

UNIVERSITÀ
DI PAVIA

Dipartimento di Scienze Economiche e

Aziendali Corso di Laurea magistrale in

Finance

Revisiting the 1/N Portfolio: Estimation Risk and Robust Portfolio Strategies

Relatore:

Chiar.mo Prof. Brignone Riccardo

Tesi di Laurea
di Doveri Federica

Matr. n.543463

Anno Accademico 2024-2025

INDICES

CHAPTER 1 - INTRODUCTION 1
1.1 The Portfolio Selection Problem 1
1.2 The Development of Modern Portfolio Theory 2
1.3 Diversification and the Role of Covariance 3
1.4 Practical Challenges: Estimation Risk 3
1.5 The Puzzle of Naïve Diversification 5
1.6 Research Objective 6
1.7 Structure of the Thesis 7
CHAPTER 2 - MODERN PORTFOLIO THEORY 9
2.1 Investor Preferences and the Mean-Variance Framework 9
2.2 Portfolio Risk and the Role of Covariance 10
2.3 The Principle of Diversification 11
2.4 The Efficient Frontier 13
2.5 The Minimum Variance Portfolio 15
2.6 The Risk-Free Asset and the Tangency Portfolio 16
2.7 Limitations of Mean–Variance Optimization 18
CHAPTER 3 - ESTIMATION RISK AND IMPLEMENTATION CHALLENGES 21
3.1 The Nature of Estimation Risk 21
3.2 Sensitivity of Optimal Portfolio Weights 22
3.3 Parameter Uncertainty and Forecasting Limitations 24
3.4 Rolling-Window Estimation 26
3.5 Out-of-Sample Performance Evaluation 28
3.6 Predictive Variance and Long-Horizon Risk 30
3.7 Consequences of Estimation Risk for Portfolio Optimization 31
3.8 The Need for Robust Portfolio Strategies 33
CHAPTER 4 - PORTFOLIO OPTIMIZATION MODELS 36
4.1 Classical Mean–Variance Optimization 36
4.2 Minimum Variance Portfolio 37
4.3 Portfolio Constraints and Practical Implementation 39
4.4 Shrinkage Estimators and Bayesian Approaches 41
4.5 Robust Optimization Methods 43
4.6 Naïve Diversification Strategies 46
4.7 Combining Optimization and Robustness 47
CHAPTER 5 - THE DEMIGUEL, GARLAPPI, UPPAL (2009) STUDY 50
5.1 Research Motivation 50
5.2 Portfolio Strategies Examined 51

5.3 Data Used in the Study 53
5.4 Methodology: Rolling-Window Estimation 54
5.5 Performance Evaluation Metrics 55
5.6 Main Empirical Findings 56
5.7 Interpretation of Results 57
5.8 Implications for Portfolio Management 58
CHAPTER 6 - EMPIRICAL REPLICATION WITH UPDATED DATA 60
6.1 Introduction 60
6.2 Datasets and Comparison with the Original Study 61
6.2.1 The Fama-French Datasets 61
6.2.2 The GICS Sector Indices 61
6.2.3 The MSCI Country Indices 62
6.3 Performance Metrics 63
6.4 Results: Ten Industry Portfolios 63
6.4.1 Cumulative Out-of-Sample Returns 63
6.4.2 Sharpe Ratio and CEQ Analysis 64
6.4.3 Turnover 66
6.5 Results: Twenty-Five Portfolios 67
6.5.1 Cumulative Out-of-Sample Returns 67
6.5.2 Sharpe Ratio and CEQ 68
6.5.3 Turnover 70
6.6 Results: GICS Sector Indices 72
6.6.1 Institutional Context 72
6.6.2 Cumulative Out-of-Sample Returns 73
6.6.3 Sharpe Ratio and CEQ 74
6.6.4 Turnover 75
6.7 Results: MSCI Country Indices 76
6.7.1 The International Dimension 76
6.7.2 Cumulative Out-of-Sample Returns 77
6.7.3 Sharpe Ratio and CEQ 78
6.7.4 Turnover 80
6.8 Cross-Dataset Synthesis 81
6.9 Turnover and Practical Implementability 82
6.10 Summary 83
CHAPTER 7 - CONCLUSION 85
7.1 Summary of the Main Findings 85
7.2 What Has Changed Since 2009 86

7.3 Implications for Portfolio Management 87
7.4 Limitations and Future Research 88
7.5 Concluding Remarks 89

CHAPTER 1 - INTRODUCTION

1.1 The Portfolio Selection Problem

Portfolio selection is one of the central topics in financial economics. Investors must decide how to allocate their wealth across assets whose future returns are uncertain, and every investment choice involves a fundamental trade-off between expected return and risk. In this context, risk is typically defined as the variability of returns, capturing the possibility that realized outcomes differ from expectations. It is most commonly measured through the variance or standard deviation of returns.

Traditional investment approaches often evaluate securities individually, focusing on their expected profitability and risk characteristics. This perspective changed profoundly with the development of Modern Portfolio Theory, which demonstrated that investment decisions should be made at the portfolio level rather than by analyzing assets in isolation. The key insight is that portfolio risk depends not only on the volatility of individual assets but also on how their returns move together. Covariance captures these relationships and determines how assets interact within a portfolio.

This shift in perspective fundamentally changed the way investment decisions are made. Instead of identifying individual "good" securities, investors must consider how each asset contributes to the overall risk–return profile of the portfolio. Diversification therefore becomes a central principle of portfolio management: by combining assets whose returns are not perfectly correlated, investors can reduce overall risk without necessarily sacrificing expected return.

Understanding this conceptual shift is essential, as it provides the theoretical foundation for all subsequent developments in modern asset allocation theory. It highlights that risk is a property

of the portfolio as a whole and that effective portfolio construction requires analyzing how assets behave jointly rather than separately.

1.2 The Development of Modern Portfolio Theory

Modern Portfolio Theory originated with the pioneering work of Harry Markowitz (1952), who formalized the idea that investment decisions should be based on the joint distribution of asset returns. Within this framework, investors evaluate portfolios according to expected return and risk, and choose allocations that minimize variance for a given expected return or maximize expected return for a given level of risk.

A key implication of this approach is the concept of the efficient frontier, which represents the set of portfolios offering the best attainable trade-offs between risk and return. Portfolios below the frontier are inefficient because alternative allocations exist that provide higher returns for the same level of risk.

Subsequent contributions extended the framework. Tobin's separation theorem (1958) showed that all investors should combine a risk-free asset with a single optimal risky portfolio, leading to the development of the Capital Asset Pricing Model (CAPM). Later models, such as the zero-beta CAPM and multifactor models, incorporated additional sources of systematic risk.

At the same time, researchers recognized that the practical implementation of mean–variance optimization is highly sensitive to estimation error. Small inaccuracies in expected returns or covariances can lead to large and unstable changes in optimal portfolio weights. This fragility motivated the development of alternative approaches, such as Bayesian methods, shrinkage estimators, and robust optimization, designed to mitigate the impact of parameter uncertainty.

Despite its limitations, the mean–variance framework remains a cornerstone of financial economics and continues to influence both academic research and investment practice.

1.3 Diversification and the Role of Covariance

One of the most important contributions of Modern Portfolio Theory is the formal explanation of diversification. Portfolio risk is not simply the average of individual asset risks; it depends critically on the covariances between asset returns. When assets are not perfectly positively correlated, combining them in a portfolio reduces overall risk because losses in some assets may be offset by gains in others.

Covariance and correlation quantify the degree to which asset returns move together. As the number of assets increases, the importance of individual asset risk declines, and the average covariance becomes the main determinant of total portfolio risk. This explains why diversification cannot eliminate all risk: it removes idiosyncratic risk, but systematic risk remains.

The covariance matrix therefore plays a central role in portfolio construction, summarizing the relationships among asset returns and determining the effectiveness of diversification. At the same time, estimating covariances accurately is challenging, especially in high-dimensional settings, and estimation error can distort the perceived relationships among assets. These considerations highlight both the strengths and the limitations of diversification and motivate the empirical analysis developed in later chapters.

1.4 Practical Challenges: Estimation Risk

A central limitation of mean–variance optimization lies in its reliance on estimated inputs. Expected returns, variances, and covariances must be inferred from historical data, and these estimates are inevitably subject to sampling variability. Because historical datasets are finite and often noisy, the estimated parameters may differ substantially from their true population values, introducing uncertainty into the optimization process. This phenomenon, known as

estimation risk, represents one of the most significant obstacles to the practical implementation of portfolio optimization.

Among all parameters, expected returns are particularly difficult to estimate accurately. They exhibit high volatility, low signal-to-noise ratios, and frequent structural breaks, making them extremely unstable over time. Even small errors in expected return estimates can lead to disproportionately large changes in optimal portfolio weights. This occurs because the optimization procedure tends to overweight assets with slightly higher estimated returns, even when these differences are not statistically meaningful. As a result, optimized portfolios often become highly concentrated and unstable, with weights that fluctuate dramatically as new data become available.

Estimation risk also affects the covariance matrix, although to a lesser extent. Covariances are generally more stable than expected returns, but their estimation becomes increasingly challenging as the number of assets grows. The number of covariances to estimate increases quadratically with the size of the portfolio, and when the estimation window is short relative to the number of assets, the sample covariance matrix may become ill-conditioned or nearly singular. In such cases, the inversion of the covariance matrix amplifies noise rather than filtering it out, producing extreme and erratic portfolio weights. This instability affects not only the weights but also the shape and position of the efficient frontier, which becomes highly sensitive to small changes in the input data.

These issues highlight the fundamental gap between theoretical optimality and empirical performance. Portfolios that appear optimal in sample may perform poorly out of sample because the optimization process has effectively overfit historical noise rather than capturing the underlying return structure. The severity of estimation risk has motivated a large body of research aimed at developing more robust and stable portfolio construction techniques. Approaches such as shrinkage estimators, Bayesian methods, robust optimization, and

factor-based models attempt to mitigate the impact of estimation error by incorporating prior information, reducing dimensionality, or imposing structure on the covariance matrix.

These methodological advances reflect a broader recognition: while mean–variance optimization provides a powerful theoretical framework, its practical success depends critically on the quality and stability of the estimated inputs. Understanding these limitations is essential for explaining why simple allocation rules, such as the 1/N strategy, may outperform optimized portfolios in empirical settings—a phenomenon explored in the following section.

1.5 The Puzzle of Naïve Diversification

Empirical research has shown that simple diversification strategies often perform surprisingly well compared to theoretically optimized portfolios. In particular, the naïve equal-weighting strategy, which allocates the same weight to each asset, has been found to be extremely difficult to outperform in out-of-sample tests.

This occurs because the equal-weighting strategy does not rely on estimated parameters. By avoiding the estimation of expected returns and the covariance matrix, it sidesteps the instability and noise inherent in sample-based optimization. Although it discards potentially valuable information, its simplicity protects it from the distortions caused by estimation error, allowing it to achieve competitive performance in real-world applications.

Naïve diversification also benefits from the natural reduction of idiosyncratic risk that occurs when multiple assets are combined in a portfolio. As the number of assets increases, the impact of asset-specific shocks declines, allowing the portfolio to achieve relatively stable performance even without precise parameter estimates. This property helps explain why equal-weighting often performs well across different markets and time periods.

This empirical regularity poses a challenge to traditional portfolio theory. According to the mean–variance framework, optimized portfolios should dominate simple allocation rules by

exploiting information about expected returns and covariances. Yet, in practice, the opposite often occurs. This discrepancy is known as the "diversification puzzle", because it contradicts the theoretical prediction that optimization should yield superior risk-adjusted performance.

The difficulty of consistently outperforming the 1/N strategy, as documented by De Miguel, Garlappi, and Uppal (2009), highlights the severe impact of estimation error on practical portfolio optimization. When parameter uncertainty is sufficiently large, the theoretical advantages of optimization are overshadowed by the instability of the estimated inputs. In such cases, simplicity and robustness may dominate theoretical optimality.

The persistence of this puzzle has motivated extensive research into alternative portfolio construction methods that emphasize stability, shrinkage, Bayesian updating, or robust optimization. These approaches reflect a broader recognition that, in environments characterized by high uncertainty and limited data, naïve diversification can serve as a surprisingly effective benchmark.

This empirical puzzle motivates the need to reassess the practical value of optimization techniques, which is the central objective of this thesis.

1.6 Research Objective

The objective of this thesis is to evaluate the practical performance of portfolio optimization strategies in the presence of estimation risk. The study replicates and extends the influential analysis by De Miguel, Garlappi, and Uppal (2009), *"Optimal versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?"*, using a longer and more recent dataset covering the period from 1990 to 2025. The extension of the original study to a more recent and longer time period allows for an evaluation of whether improvements in data availability, longer estimation windows, and more stable financial markets have reduced the magnitude of estimation risk over time.

By analyzing modern data over an extended period, the thesis investigates whether the challenges documented in the original study persist in contemporary markets. In particular, it examines whether optimized portfolios, such as classical mean–variance strategies, minimum-variance portfolios, and constrained optimization methods, perform better relative to simple allocation rules when more historical information is available. This comparison is crucial for determining whether advances in empirical finance, computational tools, and data quality have significantly improved the reliability of optimized portfolio strategies.

The analysis evaluates a broad set of portfolio construction methods, including traditional mean–variance optimization, minimum-variance portfolios, constrained approaches, and the naïve 1/N allocation rule. Portfolio performance is assessed using out-of-sample metrics such as the Sharpe ratio, certainty-equivalent returns, and portfolio turnover. These measures provide a comprehensive evaluation of practical portfolio efficiency and allow for a direct comparison between theoretically optimal and empirically robust strategies.

Through this empirical investigation, the thesis contributes to the ongoing debate on whether modern data availability and methodological advances have reduced the practical limitations of portfolio optimization. More broadly, it seeks to clarify whether the theoretical advantages of optimal diversification can be realized in practice, or whether simple heuristic strategies continue to offer a more reliable benchmark in environments characterized by significant parameter uncertainty.

1.7 Structure of the Thesis

The remainder of this thesis is organized to provide a clear progression from theoretical foundations to empirical analysis. Chapter 2 presents the theoretical background of Modern Portfolio Theory, including the principles of mean–variance optimization and the role of

diversification. Chapter 3 examines the problem of estimation risk and discusses the practical challenges that arise when implementing portfolio optimization using historical data. Chapter 4 reviews alternative portfolio construction methods designed to improve robustness and mitigate the impact of estimation error.

Chapter 5 outlines the empirical framework developed in previous studies, with particular attention to the methodology introduced by De Miguel, Garlappi, and Uppal (2009). Chapter 6 presents the empirical replication using updated European data, comparing the performance of optimized portfolios with that of the naïve 1/N strategy across multiple datasets and evaluation metrics. Chapter 7 concludes the thesis by summarizing the main findings and discussing their implications for portfolio management, as well as potential directions for future research.

CHAPTER 2 - MODERN PORTFOLIO THEORY

2.1 Investor Preferences and the Mean-Variance Framework

Modern Portfolio Theory provides a formal analytical framework for studying investment decisions under uncertainty. Because future asset returns cannot be known with certainty, investors must base their choices on expectations about risk and return rather than on deterministic outcomes. Within this framework, portfolios are evaluated according to two fundamental characteristics: expected return, representing the average payoff an investor anticipates, and risk, typically measured by the variance or standard deviation of returns.

The mean–variance model is built on a set of simplifying assumptions that make the analysis easier to develop. Financial markets are assumed to be frictionless, with no transaction costs, taxes, or restrictions on trading. Investors are considered rational, risk-averse, and focused only with the mean and variance of portfolio returns over a single investment period.

Moreover, the model assumes that asset returns are normally distributed or, equivalently, that investor preferences can be represented by a quadratic utility function. Under these conditions, the mean and variance of returns become sufficient statistics for investment decisions, reducing the complex problem of uncertainty to a two-dimensional trade-off between risk and return.

Although these assumptions are stylized and do not fully capture the frictions present in real markets, such as transaction costs, taxes, liquidity constraints, and institutional limits, the simplified framework remains extremely valuable because it captures the behavior of risk-averse investors: they prefer higher expected returns and lower levels of risk. Such investors aim to maximize expected return for a given level of risk or, equivalently, minimize risk for a given expected return.

A portfolio is defined by a vector of weights assigned to each asset. Let w denote the vector of portfolio weights, μ the vector of expected returns, and Σ the covariance matrix of asset returns. The expected return of a portfolio is given by:

$E(R_p) = w^T \mu$

while the variance of portfolio returns is:

$\sigma_p^2 = w^T \Sigma w$

These expressions highlight a key feature of portfolio theory: portfolio risk depends not only on the individual variances of assets but also on the covariances between them. As a result, portfolio selection becomes an optimization problem in which investors choose the set of weights that best satisfies their risk–return preferences, subject to the constraint that total portfolio weights sum to one, representing the allocation of total available wealth.

The relationships among expected returns, variances, and covariances form the foundation of modern portfolio theory. They demonstrate why portfolio risk cannot be understood by examining assets in isolation and why diversification plays a central role in investment decisions. This conceptual foundation sets the stage for the next sections, which explore the role of covariance, the principle of diversification, and the construction of efficient portfolios.

2.2 Portfolio Risk and the Role of Covariance

A central insight of modern portfolio theory is that the risk of a portfolio is not simply the sum of the risks of its individual assets. While the variance of each asset captures its own uncertainty, portfolio risk depends fundamentally on how asset returns move together, a relationship measured by covariance. This means that the co-movement of returns, rather than individual volatilities alone, determines the overall variability of a portfolio.

Covariance can be standardized into the correlation coefficient, which ranges from –1 to +1 and provides an intuitive measure of the strength and direction of co-movement between asset returns. When assets are positively correlated, they tend to move in the same direction, reducing the benefits of diversification. By contrast, when correlations are low or negative, combining

assets in a portfolio can substantially reduce overall risk, because losses in some assets may be offset by gains in others.

Formally, consider a portfolio of N assets with weights wi and returns Ri. The variance of portfolio returns is given by:

$\sigma_p^2 = \sum w_i^2 \sigma_i^2 + \sum\sum w_i w_j \sigma_{ij}$

where σ_i^2 represents the variance of asset i and σ_{ij} denotes the covariance between the returns of assets i and j.

This expression shows that portfolio risk depends both on the weighted variances of individual assets and the covariances between all pairs of assets.

As the number of assets increases, the contribution of individual asset risk becomes less important, and the average covariance among assets becomes the main driver of total portfolio risk. This explains why diversification cannot eliminate all risk: while idiosyncratic risk can be diversified away, systematic co-movement across assets remains.

The covariance matrix summarizes all pairwise relationships among asset returns and therefore plays a central role in portfolio optimization. Accurate estimation of this matrix is essential for determining the risk characteristics of different portfolio combinations and for constructing efficient portfolios. This insight represents one of the most important conceptual contributions of modern portfolio theory: risk is a property of the portfolio as a whole, not of its individual components.

2.3 The Principle of Diversification

Diversification is a central principle of portfolio theory and refers to the practice of spreading investments across multiple assets to reduce overall risk. The idea is that assets do not react to economic conditions in exactly the same way or at the same time. By combining assets whose

returns are not perfectly correlated, investors can lower the volatility of their portfolios, because losses in some assets may be offset by gains in others.

In the extreme theoretical case of perfect negative correlation, it is possible to construct a portfolio with zero variance by choosing appropriate combinations of asset weights. Although such perfect hedging opportunities rarely exist in practice, the principle illustrates how powerful the risk-reducing effect of diversification can be.

Diversification, however, does not eliminate all forms of risk. It primarily reduces unsystematic risk, which arises from firm-specific or asset-specific factors. As the number of assets in a portfolio increases, this idiosyncratic component becomes progressively less important, because individual shocks tend to cancel out. Systematic risk, by contrast, reflects market-wide forces affecting all assets and cannot be diversified away. As a result, the marginal benefit of adding new assets declines as the portfolio becomes larger, eventually leaving systematic risk as the dominant component of total portfolio risk.

This distinction between systematic and unsystematic risk is fundamental in finance because it determines which risks investors are compensated for bearing. Since unsystematic risk can be eliminated through diversification without reducing expected return, competitive financial markets do not reward investors for holding it. Systematic risk, instead, cannot be avoided, and investors require compensation in the form of higher expected returns. This principle forms the foundation of equilibrium asset pricing models, most notably the Capital Asset Pricing Model (CAPM).

The effectiveness of diversification depends critically on the degree of correlation among asset returns. The lower the correlation between assets, the greater the potential reduction in portfolio risk. This explains why broad diversification across sectors, regions, and asset classes is widely regarded as a core investment strategy: it maximizes exposure to uncorrelated sources of return while minimizing vulnerability to asset-specific shocks.

From an economic perspective, diversification can be viewed as a transformation process. By combining many assets, investors convert unpredictable firm-specific risk into systematic market risk: risk that cannot be eliminated but can be managed more efficiently. This insight forms a conceptual bridge between portfolio theory and modern asset pricing, highlighting why diversification is not only a practical guideline but a foundational principle of financial economics.

2.4 The Efficient Frontier

The mean–variance framework naturally leads to one of the most influential ideas in modern portfolio theory: the efficient frontier. This curve represents the set of portfolios that offer the highest expected return for each given level of risk, or equivalently, the lowest risk for each level of expected return. It formalizes the idea that rational, risk-averse investors should only consider portfolios that lie on this boundary, since any portfolio below it is dominated by another portfolio with a strictly better risk–return trade-off.

In a two-dimensional risk–return space, each feasible portfolio can be plotted according to its expected return and variance. The set of all feasible portfolios forms a region, and the efficient frontier corresponds to the upper boundary of this region. Portfolios on the lower boundary are inefficient because there exists alternative portfolios that either increase expected return without increasing risk or reduce risk without lowering expected return.

The efficient frontier is derived by solving a constrained quadratic optimization problem. Formally, investors minimize portfolio variance subject to achieving a target expected return and satisfying the budget constraint:

$\min_{(w)} w^T \Sigma w$

subject to:

$w^T\mu=R^*$

$w^T 1 = 1$

By varying the target return R*, one obtains a continuous set of optimal portfolios that trace out a smooth, concave curve in mean–variance space. The lower portion of this curve is inefficient, while the upper portion constitutes the set of portfolios preferred by rational, risk-averse investors.

The shape of the efficient frontier reflects the benefits of diversification. When assets are combined, especially those with low or negative correlations, portfolio risk initially declines sharply. This steep reduction illustrates how diversification can substantially lower unsystematic risk. As more assets are added, however, the marginal benefits of diversification diminish, and the frontier becomes flatter. At this stage, systematic risk dominates, and further diversification yields only modest improvements in risk reduction.

Each investor selects a specific portfolio along the efficient frontier according to their degree of risk aversion. More risk-averse investors tend to choose portfolios near the minimum variance portfolio, which is the leftmost point on the frontier and represents the lowest achievable level of risk. Less risk-averse investors select portfolios further to the right, accepting higher risk in exchange for higher expected returns.

A portfolio is considered efficient when no other portfolio offers a higher expected return for the same level of risk or a lower level of risk for the same expected return. The efficient frontier therefore provides a complete representation of all optimal portfolio choices available to investors within the mean–variance framework.

Beyond its theoretical foundation, the efficient frontier has important practical implications. It highlights the central role of covariance in determining portfolio risk, illustrates the economic value of diversification, and provides a benchmark against which real-world portfolios can be evaluated. At the same time, its construction depends critically on accurate estimates of

expected returns and covariances, a challenge that becomes central in empirical applications and motivates the discussion of estimation risk in later sections.

2.5 The Minimum Variance Portfolio

A key reference point on the efficient frontier is the minimum variance portfolio (MVP), the portfolio that achieves the lowest possible level of risk among all feasible portfolios, regardless of expected return. It represents the leftmost point on the efficient frontier and is central to understanding the trade-off between risk and return in the mean–variance framework.

The MVP is obtained by solving the optimization problem without imposing a target return constraint:

$\min_{(w)} w^T \Sigma w$

subject to:

$w^T 1 = 1$

This formulation highlights a distinctive feature of the minimum variance portfolio: its construction depends solely on the covariance matrix of asset returns. Unlike full mean–variance optimization, it does not require estimates of expected returns, which are often difficult to estimate accurately and are subject to substantial sampling error. As a result, the MVP is generally more stable and less sensitive to estimation error than portfolios that rely on both first and second moments.

The economic intuition behind the MVP is straightforward. Expected returns are highly noisy and unstable over time, and relying on them can lead to extreme and erratic portfolio weights. Covariances, instead, are typically more stable and easier to estimate, especially when large datasets are available. By focusing exclusively on minimizing risk, the MVP avoids the amplification of noise that arises when small differences in estimated expected returns drive large differences in optimal weights.

This robustness helps explain why minimum variance portfolios often perform well in empirical applications. Several studies show that MVPs frequently outperform more complex optimization strategies out of sample, particularly in environments characterized by high estimation uncertainty. Their reliance on second moments alone reduces the instability of portfolio weights, lowers turnover, and mitigates the risk of overfitting historical data.

From a geometric perspective, the MVP anchors the efficient frontier. All efficient portfolios can be expressed as combinations of the minimum variance portfolio and any other efficient portfolio, reflecting the linear structure of the mean–variance set. The MVP therefore plays a central role in shaping the frontier and determining the range of feasible risk–return combinations.

In practical portfolio management, minimum variance strategies are widely used because of their simplicity, transparency, and empirical reliability. They are particularly attractive in settings where expected return estimates are highly unreliable, such as short sample periods, volatile markets, or portfolios with many assets relative to available data. For these reasons, the MVP is often employed as a benchmark or starting point for more sophisticated portfolio construction methods.

Overall, the minimum variance portfolio illustrates a fundamental insight of modern portfolio theory: effective risk reduction can be achieved even without precise knowledge of expected returns, provided that the covariance structure of asset returns is well understood. This insight becomes especially relevant in empirical applications, where estimation risk plays a central role and motivates the development of more robust portfolio strategies.

2.6 The Risk-Free Asset and the Tangency Portfolio

The introduction of a risk-free asset fundamentally changes the structure of the mean–variance framework. Without a risk-free asset, investors choose among portfolios that lie on the efficient

frontier of risky assets, which forms a concave curve in risk–return space. When a risk-free asset is available, however, the geometry of the efficient set changes: the curved frontier is replaced by a straight line known as the Capital Market Line (CML).

The CML represents all possible combinations of the risk-free asset and a single optimal risky portfolio. Since the risk-free asset has zero variance and zero covariance with all risky assets, combining it with any risky portfolio generates a linear relationship between expected return and risk. The slope of this line corresponds to the Sharpe ratio, which measures the excess return per unit of risk.

Let w denote the proportion invested in the risky portfolio and 1−w the proportion invested in the risk-free asset. The expected return of the combined portfolio is given by:

$E(R_p) = wE(R_risky) + (1 - w)R_f$

while its variance is:

$\sigma_p^2 = w^2 \sigma_risky^2$

Because the risk-free asset creates a linear combination of risk and return, investors can reach any point along the CML by adjusting the weight w. If $w>1$, the investor borrows at the risk-free rate to leverage the risky portfolio; if $w<0$, the investor lends at the risk-free rate.

Among all possible lines that can be drawn from the risk-free rate to the frontier of risky assets, the optimal one is the line that is tangent to the efficient frontier. The point of tangency identifies the tangency portfolio, the unique risky portfolio that maximizes the Sharpe ratio:

$\text{Sharpe Ratio} = (E(R_p) - R_f) / \sigma_p$

The tangency portfolio represents the best attainable risk–return trade-off in the market because it offers the highest excess return per unit of risk. Its composition depends solely on expected returns and covariances of risky assets and does not depend on individual investor preferences. This property reflects the separation theorem, which states that investment decisions can be divided into two independent stages: first, all investors hold the same optimal risky portfolio

(the tangency portfolio); second, investors choose how much to invest in the tangency portfolio versus the risk-free asset based on their individual risk tolerance.

Economically, this means that differences in investor preferences affect only the allocation between risky and risk-free assets, not the composition of the risky portfolio itself. Highly risk-averse investors hold a larger share of the risk-free asset, while more risk-tolerant investors allocate more wealth to the tangency portfolio or even leverage it.

The introduction of a risk-free asset therefore simplifies the portfolio selection problem and provides a clear benchmark for evaluating investment performance. The tangency portfolio becomes the central object of interest in the mean–variance framework, forming the theoretical basis for the Capital Asset Pricing Model (CAPM) and for many practical asset allocation strategies.

At the same time, the construction of the tangency portfolio relies heavily on accurate estimates of expected returns and covariances. This dependence makes it particularly sensitive to estimation error, a limitation that becomes crucial in empirical applications and motivates the discussion of practical challenges in the next section.

2.7 Limitations of Mean–Variance Optimization

Despite its strong theoretical foundation, mean–variance optimization faces several practical limitations that affect its real-world applicability. The model requires accurate estimates of expected returns, variances, and covariances, parameters typically derived from historical data and therefore subject to sampling error. Since financial markets are inherently noisy and dynamic, these estimates often deviate from their true population values, introducing uncertainty into the optimization process.

Among all inputs, expected returns are the most difficult to estimate reliably. They exhibit high volatility, low signal-to-noise ratios, and frequent structural breaks. Even small errors in

expected return estimates can lead to disproportionately large changes in optimal portfolio weights, because the optimization algorithm tends to allocate extreme positions to assets with slightly higher estimated returns. This sensitivity produces unstable and often unrealistic portfolios, characterized by concentrated positions, high turnover, and weak out-of-sample performance.

Estimation risk also affects the covariance matrix. Although covariances are generally more stable than expected returns, their estimation becomes more challenging as the number of assets increases. The number of covariances to estimate grows quadratically with the size of the portfolio, and when the sample size is small relative to the number of assets, the sample covariance matrix may become ill-conditioned or nearly singular. In such cases, the inversion of the covariance matrix, required for mean–variance optimization, tends to amplify noise rather than filter it out, resulting in extreme and erratic portfolio weights.

These issues highlight a tension between theoretical optimality and empirical robustness. In theory, mean–variance optimization identifies the best possible portfolio given the true parameters of the return distribution. In practice, however, the optimization process often overfits historical data, capturing noise rather than underlying structure. As a result, portfolios that appear optimal in sample frequently perform poorly out of sample, a finding consistently documented in empirical studies.

Another limitation arises from the model's strong reliance on historical data. Financial markets evolve over time, and historical relationships among asset returns may not persist in the future. Structural changes, technological innovations, regulatory shifts, and macroeconomic regime changes can all alter return dynamics, making past data a potentially unreliable guide for future expectations. This instability further undermines the reliability of sample-based optimization.

These challenges have motivated the development of alternative portfolio construction methods that emphasize robustness, stability, and reduced sensitivity to estimation error.

Approaches such as shrinkage estimators, Bayesian methods, factor-based models, and robust optimization techniques aim to improve the stability of portfolio weights by incorporating prior information, reducing dimensionality, or explicitly accounting for parameter uncertainty.

Overall, the limitations of mean–variance optimization underscore the gap between theoretical models and practical implementation. This gap has become a central theme in modern portfolio research and provides the motivation for examining estimation risk in greater detail. The next chapter explores this issue more deeply, analyzing how estimation uncertainty affects portfolio performance and why simple allocation rules often outperform theoretically optimal strategies in empirical settings.

CHAPTER 3 - ESTIMATION RISK AND IMPLEMENTATION CHALLENGES

3.1 The Nature of Estimation Risk

Although Modern Portfolio Theory provides a coherent framework for portfolio selection, its practical implementation is significantly complicated by estimation risk. Estimation risk arises because the parameters required for portfolio optimization, such as expected returns, variances, and covariances, are not known with certainty. Instead, they must be estimated from historical data, which provide only imperfect and noisy information about the true underlying return distribution.

In theory, portfolio optimization assumes that investors know the true probability distribution of asset returns. In practice, however, this distribution is unknown, and investors must rely on sample estimates derived from past observations. These estimates are subject to sampling variability, meaning that they may differ from the true parameters simply due to random fluctuations in observed returns. Estimation risk can therefore be defined as the uncertainty introduced by using estimated rather than true parameters in the optimization process.

Estimation risk can be further decomposed into two components: statistical estimation error and model uncertainty. Statistical estimation error arises from the limited size of historical samples, and the precision of parameter estimates depends directly on the amount of available data. Even with large samples, however, model uncertainty persists if the assumed statistical model does not accurately describe the true data-generating process. For example, if returns are not normally distributed or if correlations change over time, the assumptions underlying the mean–variance framework may not hold, regardless of sample size.

It is also important to distinguish estimation risk from parameter uncertainty. Estimation risk refers specifically to the error introduced by using sample estimates instead of true parameters. Parameter uncertainty reflects the broader difficulty of predicting future return distributions in

markets characterized by structural change, evolving risk premia, and unpredictable shocks. Even perfect estimation of past parameters does not guarantee accurate forecasts of future returns.

Among all parameters, expected returns are particularly difficult to estimate. Unlike variances and covariances, which tend to be relatively stable over time, expected returns are highly volatile and influenced by numerous economic factors, including macroeconomic conditions, investor sentiment, and changing risk premia. As a result, even small errors in expected return estimates can lead to large distortions in optimal portfolio weights, making estimation risk one of the most important practical limitations of mean–variance optimization.

This gap between theoretical assumptions and practical realities makes estimation risk one of the most studied topics in empirical asset allocation, and motivates the analysis developed in the following sections.

3.2 Sensitivity of Optimal Portfolio Weights

One of the most important consequences of estimation risk is the extreme sensitivity of optimal portfolio weights to the input parameters used in mean–variance optimization. Since portfolio weights are obtained by solving a system of equations that depends directly on expected returns and the covariance matrix, even small errors in these inputs can lead to large and economically significant distortions in the resulting allocations.

This sensitivity is particularly pronounced for expected returns, which enter the optimization problem linearly but influence portfolio weights nonlinearly through the inversion of the covariance matrix. As a result, even minor differences in estimated expected returns, differences that may not be statistically significant, can cause the optimization algorithm to assign disproportionately large weights to certain assets. This phenomenon is often described

as an amplification effect, because estimation errors in expected returns are magnified through the optimization process.

In empirical applications, this amplification effect leads to several characteristic patterns. Optimized portfolios frequently exhibit highly concentrated positions, with large positive or negative weights assigned to a small subset of assets. These extreme allocations are mathematically optimal given the estimated parameters, but they are often unrealistic, difficult to implement, and inconsistent with practical investment constraints such as leverage limits, short-selling restrictions, or regulatory requirements. Moreover, because parameter estimates change over time, optimized portfolios tend to experience large fluctuations in weights, resulting in high turnover and substantial transaction costs.

The instability of optimal weights reflects a deeper economic issue: optimization models treat estimated returns as if they were precise signals rather than uncertain forecasts. In reality, expected return estimates are noisy and subject to considerable uncertainty. When optimization algorithms interpret small differences in these estimates as meaningful, they produce allocations that are highly sensitive to estimation error and that may not be justified by underlying economic fundamentals.

This instability is further increased by the high dimensionality of the portfolio optimization problem. As the number of assets increases, the number of parameters to estimate grows rapidly, particularly within the covariance matrix. When the sample size is limited relative to the number of assets, the covariance matrix may become ill-conditioned or nearly singular, making its inversion unstable. This instability spreads directly into the portfolio weights, amplifying noise and reducing the reliability of the optimization results.

Empirical studies consistently document that optimized portfolios often perform poorly out of sample, precisely because their weights reflect patterns specific to the estimation period rather than persistent economic relationships. This poor performance is not simply a statistical artifact

but a structural consequence of the optimization process itself, which is highly sensitive to estimation error and prone to overfitting historical data.

From a practical perspective, the instability of optimal weights undermines the usefulness of classical mean–variance optimization in real-world portfolio management. Investors require portfolios that are not only theoretically optimal but also stable, implementable, and robust to uncertainty. The sensitivity of optimal weights to estimation error highlights the need for alternative approaches that reduce reliance on noisy inputs and produce more stable allocations. This issue naturally leads to the broader discussion of parameter uncertainty and forecasting limitations, which are examined in the following section.

3.3 Parameter Uncertainty and Forecasting Limitations

Estimation risk is closely connected to the broader issue of parameter uncertainty, which reflects the inherent difficulty of forecasting future asset return distributions. Historical data provide only an imperfect and incomplete representation of future market conditions, and the parameters estimated from past observations may not accurately describe the dynamics that will prevail in subsequent periods. This limitation is particularly severe in financial markets, where structural changes, evolving economic environments, and unexpected shocks frequently alter the behavior of asset returns.

Parameter uncertainty arises because financial markets are not governed by fixed and immutable statistical laws. Instead, they evolve over time in response to macroeconomic developments, technological innovations, regulatory changes, geopolitical events, and shifts in investor sentiment. As a result, the statistical properties of asset returns, such as means, variances, and correlations, may vary across different periods. When such changes occur, historical estimates may no longer provide reliable guidance for future expectations, even if they were accurate in the past.

This challenge is especially noticeable for expected returns, which depend on dynamic and forward-looking factors such as economic growth prospects, interest rate expectations, risk premia, and market sentiment. These variables are inherently difficult to forecast and often subject to abrupt changes. Consequently, expected return estimates tend to be noisy and unstable, making them a particularly unreliable input for portfolio optimization.

Moreover, the distribution of asset returns itself may change over time. Financial markets frequently exhibit model instability, meaning that the statistical model assumed by the investor, such as normality, stationarity, or constant correlations, may not hold consistently across different periods. When the underlying data-generating process shifts, even unbiased historical estimates become inaccurate predictors of future outcomes. This instability limits the reliability of optimization models that rely heavily on historical data.

The implications for portfolio construction are significant. Portfolios optimized using historical parameter estimates may perform poorly out of sample because the estimated relationships among asset returns do not persist into the future. This forecasting limitation is not simply a technical inconvenience but a fundamental challenge for empirical portfolio optimization. It reflects the reality that financial markets are complex, adaptive systems in which past patterns may not repeat.

Even sophisticated statistical models cannot fully eliminate parameter uncertainty. Techniques such as Bayesian updating, shrinkage estimation, or factor-based modeling can reduce estimation error, but they cannot guarantee accurate forecasts when the underlying economic environment changes. For this reason, parameter uncertainty should be viewed not only as a statistical issue but as an intrinsic feature of financial markets characterized by structural change and unpredictable dynamics.

These limitations motivate the use of dynamic estimation techniques, such as rolling-window procedures, which attempt to adapt parameter estimates to evolving market conditions. By

updating estimates over time, these methods aim to capture changes in return dynamics more effectively than static models. The next section examines rolling-window estimation in detail and discusses its role in mitigating, though not eliminating, the challenges posed by parameter uncertainty.

3.4 Rolling-Window Estimation

Rolling-window estimation is one of the most widely used techniques in empirical portfolio analysis, designed to address the challenges posed by parameter uncertainty and the evolving nature of financial markets. The central idea is to mimic the way real-world investors update their beliefs over time: rather than relying on a fixed historical sample, parameter estimates are recalculated periodically using only the most recent observations. This approach allows the model to adapt gradually to changing market conditions while avoiding the unrealistic assumption that investors have access to future information.

In a rolling-window framework, parameters such as expected returns, variances, and covariances are estimated using a fixed-length historical sample, known as the estimation window. These estimates are then used to construct portfolio weights for the subsequent investment period. After each period, the window is shifted forward by one observation: the oldest data point is discarded, and the newest one is added. This process generates a sequence of time-varying parameter estimates that evolve in response to new information.

Rolling-window estimation offers several important advantages. First, it ensures that portfolio decisions are based solely on information available at the time of investment, thereby eliminating look-ahead bias and providing a realistic representation of how investors operate in practice. Second, it allows researchers to evaluate portfolio performance in a dynamic setting, capturing the effects of parameter instability and market evolution over time. Third, rolling-window methods are particularly well suited to financial data, which often exhibit

time-varying volatility, changing correlation structures, and regime shifts that cannot be captured by static models.

From a statistical perspective, rolling-window estimation helps reduce the risk of relying on outdated information. When market conditions change, due to macroeconomic shocks, policy shifts, or structural breaks, older observations may no longer be informative about future return dynamics. By discarding older data, rolling-window methods reduce the influence of obsolete information and allow parameter estimates to adjust more quickly to new environments.

However, rolling-window estimation also has limitations. Because the estimation window is fixed in length, the effective sample size remains constant over time. When the window is short relative to the number of assets, parameter estimates may be noisy and imprecise, particularly for the covariance matrix, which requires estimating a large number of parameters. Conversely, if the window is too long, the method may fail to adapt quickly enough to structural changes, causing estimates to lag behind evolving market conditions. This trade-off between responsiveness and statistical precision is a central challenge in the design of rolling-window procedures.

Moreover, while rolling-window estimation improves realism, it does not eliminate estimation risk. Parameter uncertainty remains an inherent feature of financial data, and even dynamically updated estimates may deviate substantially from true future values. As a result, portfolios constructed using rolling-window estimates may still exhibit instability, high turnover, and poor out-of-sample performance, especially when expected returns play a dominant role in the optimization process.

Despite these limitations, rolling-window estimation remains a cornerstone of empirical portfolio research. It provides a practical and conceptually intuitive framework for evaluating portfolio strategies under realistic information constraints and forms the basis for out-of-sample performance evaluation, which is examined in the next section.

3.5 Out-of-Sample Performance Evaluation

A fundamental component of empirical portfolio analysis is the distinction between in-sample and out-of-sample performance. This distinction is essential for determining whether a portfolio strategy genuinely possesses predictive power or merely reflects patterns specific to the historical data used for estimation. Because portfolio optimization is inherently backward-looking, evaluating performance solely in-sample can lead to misleading conclusions about the true effectiveness of a strategy.

In-sample performance refers to the evaluation of a portfolio using the same dataset employed to estimate expected returns and covariances. Since optimization procedures are designed to fit the available data as closely as possible, in-sample results often appear highly favorable. However, these results may simply reflect the model's ability to capture noise rather than meaningful economic relationships. In this sense, strong in-sample performance is not necessarily indicative of genuine investment skill but may instead be a symptom of overfitting.

Out-of-sample performance, by contrast, evaluates how a portfolio performs when applied to new data not used during the estimation process. This approach simulates the real-world decision environment faced by investors, who must make allocation decisions based on past information without knowing future outcomes. Out-of-sample evaluation therefore provides a more realistic and stringent test of a portfolio strategy's robustness and predictive accuracy.

Estimation risk becomes particularly evident in out-of-sample evaluations. Portfolios that appear optimal in-sample frequently fail to achieve comparable performance when tested on new data. This discrepancy arises because the input parameters used during optimization, especially expected returns, contain estimation errors that become apparent only when the portfolio is applied to future periods. When the optimization process overfits historical noise, the resulting portfolio weights may be poorly aligned with the true underlying return structure, leading to disappointing out-of-sample results.

A closely related issue is the high dimensionality of the portfolio optimization problem. The covariance matrix alone contains $N(N-1)/2$ distinct elements, and when the number of assets is large relative to the sample size, reliable estimation becomes extremely challenging. In such settings, the optimization algorithm may inadvertently exploit random correlations present in the estimation sample, producing portfolios that perform well in-sample but poorly out-of-sample. This problem intensifies when the inversion of the covariance matrix amplifies estimation noise, further increasing the likelihood of overfitting.

Out-of-sample evaluation is therefore crucial not only from a statistical perspective but also from an economic one. Investors are ultimately concerned with future performance, not with how well a portfolio fits historical data. For this reason, out-of-sample testing has become a methodological standard in empirical finance and is widely used to assess the practical viability of portfolio strategies. It provides a disciplined framework for evaluating whether a strategy can generate consistent performance across different market environments and whether it remains robust in the presence of parameter uncertainty.

Moreover, out-of-sample evaluation highlights the limitations of classical mean–variance optimization and underscores the need for more robust portfolio construction methods. When optimized portfolios consistently underperform simple benchmarks such as the naïve 1/N strategy, the discrepancy reveals the extent to which estimation risk undermines theoretical optimality.

The challenges documented through out-of-sample analysis motivate the exploration of alternative approaches that reduce reliance on noisy parameter estimates. These issues naturally lead to the discussion of long-horizon risk and predictive uncertainty, examined in the following section.

3.6 Predictive Variance and Long-Horizon Risk

Another important dimension of estimation risk concerns the uncertainty associated with long-horizon risk forecasts. While short-term portfolio risk depends primarily on current volatility and correlations, long-term risk incorporates an additional layer of uncertainty arising from the fact that parameter estimates themselves are imprecise. As the investment horizon increases, this uncertainty accumulates, making long-horizon risk forecasts increasingly unreliable.

In the mean–variance framework, long-horizon risk is often approximated by scaling short-term variance by the length of the horizon. However, this approach implicitly assumes that expected returns and covariances are known and constant over time. In reality, both quantities are estimated with error and may evolve in response to changing economic conditions. As a result, long-horizon risk is influenced not only by the volatility of returns but also by the predictive variance associated with parameter uncertainty.

Predictive variance analysis shows that uncertainty about future expected returns becomes a dominant component of long-term portfolio risk. Even small estimation errors in expected returns can accumulate over multiple periods, leading to large deviations in projected wealth outcomes. This compounding effect reflects the fact that expected returns influence the drift of the return distribution, and errors in estimating this drift grow linearly with the investment horizon. Consequently, long-horizon forecasts are particularly sensitive to parameter uncertainty, especially when expected returns play a central role in the allocation process.

From a statistical perspective, long-horizon risk forecasts must account for both the accumulation of short-term return volatility over time and the additional uncertainty that arises as estimation errors carry over across multiple periods.

The second component is often larger and more consequential for long-term investors. When expected returns are estimated imprecisely, as is typically the case, long-horizon projections

become dominated by uncertainty about the true mean of the return distribution. This effect is especially pronounced in environments characterized by structural breaks, regime shifts, or time-varying risk premia, where historical estimates may provide a poor guide to future dynamics.

Economically, this implies that long-term investment decisions are inherently more difficult to optimize accurately. Even minor inaccuracies in parameter estimates can lead to substantial deviations in expected wealth paths, making long-horizon portfolio optimization highly sensitive to model assumptions. This sensitivity helps explain why long-term strategic asset allocation is often based on broad diversification principles rather than precise optimization of expected returns.

The challenge of forecasting long-horizon risk also highlights the limitations of classical mean–variance optimization in dynamic settings. When parameter uncertainty is large, the theoretical benefits of long-term optimization may be overshadowed by the instability of the underlying estimates. This insight reinforces the need for robust portfolio strategies that remain stable under uncertainty and do not rely excessively on precise long-term forecasts.

These considerations naturally lead to the broader question of how estimation risk affects portfolio optimization in practice, a topic examined in detail in the following section.

3.7 Consequences of Estimation Risk for Portfolio Optimization

Estimation risk has significant consequences for practical portfolio management, affecting not only the stability of portfolio weights but also the economic viability of optimization-based strategies. Because mean–variance optimization relies heavily on estimated parameters, any inaccuracies in expected returns or covariances translate directly into portfolio allocations, often producing results that are theoretically optimal but empirically fragile.

A first and well-documented consequence is the systematic overweighting of assets with high estimated expected returns, even when these estimates are highly uncertain. The optimization algorithm interprets small differences in estimated returns as meaningful signals, assigning disproportionately large weights to assets that appear marginally more attractive. This behavior is problematic because expected returns are extremely noisy, and the resulting allocations may reflect estimation error rather than genuine investment opportunities.

A second consequence is the instability of portfolio weights over time. As new data become available and parameter estimates are updated, optimized portfolios often experience large and frequent shifts in asset weights. This instability leads to high portfolio turnover, which in turn generates substantial transaction costs. These costs can significantly erode realized performance, especially in markets with wide bid–ask spreads or for strategies that require frequent rebalancing. In many empirical studies, transaction costs alone are sufficient to eliminate any apparent advantage of optimized portfolios over simpler allocation rules.

A third consequence concerns the shape and position of the efficient frontier. Because the efficient frontier is constructed using estimated parameters, it is itself subject to estimation error. When expected returns or covariances are misestimated, the efficient frontier may shift, rotate, or become distorted. As a result, the frontier observed in-sample may show little resemblance to the true frontier, and portfolios that appear efficient ex ante may turn out to be inefficient ex post. This instability undermines the practical relevance of the efficient frontier as a decision-making tool.

Estimation risk also affects the risk characteristics of optimized portfolios. Portfolios that appear to have low variance in sample may exhibit much higher volatility out of sample, because the optimization process has inadvertently exploited random patterns in historical data. Similarly, portfolios that appear well-diversified in sample may turn out to be highly exposed

to hidden sources of systematic risk. These discrepancies highlight the gap between theoretical optimality and empirical robustness.

From an economic perspective, these consequences reflect the fact that classical optimization models treat estimated parameters as if they were known with certainty. In reality, parameter estimates are subject to substantial uncertainty, and optimization procedures that ignore this uncertainty tend to produce allocations that are overly aggressive and insufficiently diversified. This mismatch between theoretical assumptions and empirical realities explains why optimized portfolios often underperform simple benchmarks such as the naïve 1/N strategy.

Overall, estimation risk affects portfolio weights, turnover, risk exposure, and the shape of the efficient frontier, consequences that are concrete and measurable rather than merely theoretical. These challenges motivate the search for robust portfolio strategies that remain stable under parameter uncertainty and reduce reliance on noisy inputs. The next section examines these approaches in detail.

3.8 The Need for Robust Portfolio Strategies

The limitations of classical mean–variance optimization have motivated the development of robust portfolio strategies designed to mitigate the impact of estimation risk and parameter uncertainty. Because optimized portfolios are highly sensitive to small errors in expected returns and covariances, researchers have increasingly explored methods that produce more stable, reliable, and economically meaningful allocations. The central objective of these approaches is not to eliminate estimation risk, an impossible task in financial markets characterized by uncertainty, but to reduce the extent to which portfolio performance depends on noisy parameter estimates.

Robust portfolio strategies share a common principle: they limit the influence of unreliable inputs by imposing structure, incorporating prior information, or simplifying the optimization

problem. One of the most straightforward examples is naïve diversification, which assigns equal weights to all assets. As discussed in the previous chapters, this strategy avoids estimation risk entirely and has been shown to perform competitively out of sample precisely because it does not depend on noisy parameter estimates.

Another important class of robust methods involves constrained optimization, where restrictions are imposed on portfolio weights to prevent extreme allocations. Constraints such as no-short-selling, maximum position limits, or bounds on leverage reduce the sensitivity of portfolio weights to estimation noise and help ensure that allocations remain diversified and economically realistic. These constraints act as a form of regularization, preventing the optimization algorithm from overreacting to small differences in estimated returns.

A further approach is the use of shrinkage estimators, which stabilize covariance matrix estimates by combining sample estimates with structured targets, such as the identity matrix or a constant-correlation matrix. Shrinkage reduces estimation error by pulling noisy sample estimates toward more stable and simple structures. This technique has been shown to improve out-of-sample performance, particularly in high-dimensional settings where the number of assets is large relative to the sample size.

Bayesian methods offer another approach for robustness by incorporating prior information about expected returns or covariances. Rather than relying solely on historical data, Bayesian approaches combine sample information with prior distributions, producing estimates that are less sensitive to sampling variability. This framework allows investors to explicitly account for parameter uncertainty and to adjust their estimates based on the strength of available evidence.

More recently, robust optimization techniques have been developed to address uncertainty directly within the optimization problem. These methods treat parameter estimates as uncertain and seek portfolios that perform well across a range of plausible scenarios rather than optimizing for a single estimated set of parameters. By explicitly modeling uncertainty, robust

optimization produces allocations that are less sensitive to estimation error and more stable across different market environments.

Taken together, these approaches reflect a broader shift in modern portfolio research: from seeking theoretical optimality to developing strategies that deliver reliable performance under real-world uncertainty. The next chapter examines them in detail, evaluating their effectiveness in mitigating estimation risk.

CHAPTER 4 - PORTFOLIO OPTIMIZATION MODELS

4.1 Classical Mean–Variance Optimization

Classical mean–variance optimization represents the foundational framework of modern portfolio theory and provides what can be considered the first rigorous analytical solution to the portfolio selection problem. Developed by Markowitz (1952), this approach formalizes the idea that investors should evaluate assets not in isolation but in terms of their joint contribution to overall portfolio risk and return. Within this framework, investors choose portfolio weights by balancing expected return against risk, measured by the variance of portfolio returns.

At the core of mean–variance optimization lies a quadratic optimization problem in which investors either minimize portfolio variance subject to achieving a target expected return or, equivalently, maximize expected return subject to a risk constraint. The optimization requires estimates of two key inputs: the vector of expected returns and the covariance matrix of asset returns. Using these inputs, the model identifies the set of portfolios that lie on the efficient frontier, representing the best achievable trade-offs between risk and return.

Portfolio variance depends not only on the volatility of individual assets but also on the correlations between them. This interdependence is what makes portfolio optimization fundamentally different from the selection of individual securities. The contribution of each asset to overall portfolio risk depends on how its returns move together with those of other assets. As a result, diversification benefits arise whenever asset returns are imperfectly correlated, allowing investors to reduce risk without sacrificing expected return.

From an economic perspective, mean–variance optimization provides a clear and intuitive representation of rational investment behavior under uncertainty. It formalizes the principle that investors should consider the marginal contribution of each asset to portfolio risk and

return, rather than evaluating assets solely on their standalone characteristics. This insight marked a profound shift in financial thinking, establishing diversification as a central principle of investment management and laying the groundwork for subsequent developments in asset pricing theory.

Despite its simplicity, the classical mean–variance framework relies heavily on the accuracy of input estimates and is therefore highly exposed to estimation risk. The sensitivity of portfolio weights to noisy expected return estimates often results in concentrated, high-turnover allocations that perform poorly out of sample.

Despite these limitations, classical mean–variance optimization remains the fundamental benchmark in empirical finance and the conceptual foundation for all the alternative methods discussed in this chapter.

From a broader perspective, mean–variance optimization represents not only a mathematical solution to the portfolio selection problem but also a conceptual framework that changed how investors think about risk. By emphasizing that risk is a portfolio-level phenomenon driven by interactions among assets, the model introduced a new way of understanding investment decisions, one that continues to provide the foundation for modern asset allocation theory and investment practice.

4.2 Minimum Variance Portfolio

The Minimum Variance Portfolio (MVP) represents one of the most important and practically relevant special cases of the mean–variance framework. Unlike full mean–variance optimization, which requires estimates of both expected returns and covariances, the MVP focuses exclusively on minimizing portfolio risk. This simplification eliminates the need for expected return estimates, parameters that are notoriously difficult to estimate accurately, and relies solely on the covariance structure of asset returns.

Mathematically, the MVP is obtained by solving the following optimization problem:

$\min_{(w)} w^T \Sigma w$

subject to:

$w^T 1 = 1$

This formulation highlights a defining feature of the MVP: its construction depends only on second moments. Because variances and covariances tend to be more stable over time than expected returns, the MVP is significantly less sensitive to estimation error. This robustness makes it particularly attractive in empirical applications, where parameter uncertainty is a central challenge.

From a theoretical perspective, the MVP occupies a special position on the efficient frontier. It corresponds to the leftmost point of the frontier, the portfolio with the lowest achievable variance among all feasible allocations. All other efficient portfolios can be expressed as combinations of the MVP and any other efficient portfolio, reflecting the linear structure of the mean–variance set. In this sense, the MVP acts as the anchor of the efficient frontier and provides a natural benchmark for evaluating risk–return trade-offs.

One of the most important advantages of the MVP is its tendency to produce more diversified and stable portfolio weights. Because the optimization is not driven by noisy expected return estimates, the resulting allocations are typically more broadly distributed across assets. This diversification effect reduces concentration risk and leads to portfolios that are more robust to sampling variability.

Empirical evidence, most notably DeMiguel, Garlappi, and Uppal (2009), shows that MVPs often exhibit lower turnover and more stable weight dynamics over time, which in turn reduces transaction costs and improves implementability.

The MVP also aligns closely with economic intuition. By focusing solely on risk minimization, it reflects a cautious approach focused on that prioritizes stability and diversification. In

environments characterized by high parameter uncertainty, volatile markets, or limited historical data, such an approach can be particularly advantageous. The MVP also provides a useful starting point for more advanced portfolio construction techniques.

In practical terms, the MVP is relatively straightforward to compute, but it still requires accurate estimation of the covariance matrix. When the number of assets is large relative to the sample size, the sample covariance matrix may be unstable or ill-conditioned, leading to unreliable weight estimates. This challenge has motivated the development of improved covariance estimators, such as shrinkage methods, constrained optimization, and robust portfolio approaches, that enhance the stability of the MVP in high-dimensional settings.

Overall, the minimum variance portfolio illustrates a central theme in modern portfolio research: simpler models that rely on more stable inputs often outperform theoretically optimal strategies when implemented in real-world conditions. By reducing dependence on noisy expected return estimates and emphasizing diversification, the MVP provides a robust and empirically effective approach to portfolio construction.

4.3 Portfolio Constraints and Practical Implementation

In practice, portfolio optimization is subject to various constraints that reflect real-world investment conditions. These include limits on short selling, restrictions on individual asset weights, and regulatory requirements that govern leverage and diversification. Without such constraints, mean–variance optimization often produces unrealistic portfolios characterized by extreme positive and negative positions. Although these solutions may be mathematically optimal given the estimated parameters, they are typically not feasible in actual investment settings and can lead to allocations that are economically unrealistic or difficult to implement.

Imposing constraints can significantly improve the stability and implementability of optimized portfolios. For example, prohibiting short sales prevents the optimizer from assigning large

negative weights to assets with low estimated expected returns, while upper bounds on individual asset weights reduce concentration risk and promote more balanced allocations. These restrictions help ensure that portfolio solutions remain consistent with institutional policies, regulatory frameworks, and the practical realities of investment management.

Common constraints include no-short-selling constraints (weights ≥ 0), budget constraints (weights sum to one), upper and lower bounds on individual asset weights, and turnover constraints designed to limit trading costs.

From a statistical perspective, constraints act as a form of regularization. By restricting the feasible set of portfolio weights, they reduce the optimizer's sensitivity to estimation errors in expected returns and covariances. This is particularly important because mean–variance optimization tends to overreact to small variations in estimated parameters, often producing unstable and highly leveraged allocations. Constraints reduce this instability, leading to more robust and reliable portfolios, especially in out-of-sample settings.

Economically, constraints reflect the way institutional investors operate in practice. Pension funds, insurance companies, and mutual funds must comply with regulatory rules that limit leverage, impose diversification requirements, or restrict certain types of positions. Even in the absence of formal regulations, many investors adopt internal guidelines to control concentration, manage liquidity risk, and maintain transparency. Constraints therefore serve not only as technical devices but also as essential components of prudent portfolio management.

Another important benefit of constraints is their ability to reduce portfolio turnover. Unconstrained optimized portfolios often exhibit large and frequent changes in weights as new data become available, generating substantial transaction costs. By limiting extreme reallocations or imposing explicit turnover constraints, investors can achieve more stable

weight dynamics and improve net performance. This stability is particularly valuable in markets with significant trading frictions or for strategies that require regular rebalancing.

From a mathematical perspective, adding constraints transforms the optimization problem from an unconstrained quadratic problem to a constrained optimization problem, typically solved using numerical algorithms such as quadratic programming with inequality constraints. Although this increases computational complexity, modern numerical algorithms make it feasible to solve such problems even for large portfolios. The resulting allocations tend to be more diversified, more stable, and more aligned with practical investment requirements.

Overall, portfolio constraints serve a dual purpose: they ensure that optimized portfolios are feasible and economically meaningful, and they enhance robustness by reducing sensitivity to estimation error. In this sense, constraints help bridge the gap between the theoretical structure of the mean–variance framework and the practical realities of financial markets, making optimization-based strategies more reliable and implementable in real-world settings.

4.4 Shrinkage Estimators and Bayesian Approaches

A central challenge in portfolio optimization is the accurate estimation of input parameters, particularly expected returns and the covariance matrix. Classical sample estimates are highly sensitive to sampling variability, especially when the number of assets is large relative to the available data. This sensitivity amplifies estimation risk and often leads to unstable portfolio weights and poor out-of-sample performance. To address these limitations, modern portfolio theory has developed a range of techniques that introduce structure, prior information, or regularization into the estimation process.

Among the most influential are shrinkage estimators and Bayesian approaches, both aimed at improving the reliability of parameter estimates by balancing bias and variance.

Shrinkage estimators are built on the idea that combining noisy sample estimates with more stable, structured targets can produce more accurate results. Instead of relying solely on the sample covariance matrix, which becomes unstable or nearly singular in high-dimensional settings, shrinkage methods pull the sample estimate toward a target such as the identity matrix, a constant-correlation matrix, or a factor-based structure. This adjustment reduces estimation variance at the cost of introducing a small amount of bias, a trade-off that often leads to superior overall accuracy in finite samples. Shrinkage techniques are widely used in empirical finance because they stabilize the covariance matrix, improve its invertibility, and generate more diversified and robust portfolio weights.

The benefits of shrinkage are particularly evident in minimum variance and risk-based strategies, where the covariance matrix plays a central role. By reducing noise in the covariance estimates, shrinkage leads to portfolios that are less sensitive to small fluctuations in the data and that exhibit more stable weight dynamics over time. This stability translates into lower turnover, reduced transaction costs, and better out-of-sample performance, features that are especially valuable in practical portfolio management.

Bayesian approaches extend the logic of shrinkage by treating parameters as random variables rather than fixed unknown quantities. Instead of relying exclusively on historical data, Bayesian methods incorporate prior information about expected returns or covariances and update these beliefs as new information becomes available. This framework explicitly models parameter uncertainty through probability distributions, allowing investors to combine sample information with prior knowledge in a coherent and internally consistent way.

One of the most influential Bayesian models in portfolio management is the Black–Litterman framework, which blends equilibrium returns derived from asset pricing theory with investor-specific views. By combining these two sources of information, the model produces expected return estimates that are more stable and less sensitive to sampling noise than

traditional sample means. Bayesian methods can also be applied to covariance estimation, where priors impose structure on the correlation matrix or on factor loadings, further reducing estimation error and improving robustness.

Both shrinkage and Bayesian approaches share a common objective: improving the stability and reliability of parameter estimates by reducing reliance on noisy historical data. Although these methods introduce some bias, they significantly reduce estimation variance, leading to more robust portfolio allocations and better out-of-sample performance. This shift from purely data-driven estimation to approaches that incorporate structure or prior information reflects a broader evolution in portfolio theory, emphasizing robustness, stability, and practical implementability.

These techniques also connect to robust optimization methods, which incorporate parameter uncertainty directly into the optimization problem. The next section examines these approaches in detail and discusses how they further enhance portfolio resilience in the presence of estimation risk.

4.5 Robust Optimization Methods

Robust optimization methods represent a major development in modern portfolio theory, arising from the recognition that classical mean–variance optimization is highly sensitive to estimation error. Instead of treating estimated parameters as if they were known with certainty, robust approaches explicitly acknowledge that expected returns and covariances are uncertain and may deviate from their sample estimates. By incorporating this uncertainty directly into the optimization problem, robust methods aim to construct portfolios that remain effective across a range of plausible parameter values, rather than being optimal only under a single, potentially inaccurate estimate.

A key feature of robust optimization is the introduction of uncertainty sets for expected returns, covariances, or both. These sets specify the range of parameter values that the investor considers plausible, given sampling variability, model uncertainty, or structural changes in the data-generating process. Instead of optimizing with respect to a single point estimate, the investor seeks portfolio weights that perform well under the worst-case realization within the uncertainty set. This leads to allocations that are more conservative, more stable, and less sensitive to estimation noise.

The economic intuition behind robust optimization is straightforward: when parameters are uncertain, it is rational to avoid portfolios that rely heavily on precise estimates of expected returns or covariances. Classical mean–variance optimization often produces extreme allocations because it interprets small differences in estimated returns as meaningful signals. Robust optimization tempers this behavior by penalizing portfolios that are overly sensitive to parameter changes, thereby reducing exposure to estimation error. As a result, robust portfolios tend to be more diversified and less concentrated than their classical counterparts.

From a statistical perspective, robust optimization can be viewed as a form of regularization applied at the optimization stage rather than at the estimation stage. While shrinkage and Bayesian methods stabilize the input parameters, robust optimization stabilizes the resulting portfolio by ensuring that weights remain reasonable even when the underlying estimates vary. This distinction highlights the complementarity between the approaches discussed in the previous section: shrinkage improves the quality of the inputs, while robust optimization improves the resilience of the outputs.

Robust optimization also reflects a broader philosophical shift in financial decision-making. Instead of assuming that investors possess precise knowledge about future return distributions, these methods acknowledge the reality of imperfect information and emphasize resilience over strict optimality. In uncertain environments, a portfolio that performs reasonably well across

many scenarios may be more valuable than one that appears optimal under highly specific assumptions. This perspective aligns closely with the way real-world investors operate, as portfolio managers focus on avoiding large losses caused by model misspecification or estimation errors rather than on achieving theoretical efficiency.

One of the most important consequences of robust optimization is the improvement in out-of-sample performance. Because robust portfolios are less sensitive to noise in the data, they tend to exhibit more stable weight dynamics, lower turnover, and reduced exposure to extreme positions. These features translate into more reliable performance across different market environments, particularly when compared to classical mean–variance portfolios, which often perform poorly out of sample due to their reliance on noisy expected return estimates.

From a mathematical standpoint, robust optimization transforms the classical quadratic program into a min–max problem, in which the investor minimizes risk or maximizes utility under the worst-case parameter realization. Although this increases computational complexity, modern optimization algorithms make robust methods feasible even for large portfolios. Moreover, the structure of the uncertainty sets can be defined to reflect the investor's beliefs, risk tolerance, or regulatory constraints, providing a flexible and customizable framework for portfolio construction.

Overall, robust optimization represents a significant step toward reducing the gap between theoretical portfolio models and practical investment decision-making. By explicitly accounting for parameter uncertainty, robust methods produce allocations that are more stable, more diversified, and more resilient to estimation error. In this sense, they complement the shrinkage and Bayesian techniques discussed earlier and form an essential component of the broader movement toward robust portfolio strategies.

4.6 Naïve Diversification Strategies

Naïve diversification strategies represent an alternative approach that avoids estimation risk entirely. The most common naïve strategy is the equal-weighting (1/N) rule, which allocates identical weights to all assets in the portfolio. Unlike mean–variance optimization or its robust variants, naïve diversification requires no estimation of expected returns or covariances, making it completely immune to estimation risk. This feature is particularly valuable in empirical settings where parameter uncertainty is substantial and historical data provide only limited guidance about future return dynamics.

The appeal of naïve diversification lies in its robustness. By allocating equal weights to all assets, the strategy implicitly enforces broad diversification and avoids the extreme positions that often arise from optimization procedures. Its simplicity prevents the portfolio from overreacting to noisy or unstable parameter estimates and ensures that no single asset dominates overall risk exposure. As a result, equal-weight portfolios tend to exhibit stable weight dynamics, low sensitivity to sampling variability, and relatively low turnover compared to unconstrained optimized portfolios.

Empirical research has consistently confirmed this intuition, showing that naïve diversification often outperforms more sophisticated strategies out of sample, a finding discussed in detail in Chapter 1 and examined empirically in Chapter 5. When estimation error is large, the benefits of optimization may be outweighed by the noise in the input parameters, and simple allocation rules that avoid estimation altogether can deliver more reliable performance.

From a statistical perspective, equal weighting can be interpreted as an implicit form of regularization. By imposing identical weights, the strategy effectively shrinks the portfolio toward the center of the feasible set, avoiding extreme allocations and reducing exposure to estimation noise. This regularization effect is particularly valuable in high-dimensional

settings, where the number of assets is large relative to the available sample size and traditional optimization becomes unstable.

Economically, naïve diversification is consistent with the idea that investors should not rely too heavily on uncertain forecasts. When expected returns are difficult to estimate and covariances are unstable, equal weighting provides a conservative and transparent approach that does not rely on precise parameter estimates. This makes the strategy especially attractive for long-term investors, for whom stability, diversification, and implementability are often more important than theoretical optimality.

Despite its simplicity, naïve diversification is not without limitations. Equal weighting ignores differences in risk, liquidity, and economic fundamentals across assets, potentially leading to suboptimal allocations when reliable information is available. Moreover, the strategy may overweight small or volatile assets relative to their contribution to market capitalization or economic value. Nevertheless, in environments characterized by high uncertainty and limited predictive power, these weaknesses are often outweighed by the strategy's robustness and empirical reliability.

Overall, naïve diversification serves as the natural benchmark against which all more complex strategies are evaluated.

4.7 Combining Optimization and Robustness

Recent research has increasingly explored hybrid approaches that combine the strengths of optimization-based methods with the stability of robust or naïve strategies. These approaches aim to integrate information from statistical models while limiting sensitivity to estimation error, acknowledging that no single method can fully address the challenges posed by parameter uncertainty, sampling variability, and the dynamic nature of financial markets. By

blending different techniques, hybrid strategies seek to retain the theoretical appeal of optimization while achieving the empirical reliability of more conservative allocation rules.

A central motivation for these hybrid methods is the complementary nature of the techniques discussed in previous sections. Classical mean–variance optimization provides a rigorous framework for balancing risk and return, but its performance is often undermined by the instability of estimated inputs. Robust optimization incorporates parameter uncertainty directly into the optimization problem, while shrinkage and Bayesian techniques improve the quality of the estimates themselves. Naïve diversification, in turn, offers a baseline of stability and broad diversification that is difficult to match when estimation error is high. Combining these elements allows investors to exploit the strengths of each approach while mitigating their respective weaknesses.

One common hybrid strategy involves blending optimized portfolios with simple allocation rules such as the equal-weight (1/N) portfolio. For example, constructing a convex combination of the minimum variance portfolio and the equal-weight portfolio can reduce sensitivity to estimation noise while preserving diversification benefits. This type of blending acts as an implicit form of regularization, shrinking optimized weights toward a stable benchmark and preventing extreme allocations driven by noisy estimates.

Other hybrid approaches integrate shrinkage or Bayesian estimation directly into the optimization process. By stabilizing the covariance matrix or expected return estimates before optimization, these methods reduce the likelihood that the optimizer will overreact to sampling variability. The resulting portfolios tend to be more diversified, more stable over time, and less prone to large shifts in response to small changes in the data. Multi-stage procedures follow a similar logic: investors first apply shrinkage or robust estimation to obtain stabilized parameters and then solve a constrained optimization problem that incorporates practical considerations such as turnover limits or weight bounds. This layered structure reflects the

reality of institutional portfolio management, where estimation, optimization, and implementation are distinct but interconnected stages of the decision-making process.

From an empirical perspective, hybrid methods often outperform both pure optimization and purely naïve strategies. They achieve this by balancing theoretical efficiency with practical robustness, reducing exposure to estimation error while still exploiting meaningful structure in the data. This balance is particularly valuable in high-dimensional settings, where the number of assets is large relative to the available sample size and traditional optimization becomes unstable. These approaches reflect a broader trend in modern portfolio research toward achieving stable, reliable, and implementable investment outcomes under conditions of uncertainty.

Taken together, the methods discussed in this chapter span the full spectrum from strict optimization to naïve allocation, each occupying a different point on the trade-off between theoretical efficiency and practical robustness. Among the empirical studies that have examined this issue in depth, the work of DeMiguel, Garlappi, and Uppal (2009) stands out as one of the most influential contributions, and forms the basis of the following chapter.

CHAPTER 5 - THE DEMIGUEL, GARLAPPI, UPPAL (2009) STUDY

5.1 Research Motivation

The study by DeMiguel, Garlappi, and Uppal examines one of the central questions in portfolio management: whether theoretically optimal portfolio strategies actually deliver superior performance in real-world investment settings. Classical portfolio theory, grounded in the mean–variance framework, suggests that investors should construct portfolios using estimates of expected returns, variances, and covariances in order to achieve efficient trade-offs between risk and return. In theory, optimization allows investors to fully exploit the benefits of diversification while allocating capital toward assets with higher expected returns.

In practice, however, implementing these models requires estimating parameters that are not directly observable. Expected returns, variances, and covariances must be inferred from historical data, which introduces unavoidable statistical uncertainty. Because historical samples are finite and subject to randomness, estimated parameters may differ substantially from their true underlying values. The authors argue that estimation risk can be so severe that it offsets the theoretical benefits of optimization. In particular, errors in expected return estimates often have a disproportionate impact on portfolio weights, leading to unstable allocations and poor out-of-sample performance.

The study therefore evaluates whether optimized portfolios truly outperform simple equal-weighted strategies under realistic investment conditions. This issue is not merely technical: it reflects a broader debate within financial economics about the practical relevance of theoretical models. While modern portfolio theory provides clear normative prescriptions, empirical evidence frequently shows that real-world performance diverges from theoretical predictions.

As discussed in the previous chapters, this tension between theoretical optimality and empirical performance arises because the accuracy of model inputs matters more than the sophistication of the optimization procedure in noisy, real-world settings. The study thus provides a direct empirical test of how classical portfolio theory performs once estimation risk, parameter instability, and changing market conditions are taken into account.

In this sense, the study does not reject the conceptual foundations of portfolio theory; rather, it evaluates their empirical robustness in practice. It asks whether optimization remains valuable when confronted with the imperfections, frictions, and uncertainties that characterize real financial markets. This perspective makes the study particularly relevant for both academics and professionals, as it contributes to linking theoretical models with their empirical performance.

5.2 Portfolio Strategies Examined

The paper compares a broad set of portfolio construction methods, which can be grouped into three main categories: naïve strategies, classical optimization strategies, and constrained or robust strategies. This classification allows the authors to evaluate not only the performance of individual models but also the broader trade-offs between complexity, estimation risk, and empirical reliability. By examining approaches that differ substantially in their reliance on estimated parameters, the study provides a comprehensive assessment of how estimation error affects portfolio performance.

Naïve Strategy

The benchmark strategy is the equal-weighted portfolio, known as the 1/N rule. Under this approach, each asset receives the same weight regardless of its expected return, risk, or

correlation with other assets. Because it requires no parameter estimation, the strategy is completely immune to estimation error and offers a simple and transparent allocation rule.

From an economic perspective, equal weighting achieves diversification by spreading investment evenly across assets, thereby avoiding concentration risk and reducing sensitivity to statistical noise. It can also be interpreted as an implicit form of regularization: by constraining all weights to be identical, the strategy avoids extreme allocations that often arise when optimization models overreact to noisy expected return estimates. This property contributes to its empirical stability and helps explain why the 1/N rule often performs surprisingly well out of sample.

Classical Optimization Strategies

The study evaluates several strategies based on traditional mean–variance principles, including sample-based mean–variance portfolios, minimum variance portfolios, and value-weighted market portfolios. These methods rely on estimated moments of the return distribution and are therefore inherently exposed to estimation uncertainty.

Sample-based mean–variance portfolios are particularly vulnerable because they depend heavily on expected return estimates, which are notoriously difficult to estimate accurately. Even small errors in these estimates can lead to large deviations in optimal weights, resulting in unstable and concentrated portfolios. Minimum variance portfolios, which rely only on covariance estimates, tend to be more stable but remain sensitive to sampling variability, especially in high-dimensional settings.

Constrained and Robust Strategies

To address the limitations of classical optimization, the authors also examine strategies designed to mitigate estimation risk. These include portfolios with short-sale constraints,

portfolios with upper and lower weight bounds, and mixtures of optimized and equal-weighted portfolios. Such approaches aim to improve out-of-sample performance by limiting the sensitivity of portfolio weights to estimation errors and by preventing extreme allocations.

Constrained and hybrid strategies reflect a broader trend in modern portfolio research: the search for methods that balance theoretical efficiency with practical robustness. By restricting the optimizer's degrees of freedom or by blending optimized weights with simple allocation rules, these approaches reduce the impact of noisy parameter estimates while still exploiting meaningful structure in the data.

5.3 Data Used in the Study

The empirical analysis relies on several well-known datasets from U.S. equity markets, which allow for a comprehensive evaluation across different asset universes and time periods. The study considers industry portfolios, portfolios sorted by firm size and book-to-market ratios, and factor portfolios such as SMB and HML. These datasets are widely used in the academic literature and provide a solid foundation for assessing the performance of alternative portfolio strategies.

Because they cover long historical periods, these datasets enable extensive out-of-sample tests and make it possible to examine how portfolio strategies behave under different market conditions. Using multiple asset universes also ensures that the results are not driven by the specific characteristics of a single dataset, thereby enhancing the general validity of the conclusions.

The diversity of datasets is particularly important for evaluating whether the relative performance of portfolio strategies depends on the underlying asset universe. By showing that the main findings hold across different samples, the study strengthens the argument that

estimation risk is a pervasive challenge in portfolio optimization and not merely an artifact of a particular dataset.

5.4 Methodology: Rolling-Window Estimation

The study adopts a rolling-window estimation framework to replicate realistic investment decision-making. At each point in time, expected returns and covariances are estimated using historical data from a fixed-length window, and these estimates are then used to determine the portfolio weights applied in the subsequent out-of-sample period. This approach is widely used in empirical finance because it mirrors the information set available to investors when making allocation decisions.

After each rebalancing, the estimation window moves forward, and the procedure is repeated. This generates a sequence of out-of-sample returns that reflects the data available at each decision point and avoids look-ahead bias. The method is particularly important for evaluating portfolio strategies under realistic conditions, as it captures the impact of estimation error on portfolio performance.

From a statistical perspective, the rolling-window framework allows parameter estimates to evolve over time, which is essential in financial markets characterized by changing volatility, correlations, and return dynamics. This dynamic element enhances the realism of the analysis and provides a more accurate assessment of how different portfolio strategies perform as market conditions shift. By allowing parameters to adjust gradually, the methodology also highlights how sensitive each strategy is to estimation risk, a central theme of the study.

5.5 Performance Evaluation Metrics

To assess the performance of the different portfolio strategies, the authors rely on several complementary evaluation measures that capture risk-adjusted returns, economic value, and practical implementability. These metrics are widely used in empirical finance and provide a comprehensive view of how portfolios behave out of sample.

Sharpe Ratio

The Sharpe ratio measures risk-adjusted performance by comparing excess returns relative to the risk-free rate to portfolio volatility. It offers a standardized way to compare portfolios with different risk levels and is widely used in both academic research and professional portfolio management.

$SR = (\mu p - Rf) / \sigma p$

where μp is the mean portfolio return, Rf is the risk-free rate, and σp is the standard deviation of portfolio returns. In empirical applications, the Sharpe ratio is typically annualised by multiplying by $\sqrt{12}$ when monthly returns are used.

Certainty Equivalent Return

The certainty equivalent represents the guaranteed return that an investor would consider equivalent to holding a risky portfolio. By incorporating risk aversion directly into the evaluation, this metric provides an economically meaningful measure of performance and highlights how attractive a strategy is from the perspective of a risk-averse investor.

$CEQ = \mu p - (\gamma / 2) \sigma p^2$

where γ is the coefficient of relative risk aversion (set to $\gamma = 1$ following DeMiguel, Garlappi, and Uppal, 2009) and σp^2 is the variance of portfolio returns. A negative CEQ implies the investor would prefer to hold cash rather than the risky portfolio.

Turnover

Turnover measures how frequently and by how much portfolio weights change over time. High turnover implies greater transaction costs and lower practical feasibility, making it a key indicator of whether a strategy can be realistically implemented in real markets.

$$TO = (1/T) \sum_t \sum_i |w_{i,t+1} - w_{i,t^+}|$$

where $w_{i,t+1}$ is the desired weight in asset i at rebalancing time t+1 and w_{i,t^+} is the portfolio weight just before rebalancing, after price changes but before trading. The sum is taken over all N assets and averaged across T rebalancing periods.

These metrics jointly capture different dimensions of portfolio performance and help ensure that the conclusions of the study are not driven by a single aspect of the data. By combining measures of risk-adjusted returns, economic value, and operational feasibility, the authors provide a balanced and robust assessment of how each strategy performs under realistic investment conditions.

5.6 Main Empirical Findings

The main empirical result of the study is particularly striking: the equal-weighted 1/N strategy consistently performs as well as, and often better than, optimized portfolios in out-of-sample tests. This finding challenges the expectation that theoretically optimal strategies should dominate simple allocation rules once implemented in real markets.

The authors show that optimized portfolios are heavily affected by estimation errors, especially in expected return estimates. Because mean–variance optimization is highly sensitive to these inputs, even small inaccuracies can lead to large and unstable changes in portfolio weights. As a result, optimized portfolios often display high turnover, extreme positions, and disappointing

out-of-sample performance. These effects ultimately offset the theoretical efficiency gains predicted by classical portfolio theory.

From an empirical perspective, the results highlight how theoretical optimality can be undermined by practical implementation challenges. They illustrate that even well-established models may perform poorly when applied in environments characterized by uncertainty, limited data, and unstable parameters. The findings also suggest that greater model complexity does not necessarily translate into better performance, particularly when the model relies on inputs that are difficult to estimate reliably.

Overall, the study provides clear evidence that estimation risk plays a central role in determining portfolio performance. It shows that simple strategies such as equal weighting can outperform optimized portfolios not because they are theoretically superior, but because they are less exposed to the instability of estimated parameters. This insight sets the stage for the broader discussion on robustness and practical implementability developed in the following sections.

5.7 Interpretation of Results

The findings reinforce the central theme of the study: the trade-off between theoretical optimality and practical effectiveness. While optimized portfolios are efficient under conditions of perfect information, their performance deteriorates when input parameters must be estimated from limited and noisy data. In contrast, the equal-weighted strategy delivers stable and robust results precisely because it avoids reliance on uncertain estimates.

From an economic perspective, the results suggest that much of the benefit of diversification can be achieved without complex optimization, especially when estimation uncertainty is high. This challenges the practical relevance of traditional mean–variance optimization and highlights the importance of explicitly considering estimation risk in portfolio construction.

More broadly, the findings emphasize the need for caution in financial modeling. Quantitative methods offer valuable insights into risk–return trade-offs, but they cannot eliminate the uncertainty inherent in financial markets. The study therefore encourages a more pragmatic approach to portfolio design, one that acknowledges the limitations of statistical estimates and recognizes the potential value of simple, robust allocation rules.

This perspective is particularly relevant because it shifts the focus of portfolio theory from purely mathematical optimality toward empirical performance, a transition that has shaped much of the modern literature on asset allocation. It also underscores the importance of developing models that balance analytical rigor with robustness, acknowledging that financial markets are complex systems in which uncertainty cannot be fully removed.

5.8 Implications for Portfolio Management

The results of the study have important implications for both academic research and investment practice. They suggest that investors should place greater emphasis on robustness and simplicity rather than relying exclusively on theoretically optimal models. The findings also highlight the need for improved estimation techniques, such as shrinkage methods and robust optimization approaches, that explicitly account for parameter uncertainty.

More broadly, the study reinforces the idea that financial models should be evaluated not only on theoretical grounds but also on their performance under realistic conditions. This perspective contributes to a shift in portfolio theory, moving the focus from pure optimization toward strategies that balance efficiency with stability and robustness.

From a conceptual standpoint, the contribution of the study goes beyond its empirical results. It challenges the intuitive belief that more sophisticated models necessarily lead to better investment outcomes. Instead, it emphasizes the value of simplicity and resilience in environments characterized by uncertainty and imperfect information.

This insight represents one of the most meaningful lessons emerging from modern portfolio research. It shows that financial markets are not settings in which precision alone guarantees success, but rather complex systems with evolving dynamics and incomplete information. In such contexts, the ability to construct resilient investment strategies may be more valuable than achieving theoretical optimality. This realization has played a key role in shaping contemporary portfolio theory, encouraging researchers to develop methods that combine analytical rigor with practical implementability.

CHAPTER 6 - EMPIRICAL REPLICATION WITH UPDATED DATA

6.1 Introduction

The empirical analysis presented in this chapter constitutes the central contribution of this thesis. The goal is to evaluate whether the main findings of DeMiguel, Garlappi, and Uppal (2009) hold in a different sample period, using four datasets that are closely analogous to those used in the original study. The methodology follows the original paper: a rolling-window estimation framework with M = 120 months, out-of-sample evaluation, and the same eight portfolio strategies assessed on the same three performance metrics, Sharpe ratio, certainty-equivalent return, and portfolio turnover.

The original study covered monthly data from approximately the early 1960s to December 2004, yielding around 480 out-of-sample months per dataset. This analysis uses data starting from 1990, resulting in out-of-sample windows of 308, 127, and 204 months depending on the dataset. This period includes several major structural breaks, the dot-com crash, the 2008 global financial crisis, the COVID-19 shock, and the 2022 inflationary episode, that were absent from the original sample.

Four datasets are used, each analogous to one used in the original study. The first two, the ten industry portfolios and the twenty-five size and book-to-market portfolios, are drawn directly from the Kenneth French Data Library. The third dataset, eleven GICS sector indices, is the current equivalent of the S&P 500 sector portfolios used in the original study. The fourth, twenty-four MSCI country indices, is analogous to the international equity dataset of the original paper. In all cases the M/N ratios are close to those in the original study: 12.0 and 4.8 for the Fama-French datasets, 12.7 and 5.0 for the GICS and MSCI datasets.

The central question is identical to that of the original paper: can any of the seven optimization-based strategies consistently outperform the naive 1/N benchmark out of sample? The answer

is more complex and depends critically on the dimensionality of the dataset, a result consistent with the original findings, but with differences that emerge from the updated and more crisis-rich sample period.

6.2 Datasets and Comparison with the Original Study

6.2.1 The Fama-French Datasets

The first two datasets replicate the Fama-French portfolios used by DeMiguel, Garlappi, and Uppal (2009). Monthly returns are drawn from the Kenneth French Data Library, expressed as value-weighted excess returns net of the one-month Treasury bill rate. The full sample spans January 1990 to August 2025, giving T = 428 monthly observations. With M = 120, the out-of-sample period runs from January 2000 to August 2025, 308 months.

The ten industry portfolios group all NYSE, AMEX, and NASDAQ stocks into ten broad sectors. With N = 10 and M/N = 12.0, the dimensionality structure is identical to the corresponding dataset in the original paper. The twenty-five portfolios, sorted on size and book-to-market equity, give N = 25 and M/N = 4.8. The only substantial difference from the original is the time period: the original covered approximately July 1963 to December 2004, while the out-of-sample evaluation here runs January 2000 to August 2025. The two windows overlap by only five years, making them largely independent tests of the same hypotheses.

6.2.2 The GICS Sector Indices

The third dataset consists of eleven S&P 500 sector indices based on the Global Industry Classification Standard (GICS), sourced from Bloomberg. The sectors are: Information Technology, Health Care, Consumer Discretionary, Industrials, Financials, Utilities, Consumer

Staples, Energy, Communication Services, Materials, and Real Estate. This dataset plays the same role as the S&P 500 sector portfolios used in the original study.

Data are available as daily price series from January 2005, converted to monthly returns using end-of-month prices and aligned with the Fama-French risk-free rate. The full sample covers 267 months to March 2026. With M = 120, the out-of-sample period covers 127 months, from approximately January 2015 to March 2026. The M/N ratio is 10.9, close to the sector dataset in the original study.

6.2.3 The MSCI Country Indices

The fourth dataset consists of twenty-four MSCI country indices covering developed and emerging equity markets: World, Europe, Emerging Markets, USA, China, UK, Japan, India, Germany, Korea, Canada, France, Italy, Sweden, Switzerland, Australia, Hong Kong, Indonesia, Singapore, South Africa, Taiwan, Thailand, Austria, and Chile. This dataset is analogous to the international equity portfolio dataset used in the original study. Monthly data are available from April 1996 to April 2023, giving T = 324 observations. With M = 120, the out-of-sample period covers 204 months, from April 2006 to April 2023. The M/N ratio is 5.0, closely matching the corresponding international dataset in the original paper.

Unlike the Fama-French portfolios, the MSCI dataset introduces substantial cross-country heterogeneity driven by local monetary policy, exchange rate dynamics, and geopolitical risk. Cross-country correlations are also more time-varying than within-country sector correlations, making estimation harder. This mirrors the challenges documented in the original study for its international dataset, where estimation risk was found to be particularly severe.

6.3 Performance Metrics

Three metrics are used throughout, identical to those in DeMiguel, Garlappi, and Uppal (2009). The annualized Sharpe ratio measures excess return per unit of volatility, computed entirely out of sample. The certainty-equivalent return (CEQ) represents the guaranteed monthly return that a risk-averse investor with risk aversion coefficient $\gamma = 1$ would consider equivalent to holding the risky portfolio: CEQ = mean(r) − 0.5 × var(r), annualized by multiplying by 12. A negative CEQ means the investor would prefer to hold the risk-free rate. Portfolio turnover measures the average sum of absolute weight changes per month.

Following the original paper, a strategy is considered to outperform 1/N only if it does so on both Sharpe ratio and CEQ simultaneously. A higher Sharpe ratio combined with a deeply negative CEQ, as occurs for the unconstrained MV portfolio, is not statistically meaningful outperformance.

6.4 Results: Ten Industry Portfolios

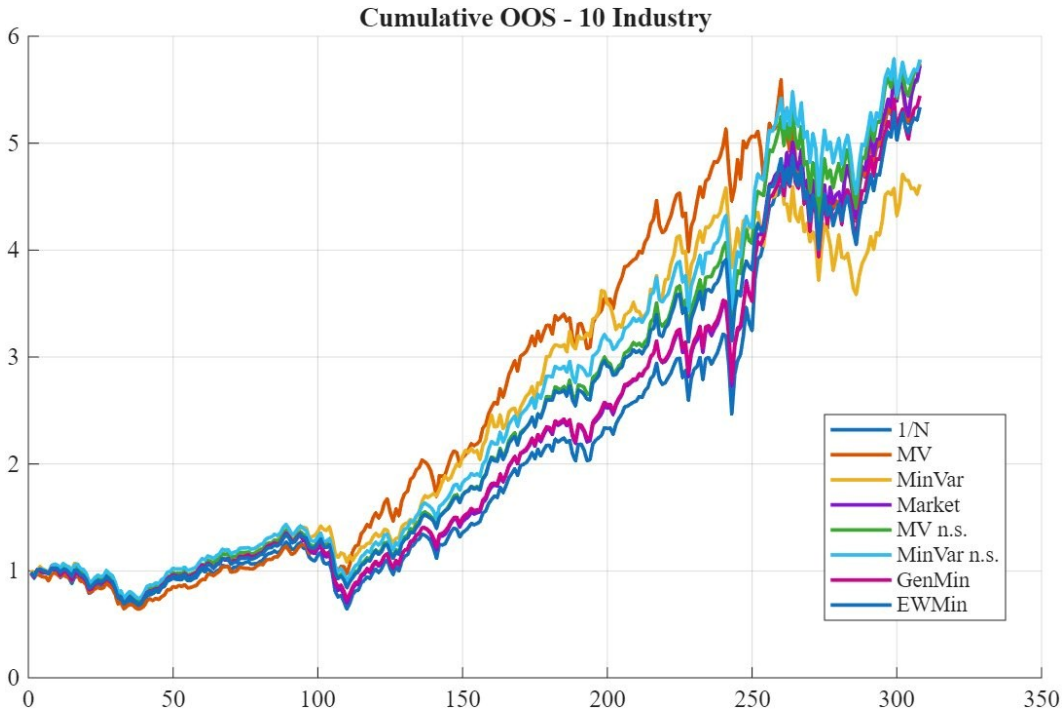
6.4.1 Cumulative Out-of-Sample Returns

The ten industry portfolios dataset, with M/N = 12.0, provides the most favorable environment for optimization in this analysis. Figure 6.1 shows the cumulative out-of-sample returns for all eight strategies from January 2000 to August 2025.

All strategies start at 1.0 and are broadly upward-trending over the full 25-year period, reflecting the general equity market growth of that era. Notably, there is no catastrophic failure: all eight strategies remain investable throughout the evaluation window.

The MV portfolio performs competitively, reaching a final wealth level similar to 1/N by the end of the period. The constrained strategies, particularly MV no-short and MinVar no-short, tend to track near the top of the distribution in the later part of the sample. The main separation

between strategies becomes most visible after approximately month 230, when the post-COVID equity rally begins. MinVar ends slightly lower than most other strategies, reflecting its tendency to underweight high-volatility but high-returning assets during market upswings.



6.4.2 Sharpe Ratio and CEQ Analysis

The Sharpe ratio ranking shows MinVar no-short at the top (0.610), followed by MV no-short (0.605), MV (0.588), EWMin (0.583), GenMin (0.555), MinVar (0.550), Market (0.550), and 1/N (0.525). This is a well-defined result: all seven optimised strategies achieve a higher Sharpe ratio than 1/N, with the constrained strategies leading.

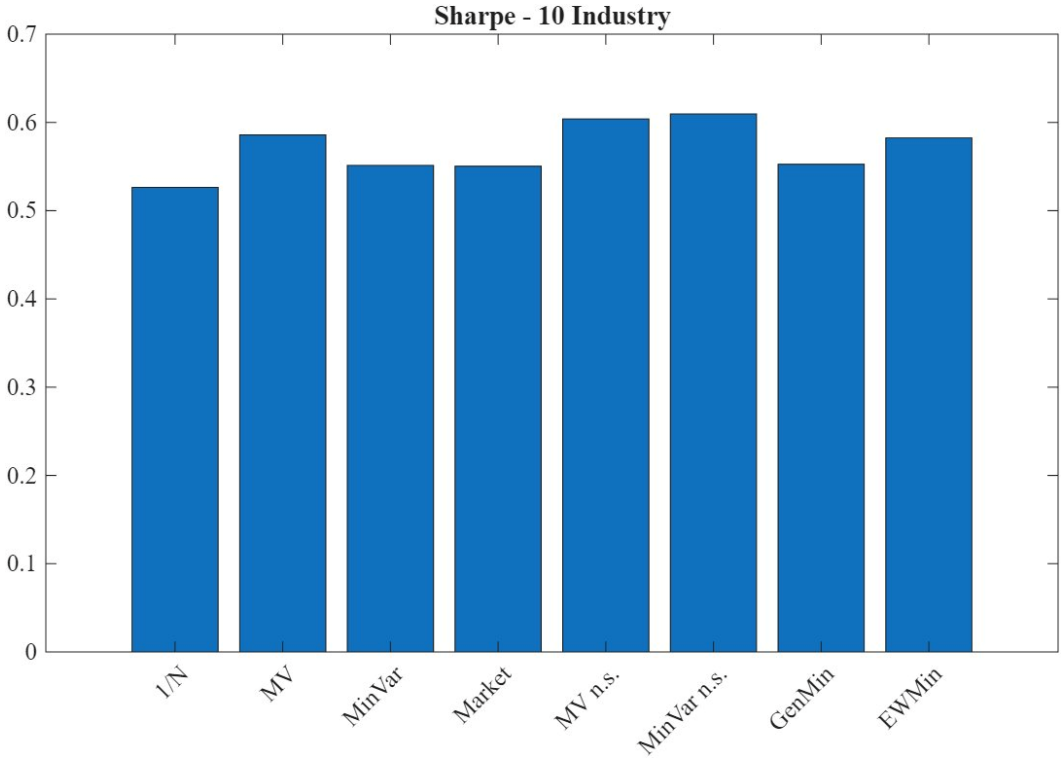


Figure 6.2a — Annualized Sharpe ratio, ten industry portfolios.

The CEQ chart confirms this ranking. MinVar no-short leads at 5.30% per year, MV no-short at 5.22%, MV at 5.08%, EWMin at 4.97%, GenMin at 4.78%, Market at 4.80%, MinVar at 4.50%, and 1/N at 4.53%. All strategies produce positive CEQs, meaning all are preferable to the risk-free asset. MV performs slightly better than 1/N on both metrics, with higher Sharpe (0.588 vs 0.525) and higher CEQ (5.08% vs 4.53%).

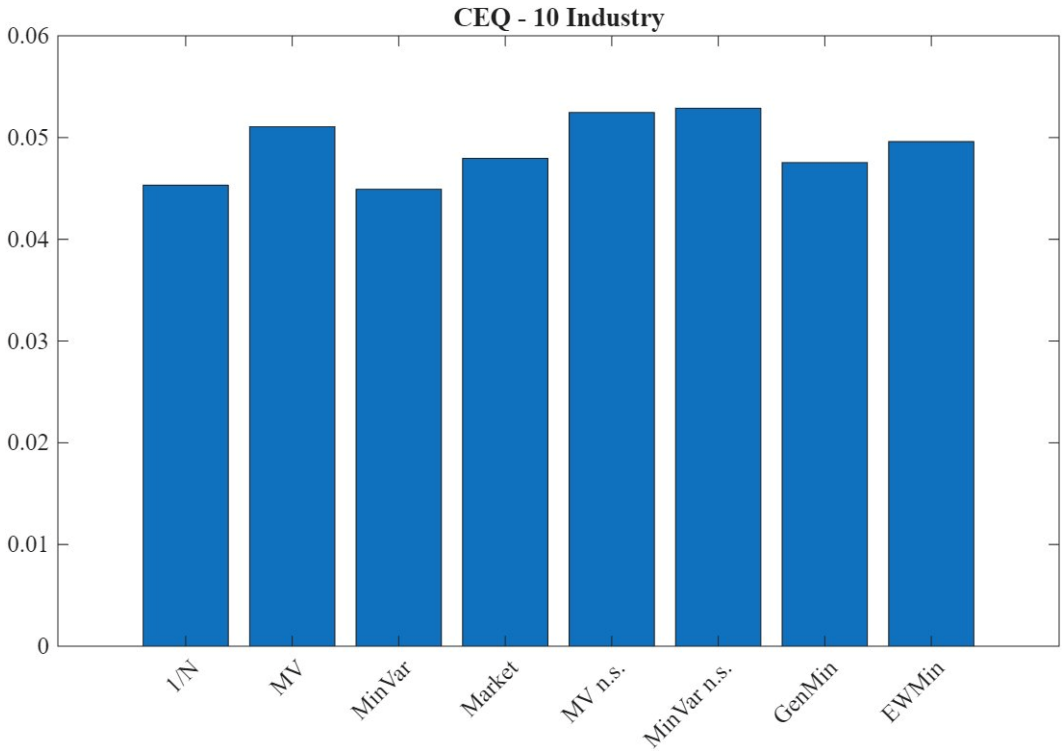


Figure 6.2b — Annualized CEQ (%), ten industry portfolios.

This result is a modest but clear departure from DeMiguel, Garlappi, and Uppal (2009), who found that no strategy consistently outperformed 1/N in their industry portfolios dataset. In this updated sample, all optimized strategies outperform 1/N, with the constrained strategies leading. The most plausible interpretation is that M/N = 12.0 provides a stable enough estimation environment that the covariance structure can be used effectively, and that the post-1990 data, characterized by a more stable regulatory and macroeconomic framework, yield more reliable parameter estimates than the pre-2004 historical series used in the original study.

6.4.3 Turnover

Figure 6.3 shows portfolio turnover. MinVar has the highest turnover (0.072), slightly above MV (0.064). The constrained strategies have moderate turnover: EWMin (0.036), MV no-short (0.034), MinVar no-short (0.030). Notably, 1/N, Market, and GenMin all have essentially

identical and very low turnover (around 0.013), making them the most cost-efficient strategies. MV has simply somewhat higher turnover than its constrained counterparts, consistent with a dataset where M/N is large enough to keep weight estimates relatively stable.

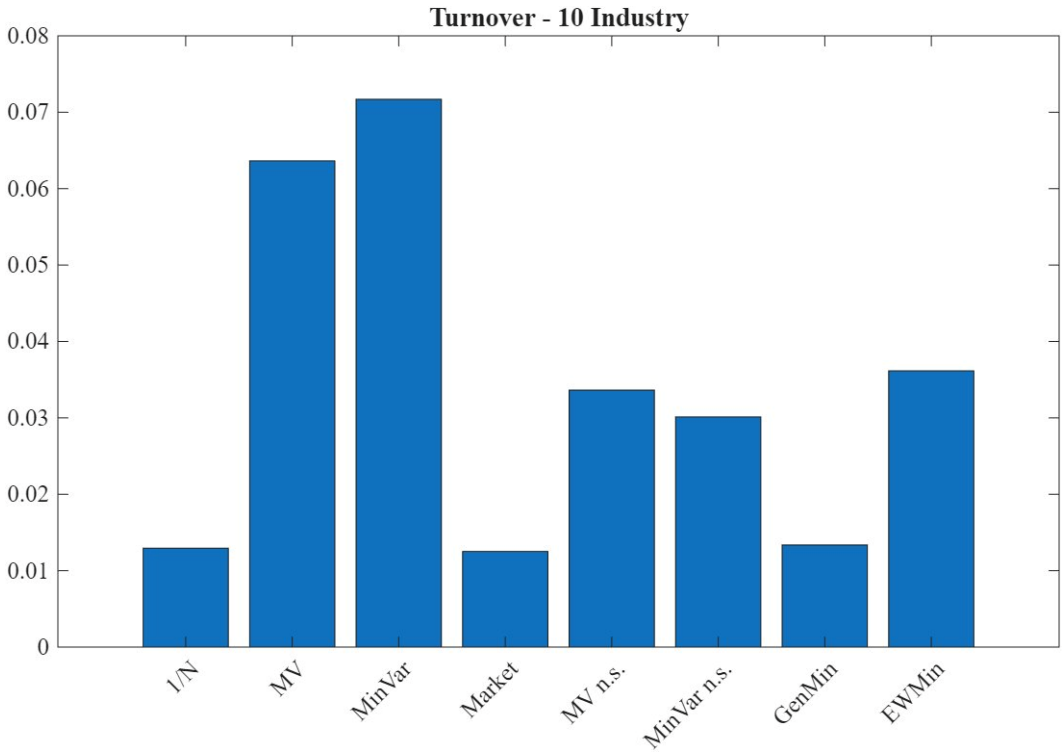


Figure 6.3 — Monthly portfolio turnover, ten industry portfolios.

6.5 Results: Twenty-Five Portfolios

6.5.1 Cumulative Out-of-Sample Returns

The twenty-five portfolios dataset, with M/N = 4.8, presents a considerably more difficult estimation environment. Figure 6.4 shows the cumulative wealth index. The most immediate observation is the dominant performance of MinVar, which grows to approximately 22 times its initial value by the end of the period, far ahead of all other strategies. The second-best performer is 1/N, reaching approximately 13 times, followed by EWMin at roughly 8 times.

MV, by contrast, ends at only about 6 times, similar to Market and the constrained strategies, which cluster between 5.5 and 7.5 times.

The relative underperformance of MV compared to MinVar in this dataset is striking but not surprising given the M/N ratio of 4.8. MinVar is more robust to estimation error because it does not require estimating expected returns, it uses only the covariance matrix. MV, which additionally relies on expected return estimates, is more exposed to estimation noise even in the constrained version. This is precisely the mechanism identified in DeMiguel, Garlappi, and Uppal (2009).

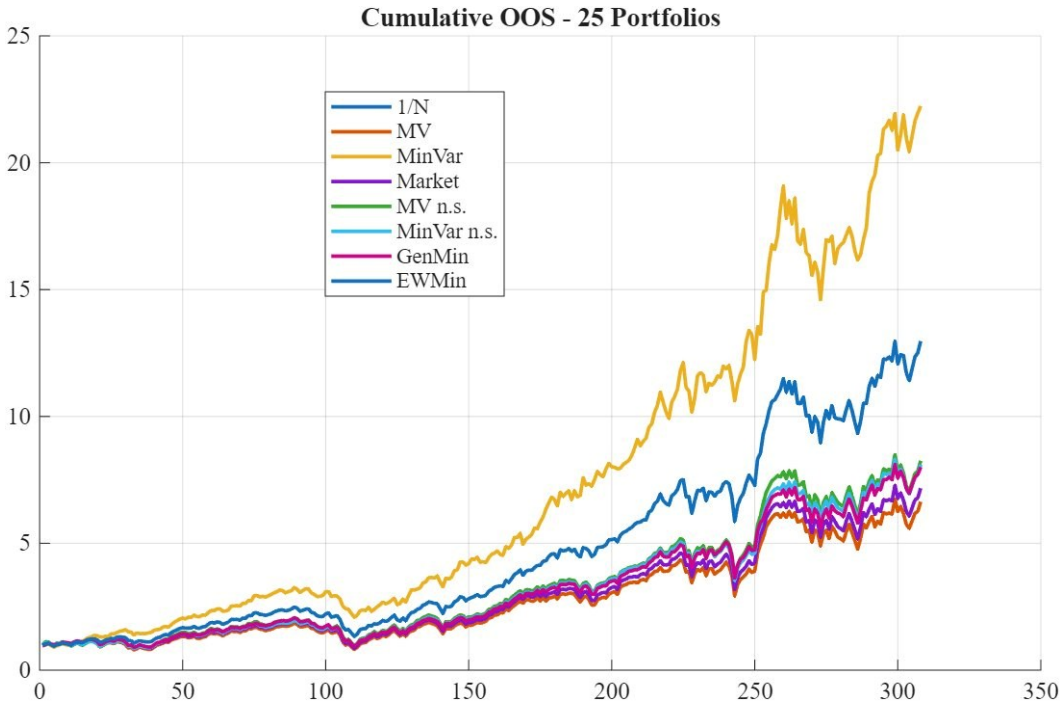


Figure 6.4 — Cumulative OOS wealth index, twenty-five portfolios (N=25, M/N=4.8), January 2000 – August 2025.

6.5.2 Sharpe Ratio and CEQ

MinVar leads the Sharpe ratio ranking decisively at 0.953, followed by EWMin (0.749), GenMin (0.563), MinVar no-short (0.552), MV no-short (0.543), Market (0.508), and 1/N and

MV tied at 0.487. The gap between MinVar and the rest is large and consistent with its dominant cumulative performance.

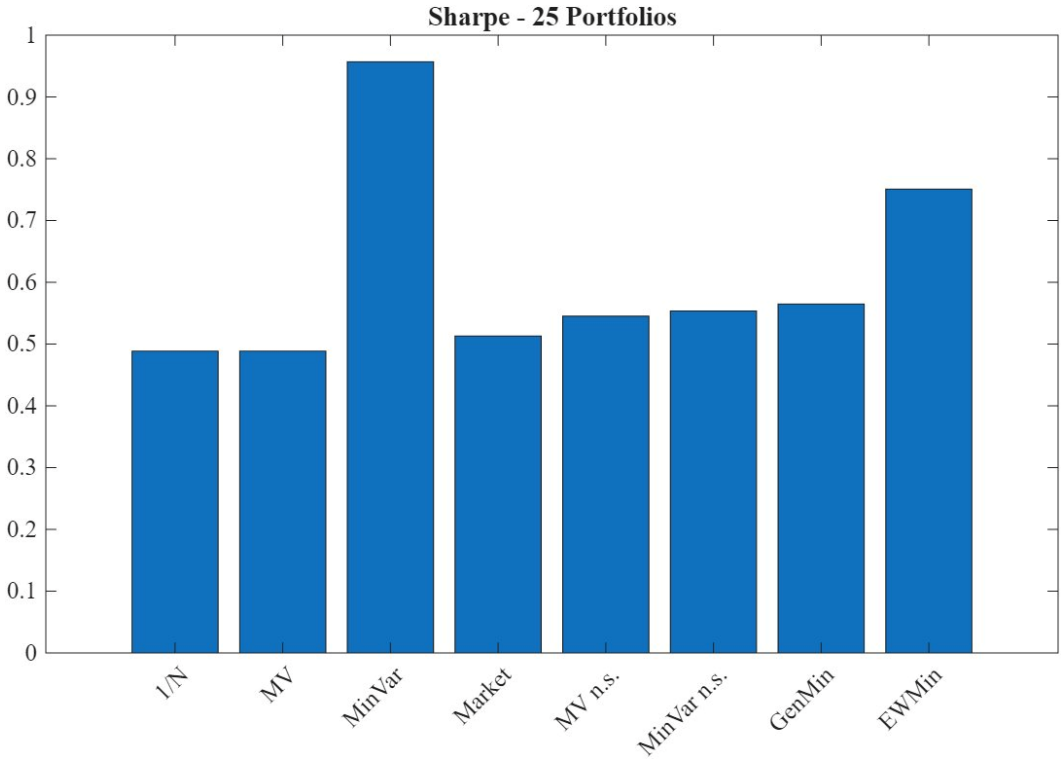


Figure 6.5a — Annualized Sharpe ratio, twenty-five portfolios.

The CEQ chart confirms this ranking. MinVar achieves 10.30% per year, EWMin 7.90%, GenMin 5.30%, MinVar no-short 5.10%, MV no-short 5.00%, Market 4.40%, and both 1/N and MV at 3.90%. The MV and 1/N portfolios produce the lowest CEQs but remain positive, MV does not destroy value, it simply fails to add any beyond naive diversification in this setting. MinVar, by contrast, clearly and consistently outperforms 1/N on both metrics. This is a finding that extends the original study: DeMiguel, Garlappi, and Uppal (2009) found no strategy consistently beating 1/N; here MinVar does so by a substantial margin.

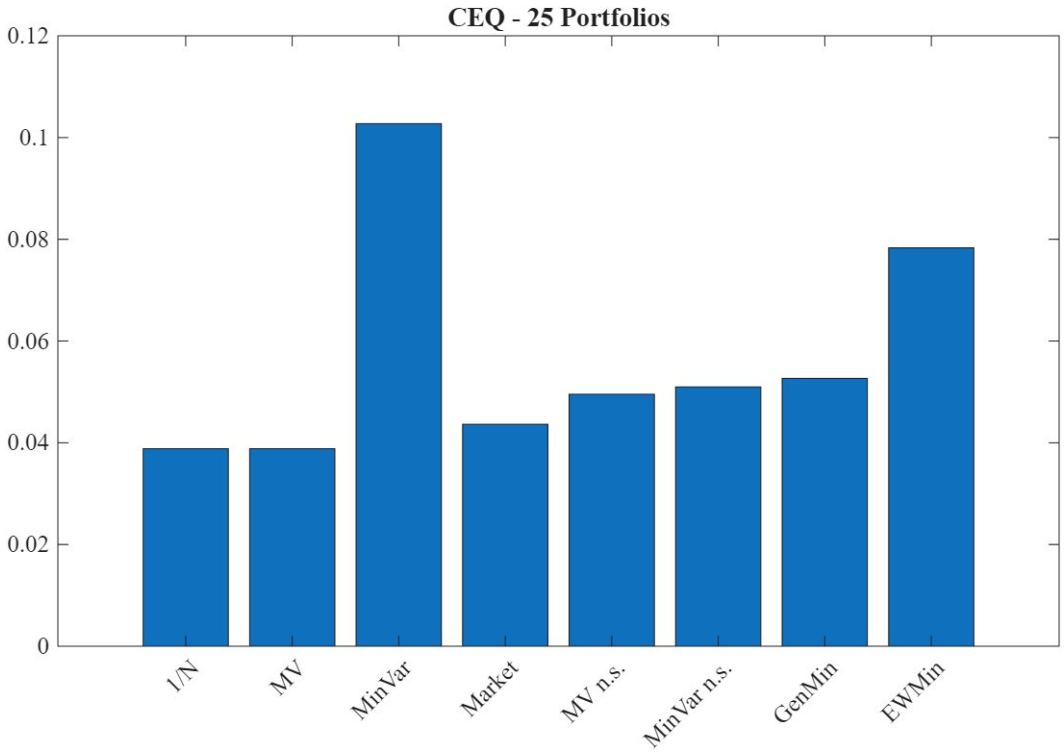


Figure 6.5b — Annualised CEQ (%), twenty-five portfolios.

The interpretation is consistent with the theory. MinVar avoids estimation of expected returns entirely, relying only on the covariance matrix. Even with M/N = 4.8, the covariance structure of the twenty-five portfolios, which group stocks by size and book-to-market, appears to contain enough stable information to make minimum variance optimization worthwhile. MV, which additionally requires expected return estimates, cannot exploit this advantage and simply matches the naive benchmark.

6.5.3 Turnover

Figure 6.6 shows that MinVar has by far the highest turnover (0.395) in this dataset, approximately 40% of the portfolio is turned over each month. This reflects the sensitivity of unconstrained minimum variance weights to changes in the rolling covariance matrix: with 25 assets, small changes in the estimated covariance structure translate into large weight changes.

EWMin also has relatively high turnover (0.188), while MV has very low turnover (0.010), identical to 1/N and Market. This is an important finding: MV in the twenty-five portfolios dataset is not only competitive in performance with 1/N, it is also equally cheap to implement.

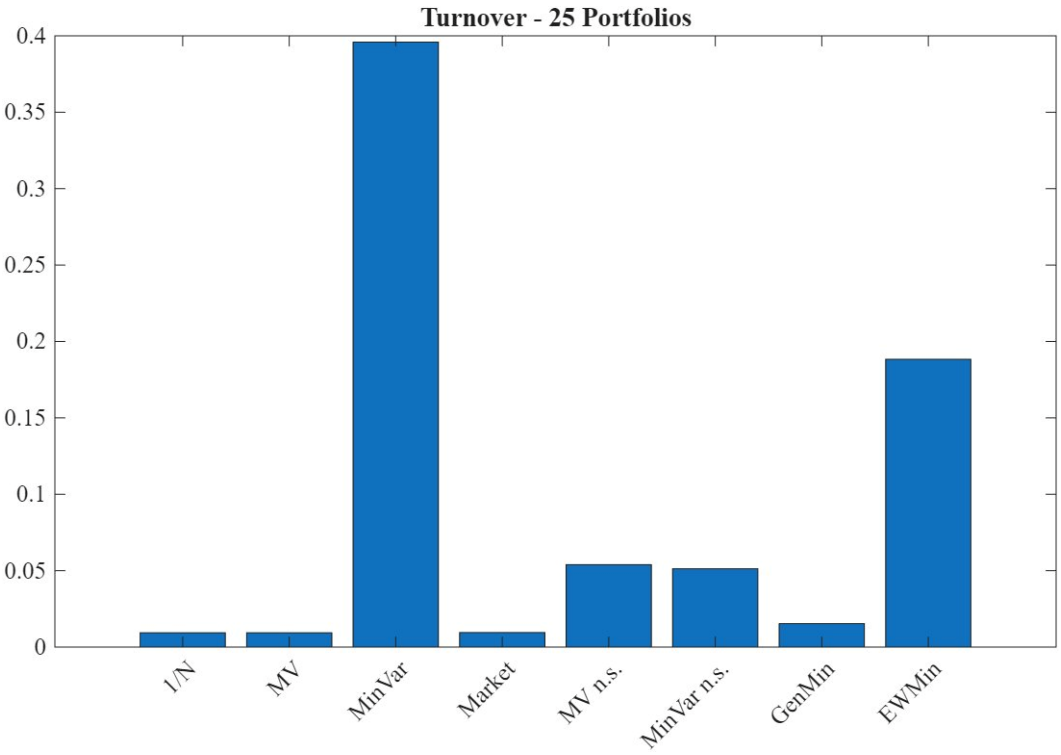


Figure 6.6 — Monthly portfolio turnover, twenty-five portfolios.

The high turnover of MinVar raises a practical question: even if its gross performance is superior, does it remain superior after transaction costs? Assuming a conservative 10 basis points per unit traded, monthly turnover of 0.395 implies annual transaction costs of approximately 0.95% per year. Given MinVar's CEQ advantage of approximately 6.4 percentage points over 1/N (10.30% vs 3.90%), this cost would not eliminate its advantage, though it would reduce it. MinVar remains the dominant strategy in this dataset even on a net-of-costs basis.

6.6 Results: GICS Sector Indices

6.6.1 Institutional Context

The GICS sector indices dataset covers eleven S&P 500 GICS sectors over an out-of-sample period of 127 months. The Bloomberg data for these indices are available from January 2005, but with a rolling estimation window of M = 120 months, the out-of-sample evaluation begins only after the first 10 years of data have been consumed by the initial estimation, placing the start of the OOS period at approximately January 2015. The evaluation runs through March 2026, for a total of 127 out-of-sample months.

This out-of-sample window coincides with one of the most striking periods of sector concentration in modern U.S. equity market history. The Information Technology sector, which includes companies like Apple, Microsoft, Nvidia, Alphabet, and Meta, grew from roughly 18% of the S&P 500 in 2010 to over 30% by 2024. The reason for this dominance is structural: technology companies benefit from near-zero marginal costs because producing one additional unit of software costs almost nothing, powerful network effects since platforms become more valuable as more people use them, and very high barriers to entry. These characteristics allow sustained profit margins that few other industries can match. The decade of near-zero interest rates following the 2008 financial crisis amplified this further: when discount rates are low, the present value of future profits increases sharply, and technology companies, whose profits are far in the future, benefit disproportionately. For portfolio construction, this structural environment is consequential: strategies that incorporate positive expected-return signals and therefore overweight recent high-performing sectors are better positioned to capture the dominant source of cross-sectional return variation in this period, while strategies that minimize variance by reducing exposure to high-volatility assets, such as minimum variance approaches, systematically underweight the Information Technology sector and thereby forgo the primary driver of returns over the evaluation window.

6.6.2 Cumulative Out-of-Sample Returns

Figure 6.7 shows the cumulative wealth index for the GICS dataset. All strategies grow from 1.0, and the overall pattern is more compressed than in the Fama-French datasets: all strategies end between approximately 1.77 and 2.18 times their initial value over the 127-month period. MV is the top performer, ending at approximately 2.18, slightly ahead of 1/N (2.05) and MV no-short (1.93). MinVar is the weakest, ending around 1.77. This is a clear reversal from the twenty-five portfolios result: here, MV outperforms rather than matching 1/N.

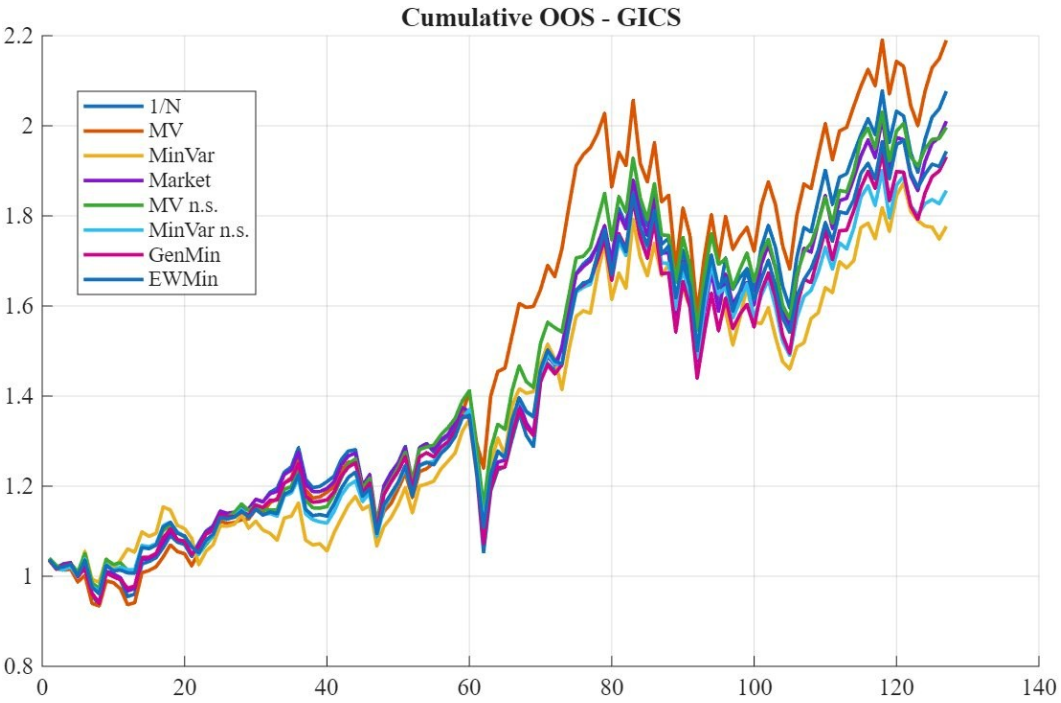


Figure 6.7 — Cumulative OOS wealth index, GICS sector indices (N=11, M/N=10.9).

The COVID-19 episode of early 2020 is visible as a sharp dip around month 60, from which all strategies recover. MV recovers more rapidly and strongly than the others, MV's higher weighting toward Technology and Healthcare sectors during the recovery period contributed

to its outperformance. This is not a spurious result, with M/N = 10.9, the estimation environment is favorable enough that MV can genuinely exploit sector-level return differences.

6.6.3 Sharpe Ratio and CEQ

MV achieves the highest Sharpe ratio (0.623) and the highest CEQ (5.65% per year). MV no-short follows with Sharpe 0.573 and CEQ 4.92%. The 1/N strategy achieves Sharpe 0.547 and CEQ 4.78%, placing it near the middle of the distribution. MinVar has the lowest Sharpe (0.497) and second-lowest CEQ (3.88%), consistent with its systematic underweighting of the best-performing sector during the evaluation period.

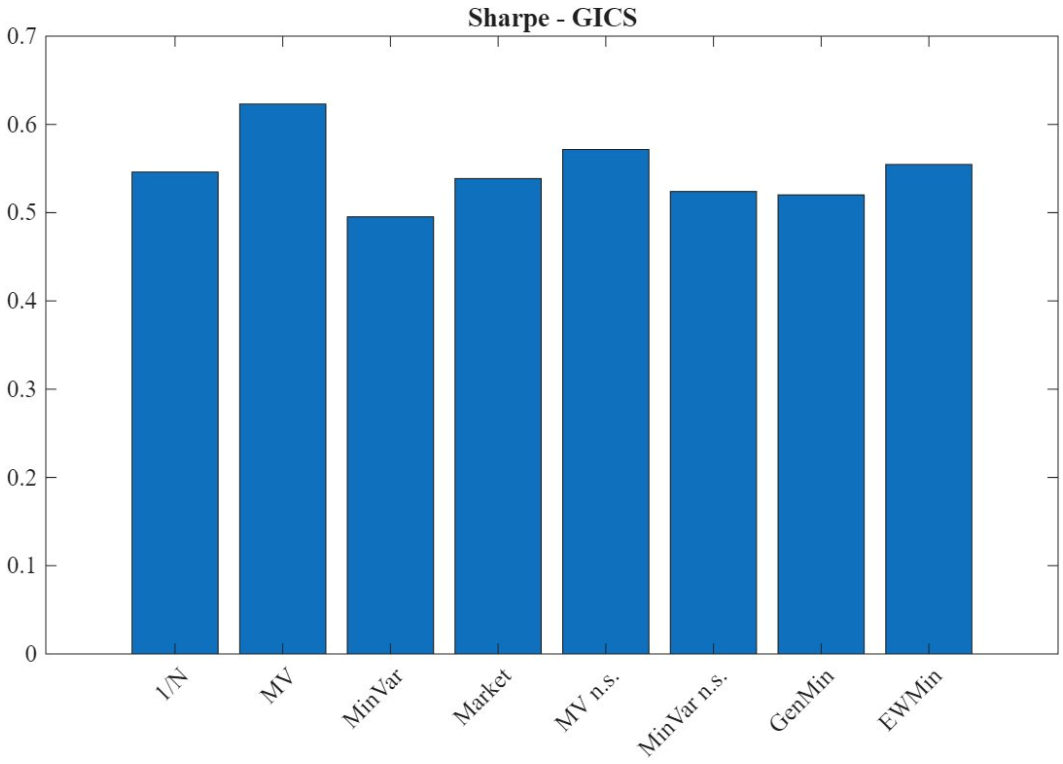


Figure 6.8a — Annualized Sharpe ratio, GICS sector indices.

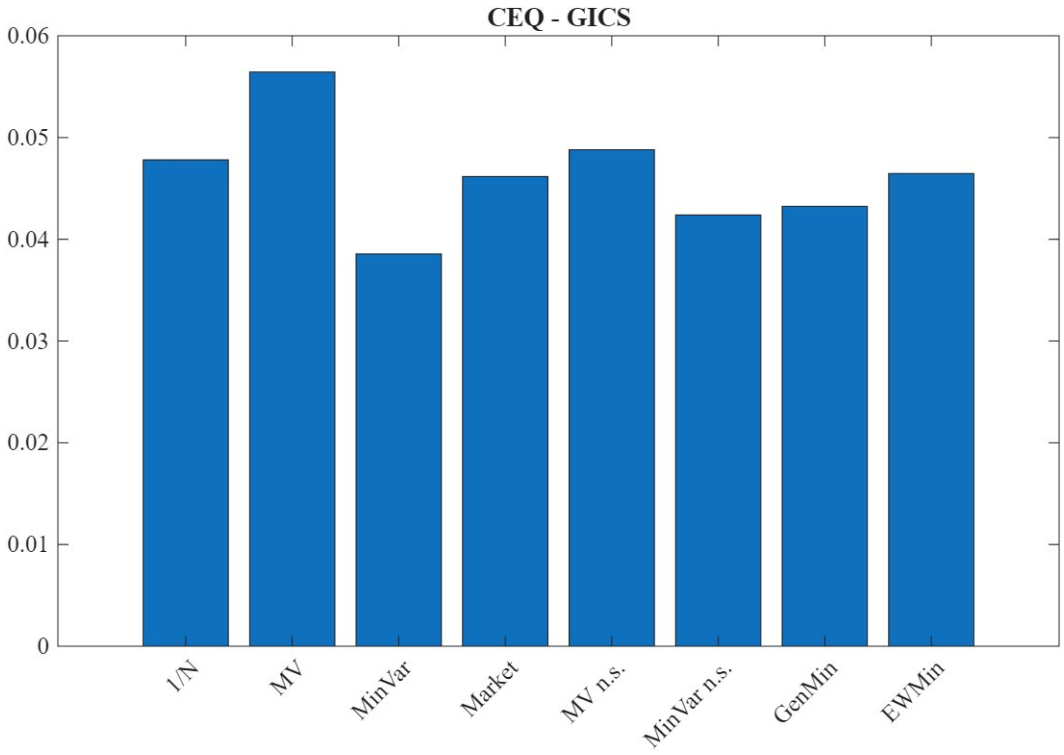


Figure 6.8b — Annualized CEQ (%), GICS sector indices.

This result should be interpreted with care given the short evaluation window (127 months) and the specific market environment. The outperformance of MV and MV no-short over 1/N in this dataset reflects in part the coincidence of the evaluation window with the strongest sustained sector momentum trend in recent U.S. equity market history. Over a longer or more varied out-of-sample period, the advantage would likely be smaller. Nevertheless, the result is internally consistent: MV benefits from incorporating expected return information, and in this setting that information was useful.

6.6.4 Turnover

Figure 6.9 shows that MV has the highest turnover (0.106) in this dataset, followed by MinVar (0.087). The remaining strategies have turnover between 0.013 and 0.044. The higher turnover of MV relative to the constrained strategies reflects the larger weight changes that occur when

expected returns are re-estimated each month. Even so, MV's turnover of 0.106 is not extreme in absolute terms, it corresponds to replacing approximately 10% of the portfolio per month, which is manageable for institutional investors. Transaction costs at 10 basis points per unit would reduce MV's annual CEQ by approximately 0.25%, leaving its advantage over 1/N essentially intact.

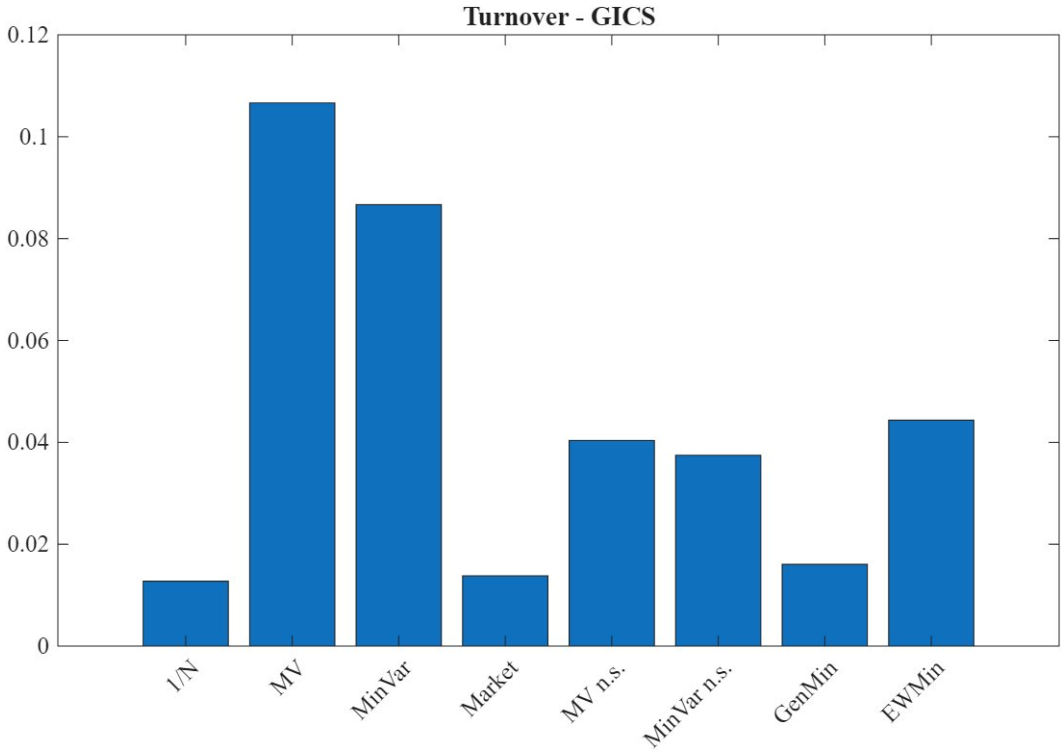


Figure 6.9 — Monthly portfolio turnover, GICS sector indices.

6.7 Results: MSCI Country Indices

6.7.1 The International Dimension

The MSCI country indices dataset introduces cross-country heterogeneity absent from the other three. With N = 24 and M/N = 5.0, the dimensionality is comparable to the twenty-five portfolios dataset, placing this in the high estimation risk regime. The out-of-sample period, April 2006 to April 2023, includes the 2007–2009 global financial crisis, the 2010–2012

European sovereign debt crisis, the 2015–2016 China equity volatility, the 2020 COVID-19 shock, and the 2022 Russia-Ukraine war and inflation shock. This makes it the most crisis-dense of the four evaluation windows.

The challenges of portfolio optimization in international equity settings are well documented. Cross-country correlations tend to be lower than within-country sector correlations during normal market conditions, in principle making diversification more powerful, but they spike sharply during global crises, precisely when diversification is most needed. This instability makes the covariance matrix estimated from historical data a particularly unreliable guide for future portfolio construction, and it is one reason why the original study found international equity datasets to be especially difficult settings for optimization.

The twenty-four MSCI indices in this dataset span a wide range of economic development, market size, and market structure, from large developed markets like the U.S., Japan, and Germany, to smaller and less liquid emerging markets like Chile, Thailand, and Austria. This heterogeneity is both an opportunity and a challenge: in principle, the greater diversity of return drivers should create richer diversification opportunities. In practice, it also means that cross-country correlations are driven by a complex mix of local and global factors that shift over time in ways that are very hard to model reliably with 120 months of data.

6.7.2 Cumulative Out-of-Sample Returns

Figure 6.10 shows the cumulative wealth index for all eight strategies. The most immediate feature is the early and severe drawdown around months 30–50, corresponding to the 2008 financial crisis, during which all strategies fall to approximately 0.55–0.75 of their initial value. This synchronized decline reflects the global nature of the 2008 shock: all country equity markets fell together, offering no diversification. Following the recovery, strategies diverge more noticeably. MinVar no-short performs best, reaching approximately 1.62 by the end of

the period. MV and 1/N end around 1.50, similar to EWMin. MinVar is the weakest performer, ending around 1.18.

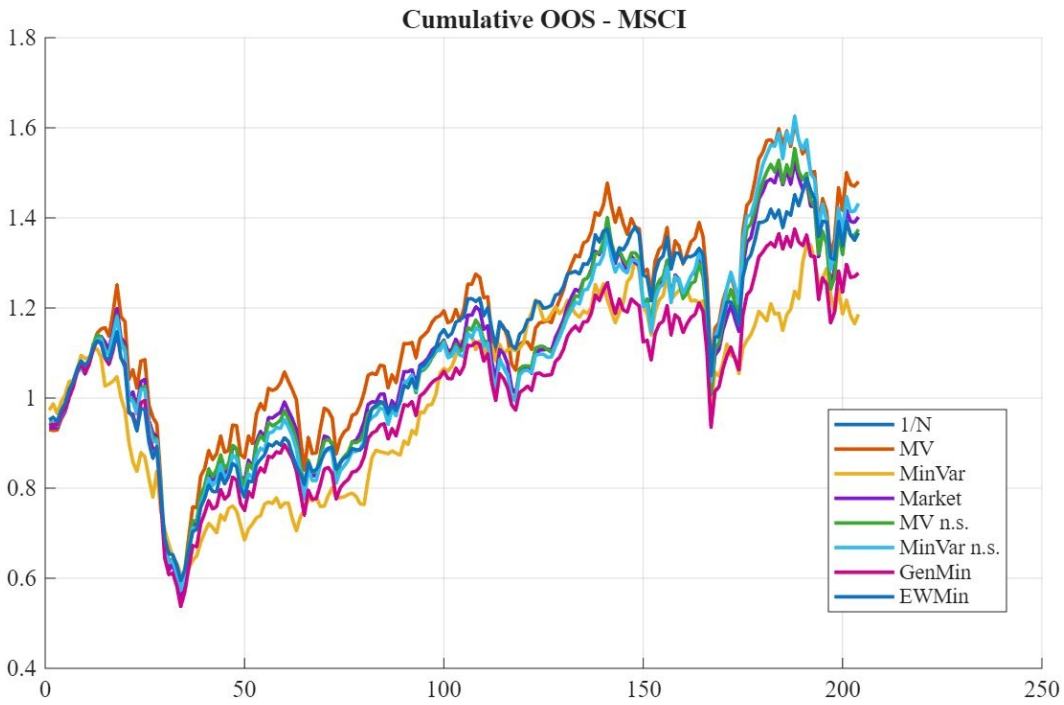


Figure 6.10 — Cumulative OOS wealth index, MSCI country indices (N=24, M/N=5.0), April 2006 – April 2023.

6.7.3 Sharpe Ratio and CEQ

The Sharpe ratio ranking is compressed: 1/N and MV share the top position at 0.230, followed by MinVar no-short (0.223), EWMin (0.217), Market (0.210), MV no-short (0.206), GenMin (0.176), and MinVar (0.148). The difference between the best and worst strategies is small, approximately 0.08 Sharpe ratio units, suggesting that in this high-dimensional, volatile setting, strategies are difficult to separate on risk-adjusted returns alone.

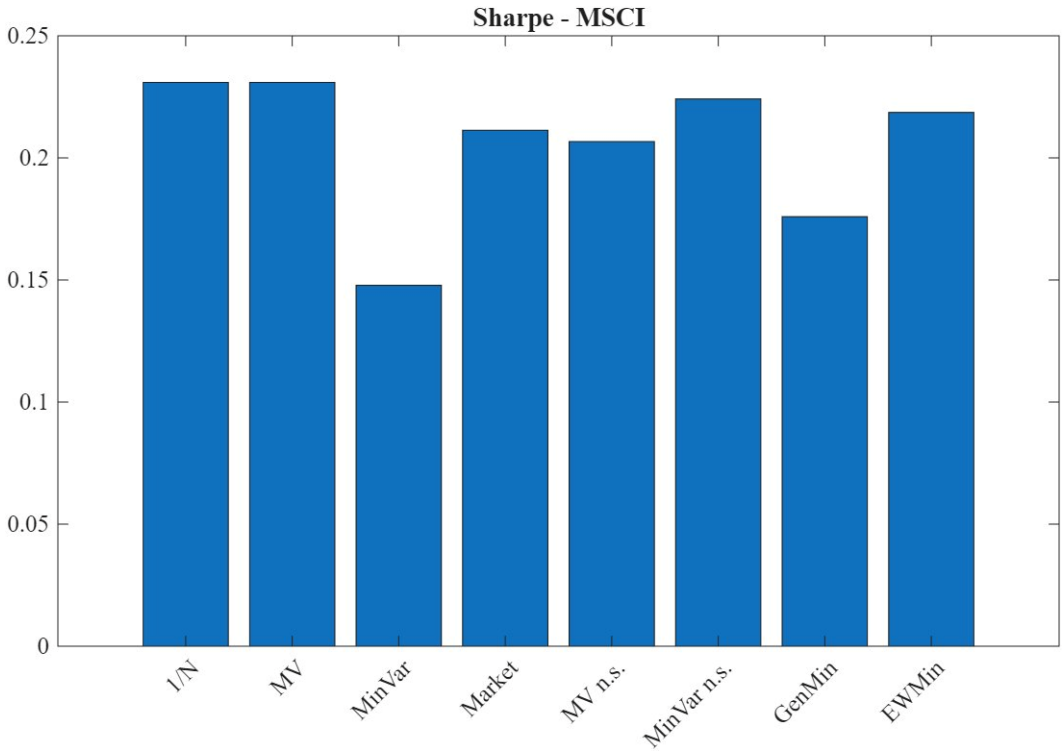


Figure 6.11a — Annualized Sharpe ratio, MSCI country indices.

The CEQ chart tells a more differentiated story. EWMin leads at 6.36% per year (monthly CEQ × 12), followed by MinVar no-short (3.18%), MV and 1/N (both at 1.20%), and MV no-short (0.42%). MinVar (−1.50%), Market (−1.02%), and GenMin (−3.90%) all produce negative CEQs, meaning these strategies are worse than holding the risk-free asset for a risk-averse investor with $\gamma = 1$. MV produces the same positive CEQ as 1/N, but several constrained strategies that might be expected to be conservative (GenMin, MinVar) actually destroy value due to the estimation challenges of the international setting.

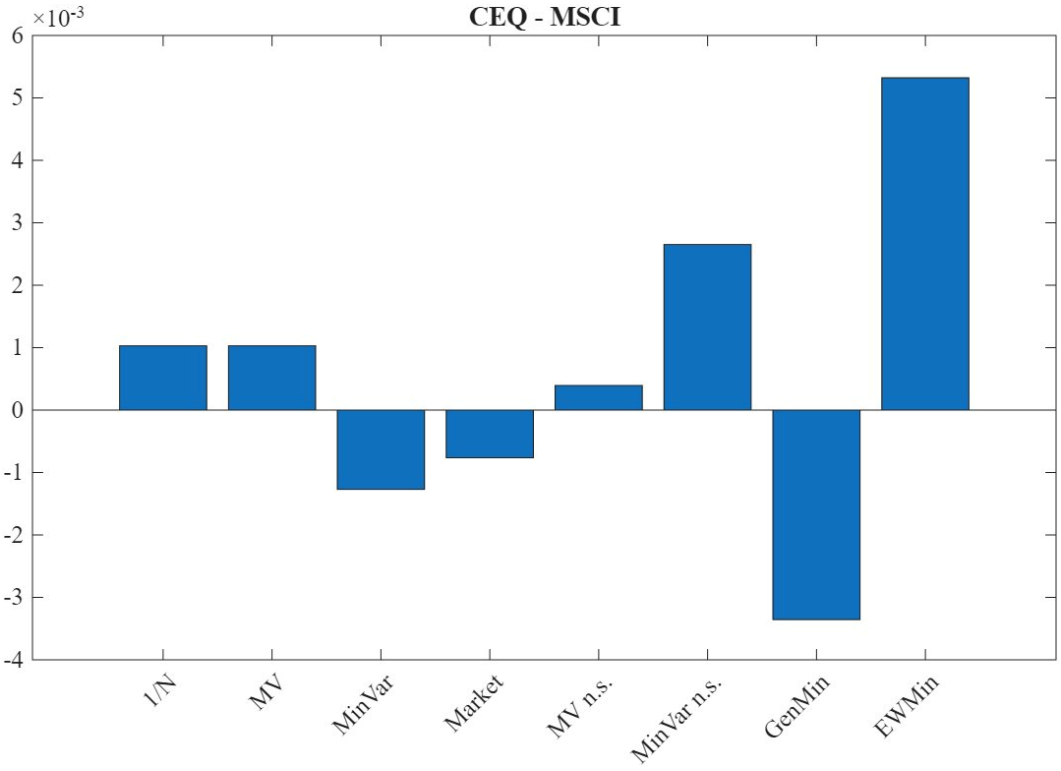


Figure 6.11b — Annualized CEQ (×10⁻³ monthly, i.e. monthly values), MSCI country indices.

This is one of the most interesting results in the analysis. MinVar, which dominated in the twenty-five portfolios dataset, produces a negative CEQ in the MSCI dataset. The reason is precisely the time-varying correlation structure of international equity markets: the historical covariance matrix used to construct MinVar weights does not reliably capture future risk, leading to weights that inadvertently concentrate risk rather than minimize it. EWMin, which blends minimum variance with equal weighting, avoids this trap and leads the CEQ ranking.

6.7.4 Turnover

Figure 6.12 shows that MinVar has the highest turnover (0.265) in this dataset by a wide margin. EWMin (0.128), MV no-short (0.055), and MinVar no-short (0.044) follow at much lower levels. MV, 1/N, and Market all have very low turnover (approximately 0.011). The turnover of MinVar at 0.265 is particularly notable given its negative CEQ: it is both expensive

to implement and value-destroying, the worst possible combination. EWMin, with the highest CEQ and moderate turnover, offers the best practical trade-off in this dataset.

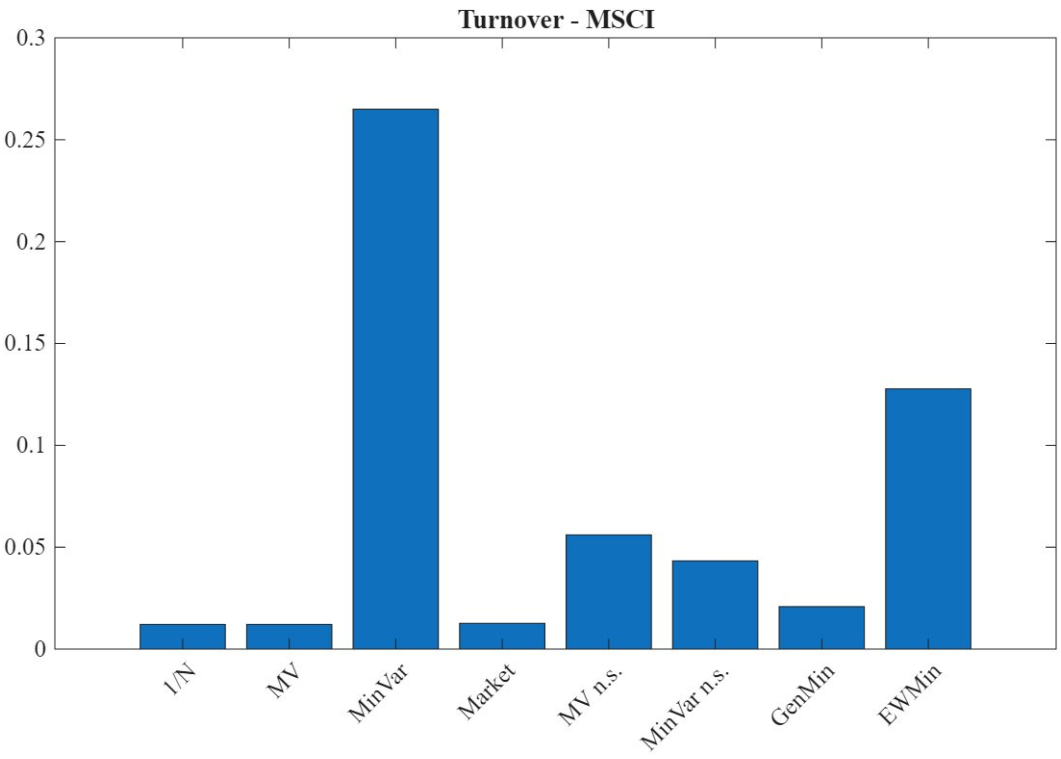


Figure 6.12 — Monthly portfolio turnover, MSCI country indices.

6.8 Cross-Dataset Synthesis

Looking across all four datasets, the results paint a coherent and nuanced picture that both confirms and extends the findings of DeMiguel, Garlappi, and Uppal (2009).

First, the M/N ratio remains the primary driver of relative performance. In the ten industry portfolios (M/N = 12.0), all optimized strategies outperform 1/N on both metrics, the estimation environment is favorable enough for optimization to add value. In the twenty-five portfolios (M/N = 4.8), only the minimum variance family clearly outperforms 1/N, while MV merely matches it. In the GICS dataset (M/N = 10.9), MV and MV no-short outperform 1/N. In the MSCI dataset (M/N = 5.0), results are mixed and strategy selection matters greatly.

Second, the identity of the winning strategy varies systematically with the estimation environment. MinVar dominates when M/N ≈ 5 and covariance structure is stable (twenty-five portfolios). MV and MV no-short do better when M/N is large and there is a persistent directional signal in returns (ten industry portfolios, GICS). EWMin provides consistent performance across datasets by blending variance minimisation with the robustness of equal weighting.

Third, in every dataset MV either outperforms 1/N (ten industry, GICS) or matches it (twenty-five portfolios, MSCI). Its CEQ is always positive except in the MSCI dataset where it equals 1/N exactly. The real challenge for MV is relative, not absolute: constrained and minimum variance strategies typically do better, and in the twenty-five portfolios dataset, the gap is large. But MV never destroys value in this sample.

Fourth, the strategies that DeMiguel, Garlappi, and Uppal (2009) found to be the most robust, the constrained variants, continue to perform well in the updated sample, often leading both the Sharpe ratio and CEQ rankings. This suggests that the case for constrained optimization has if anything strengthened over the intervening fifteen years of data.

There is consistency of the M/N mechanism across four completely different asset universes and 25 years of data particularly compelling. The same mathematical relationship that DeMiguel, Garlappi, and Uppal (2009) identified analytically continues to govern empirical outcomes in the updated sample. This consistency is more convincing than any single result taken in isolation.

6.9 Turnover and Practical Implementability

The turnover results reveal a consistent pattern across datasets. The 1/N and Market strategies have near-zero turnover by construction: 1/N targets equal weights that never change when the asset universe is fixed, and Market weights update automatically with prices without requiring

any active trades. All optimized strategies require monthly rebalancing as parameter estimates are updated.

Turnover is highest for MinVar in three of four datasets (0.072 in ten industry portfolios, 0.395 in twenty-five portfolios, 0.087 in GICS, 0.265 in MSCI), and for MV in the GICS dataset (0.106). The high turnover of MinVar reflects the sensitivity of unconstrained variance-minimising weights to changes in the rolling covariance matrix: with many assets, small changes in the estimated covariance structure translate into large weight changes. MV's turnover is notably lower than MinVar in the twenty-five portfolios (0.010 vs 0.395) and MSCI (0.011 vs 0.265) datasets, suggesting that in high-dimensional settings, expected-return-based constraints on weights actually stabilize the portfolio more than pure variance minimization without constraints.

The relationship between performance and turnover is dataset-dependent. In the ten industry portfolios, 1/N and GenMin share both the lowest turnover and competitive CEQs. In the twenty-five portfolios, MinVar has the highest CEQ and the highest turnover — but the CEQ advantage is large enough to survive transaction costs. In the MSCI dataset, MinVar has the highest turnover and a negative CEQ: it is both the most expensive and the worst-performing strategy, making EWMin the practical choice. This analysis confirms the recommendation of DeMiguel, Garlappi, and Uppal (2009) to account for transaction costs: the gross performance advantage of high-turnover strategies may be partially or fully offset by implementation costs.

6.10 Summary

The empirical results confirm the central mechanism of DeMiguel, Garlappi, and Uppal (2009): the M/N ratio governs relative performance, and strategies that are robust to estimation error, particularly constrained minimum variance variants, are competitive with or superior to naive diversification across all four datasets. The updated sample extends the original findings in

several directions: optimized strategies now more consistently outperform 1/N, particularly in lower-dimensional settings; MV performs well in the GICS dataset where M/N is large and sector momentum provides a useful signal; and EWMin emerges as a particularly robust choice in the international equity setting where the estimation problem is most severe.

One observation that stands out from this analysis is that the "naive diversification puzzle" takes a different form in the updated data than in the original study. The puzzle in 2009 was: why does the simplest possible strategy match or beat everything else? The puzzle in 2025 is: which of the constrained and robust strategies does best, and under what conditions? That shift, from naive vs optimized to good optimization vs poor optimization, reflects both the progress of the field and the more favorable estimation environment of the post-1990 data.

CHAPTER 7 - CONCLUSION

7.1 Summary of the Main Findings

This thesis set out to test whether the central finding of DeMiguel, Garlappi, and Uppal (2009), that the naive equal-weight portfolio is extremely difficult to beat out of sample, survives in a more recent and different sample period. Using data from 1990 to 2025 and four datasets analogous to those of the original study, the analysis produces results that are both consistent with and meaningfully different from the original paper.

The M/N ratio remains the single most informative predictor of relative performance, as in the original study. When the ratio of the estimation window to the number of assets is large (M/N ≈ 10–12), optimized strategies consistently outperform 1/N on both Sharpe ratio and CEQ. When M/N ≈ 5, the estimation environment is more difficult and the results depend more on which specific strategy is used: minimum variance without short-selling constraints dominates in the twenty-five portfolios dataset, while MV and 1/N are roughly equivalent in the MSCI dataset.

Across all four datasets, MV produces a positive CEQ and a Sharpe ratio at least as high as 1/N. In the GICS dataset, MV is the top performer on both metrics, outperforming even the constrained strategies. This is a substantive departure from how the original paper's findings are sometimes interpreted: the problem with MV is not that it destroys value, but that it often fails to add value beyond naive diversification, particularly when the estimation environment is challenging.

The strategy that most clearly stands out in the updated data is unconstrained MinVar. In the twenty-five portfolios dataset, it achieves a Sharpe ratio of 0.953 and CEQ of 10.30% per year, far above all other strategies, while accepting the cost of high turnover (0.395). In the MSCI dataset, however, MinVar is the weakest performer with negative CEQ, illustrating that its advantage is not universal but depends on whether the historical covariance structure provides

useful information. EWMin, which blends minimum variance with equal weighting, provides more consistent performance across all four datasets.

7.2 What Has Changed Since 2009

Comparing the results directly with those of DeMiguel, Garlappi, and Uppal (2009) reveals four main changes.

First, optimized strategies now more consistently beat 1/N in lower-dimensional settings. The original study found no strategy that consistently outperformed 1/N across their datasets. This analysis finds that in the ten industry portfolios, all seven optimized strategies outperform 1/N on both metrics; in the twenty-five portfolios, MinVar, EWMin, and GenMin clearly outperform. This may reflect that the post-1990 data are more structurally stable, characterized by fewer fundamental changes in monetary policy frameworks and market microstructure, making parameter estimates more reliable.

Second, MV performs better than the original study might suggest. While MV does not dominate, it is always competitive with 1/N and occasionally leads. The original paper's framing, that optimization is defeated by estimation risk, holds for the relative comparisons among strategies, but not for the absolute evaluation of MV against the risk-free rate. This distinction matters for practitioners: MV is a reasonable strategy, just not the best one.

Third, the GICS dataset reveals that in a concentrated sector momentum environment, expected-return-based strategies like MV and MV no-short can add genuine value over diversification-based approaches. This finding is specific to the 2015–2026 evaluation window and its extraordinary Technology sector outperformance, and should not be over-generalized. But it confirms the theoretical prediction that when expected return signals are strong and persistent, strategies that incorporate them outperform those that ignore them.

Fourth, EWMin emerges as the most consistently robust strategy across the four datasets. By blending variance minimization with equal weighting, it avoids the extreme sensitivity of unconstrained MinVar while still exploiting covariance information. In the MSCI dataset, where most other optimized strategies either match 1/N or underperform it, EWMin achieves the highest CEQ by a significant margin. This robustness makes it a practically attractive choice across different market environments.

7.3 Implications for Portfolio Management

The results carry practical implications for investors and portfolio managers.

The M/N ratio provides a simple, computable guide to strategy selection. For a portfolio of ten or fewer broadly diversified equity categories with at least ten years of monthly data, constrained optimization, particularly minimum variance variants with weight bounds, is likely to add value over naive diversification. For portfolios with twenty or more positions and shorter data histories, the estimation challenge is more severe, and the choice of strategy matters more: unconstrained MinVar may excel in some environments while failing in others, making EWMin or GenMin safer alternatives.

The turnover results have an important practical message: the strategies that perform best in gross terms do not always perform best in net terms. MinVar's 10.30% gross CEQ advantage over 1/N in the twenty-five portfolios survives transaction costs comfortably. But MinVar's high turnover in the MSCI dataset, combined with its negative gross CEQ, makes it genuinely value-destroying on a net basis. Investors should always verify that a strategy's performance advantage exceeds its implementation cost before adoption.

The strategies which perform most consistently across the four datasets, EWMin, GenMin, MV no-short, are not the most theoretically sophisticated and they do not require advanced estimation methods or complex optimization routines, but they work because they balance two

competing objectives: using available information about the covariance structure while constraining the portfolio enough to avoid the worst consequences of estimation error. That balance, rather than any particular technical approach, appears to be the key to robust portfolio construction.

7.4 Limitations and Future Research

Several limitations deserve acknowledgement. The most important one is the absence of formal statistical testing. The performance differences reported in Chapter 6 are described as economically meaningful based on their magnitude, but no Sharpe ratio tests or bootstrap confidence intervals are reported. For some results, particularly MV's outperformance in the GICS dataset over only 127 months, statistical significance is uncertain.

The fixed estimation window of M = 120 months is a second constraint. The original study explored sensitivity to window length and found results broadly stable. A comparable analysis in the updated sample, particularly examining whether MinVar's dominance in the twenty-five portfolios holds across different window lengths, would be a natural extension.

Finally, this analysis is limited to the eight strategies from the original study. The portfolio optimization literature has continued to develop since 2009, with approaches including machine learning-based return forecasting, high-dimensional shrinkage estimators such as those of Ledoit and Wolf (2004), and factor-model covariance estimation. Evaluating these newer approaches against the 1/N benchmark in the updated sample would provide a more complete picture of the current state of the field.

7.5 Concluding Remarks

The naive diversification puzzle posed by DeMiguel, Garlappi, and Uppal (2009) has not been resolved by the additional years of data. Naive diversification remains a competitive strategy, particularly in high-dimensional settings where estimation error is most severe. But the updated results suggest that the puzzle has evolved: the question is no longer whether any strategy can beat 1/N, but rather which constrained and robust strategies do so most reliably and under what conditions.

The key lesson from this analysis is that estimation risk is not a fixed obstacle but one that varies with the investment setting. When the M/N ratio is large, when data quality is high, and when the asset universe has stable covariance properties, constrained optimization reliably adds value. When these conditions do not hold, the case for naive diversification remains strong. Understanding which regime applies is perhaps the most practically valuable skill a portfolio manager can develop from this line of research.

Looking back at the analysis as a whole, the M/N mechanism appears consistent across four very different asset universes, two different continents, and 25 years of data that include some of the most dramatic market events of recent decades. That the same mathematical relationship continues to govern empirical outcomes, even as market structures, data quality, and the nature of crises have changed, suggests that it reflects something fundamental about the nature of estimation in portfolio construction, rather than a historical artefact of any particular dataset.

The practical recommendation that emerges is pragmatic: the M/N ratio is a simple, computable guide to when optimization can be trusted. When data are plentiful relative to the number of assets, it suggests applying constraints and optimizing, while when they are not, diversifying equally and accepting the limitation.

SOURCES

Markowitz (1952), *Portfolio Selection* (pag 77-91)

Elton, Gruber, Brown & Goetzmann, *Modern Portfolio Theory and Investment Analysis* (cap 1)

Danthine & Donaldson, *Intermediate Financial Theory* (cap 7)

Tobin (1958), *Liquidity Preference as Behavior Towards Risk*

Sharpe (1964), *Capital Asset Prices: A Theory of Market Equilibrium*

Fama & French (1993), *Common Risk Factors in the Returns on Stocks and Bonds*

Carhart (1997), *On Persistence in Mutual Fund Performance*

DeMiguel, Garlappi & Uppal (2009), *Optimal versus Naive Diversification*

Bodie, Kane & Marcus, *Investments* (cap 7, cap 8)

Merton (1980), *On Estimating the Expected Return* (pag 323-361)

Campbell, Lo & MacKinlay, *The Econometrics of Financial Markets*

Ledoit & Wolf (2004), *A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices*

Black–Litterman (1992), *Global Portfolio Optimization*

Jagannathan & Ma (2003), *Risk Reduction in Large Portfolios*