delatte, Hugo, and Carlo Nicolini. 2023. Skfolio. GitHub. https://github.com/skfolio/skfolio

Jacobs, Heiko, Sebastian Müller, and Martin Weber. 2014. "How Should Individual Investors Diversify? An Empirical Evaluation of Alternative Asset Allocation Policies." Journal of Financial Markets 19 (June): 62–85. https://doi.org/10.1016/j.finmar.2013.07.004

Ledoit, Olivier, and Michael Wolf. 2004. "A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices." Journal of Multivariate Analysis 88 (2): 365–411. https://doi.org/10.1016/S0047-259X(03)00096-4

López de Prado, Marcos. 2016. "Building Diversified Portfolios That Outperform Out of Sample." The Journal of Portfolio Management 42 (4): 59–69. https://doi.org/10.3905/jpm.2016.42.4.059

Markowitz, Harry. 1952. "Portfolio Selection." The Journal of Finance 7 (1): 77–91. https://doi.org/10.2307/2975974

Martin, Robert A. 2021. "PyPortfolioOpt: Portfolio Optimization in Python." Journal of Open Source Software 6 (61): 3066. https://doi.org/10.21105/joss.03066

Plyakha, Yuliya, Raman Uppal, and Grigory Vilkov. 2012. "Why Does an Equal-Weighted Portfolio Outperform Value- and Price-Weighted Portfolios?" SSRN Electronic Journal. http://dx.doi.org/10.2139/ssrn.2724535

PwC. 2024. ETF 2024: Reshaping Around the Investor. Accessed June 6, 2025. https://www.pwc.com/gx/en/industries/financial-services/publications/etf-survey.html

Rockafellar, R. Tyrrell, and Stanislav Uryasev. 2000. "Optimization of Conditional Value-at-Risk." Journal of Risk 2: 21–24. https://doi.org/10.21314/JOR.2000.038

Sharpe, William F. 1966. Mutual Fund Performance. The Journal of Business 39 (1): 119–138. https://doi.org/10.1086/294846

State Street Global Advisors. 2023. Why ETF Growth Is Booming. Featuring data from Morningstar, as of December 31, 2023. Accessed June 6, 2025. https://www.ssga.com/us/en/intermediary/insights/why-etf-growth-is-booming

TradingView. 2025. Most Traded ETFs. Accessed June 7, 2025. https://www.tradingview.com/markets/etfs/funds-most-traded/