Constructing an Optimized Investment Portfolio:

A Minimum-Variance Optimization Approach

Bachelor Thesis

BSc in Economics and Business Administration

Aarhus University — Aarhus BSS

Author: Jakub Mandak

Supervisor: Charlotte Christiansen

Submission Date: May 1, 2026

Abstract

This thesis tests whether constrained minimum-variance optimization outperforms naive 1/N diversification and the S&P 500 in terms of risk-adjusted returns in-sample. The sample is constructed of 33 large-cap US stocks from January 2005 to December 2025 and split into four sub-periods. Portfolios are constructed using Excel Solver and evaluated on the Sharpe ratio, maximum drawdown, and other return and risk metrics. The results depend on which performance metric is used. Constrained MVO records the lowest maximum drawdown in every sub-period. On the Sharpe ratio, constrained MVO leads over 1/N and the S&P 500 across the full period and in two of four sub-periods, but 1/N is competitive or ahead when market conditions reward broad exposure. The unconstrained version achieves the highest Sharpe ratio when estimated, but requires short positions and is excluded from the two shortest sub-periods. The analysis is done in-sample, so estimation error is not driving the results in the same way it would out-of-sample or in practice. This means the results show the best-case scenario for optimization, not what a real investor would have achieved.

Executive Summary

This thesis asks whether constrained minimum-variance optimization outperforms naive 1/N diversification and the S&P 500 on risk-adjusted returns in-sample, across a 21-year period from January 2005 to December 2025. This question comes from the debate that DeMiguel, Garlappi and Uppal (2009) opened. They showed that 1/N is hard to beat out-of-sample across many optimized strategies. However, that result does not say whether the theoretical advantage of mean-variance optimization from Markowitz (1952) applies when the optimizer is evaluated in-sample, rather than in an out-of-sample setting where estimation error plays a larger role. This is the gap. The thesis runs a direct in-sample comparison to address it.

The sample is 33 large-cap US stocks drawn from the S&P 500, with three stocks per GICS sector across all 11 sectors. Monthly return data comes from CRSP via WRDS. The S&P 500 Total Return Index is from Investing.com, and the analysis compares four portfolios. The first is a constrained MVO portfolio that does not allow short selling. The second is an unconstrained MVO portfolio that allows short positions. The third is a 1/N equally weighted portfolio. The fourth is the S&P 500 as the passive benchmark. Portfolio weights are solved numerically using Excel Solver with the GRG Nonlinear method. All four are evaluated over the full 21-year period and across four economically defined sub-periods: pre-crisis, crisis, post-crisis, and recent.

The results show that it depends on which metric is used. On maximum drawdown, constrained MVO records the lowest peak-to-trough decline in every single sub-period. The gap is widest during the 2008 to 2009 crisis, where constrained MVO had a maximum drawdown of -17.94%, compared with -39.24% for 1/N and -46.42% for the S&P 500. The Sharpe ratio results are more mixed. Constrained MVO leads over 1/N and the S&P 500 across the full period at 1.215, compared to 1.177 for 1/N and 0.764 for the S&P 500, and leads in the crisis and post-crisis sub-periods. However, 1/N edges ahead in the pre-crisis expansion market and in the recent period, where broad exposure to high-growth sectors, especially Information Technology and Communication Services, drove returns. The unconstrained MVO is the strongest Sharpe ratio performer in every sub-period where it can be estimated. Because unconstrained MVO is excluded from the two shortest sub-periods and requires short positions, constrained MVO is the more relevant version for answering the research question and for most investors in practice.

The sector weight analysis explains the drawdown results. During the 2008 to 2009 crisis, constrained MVO put 71.6% into Utilities and Consumer Staples, the two most defensive sectors. In the post-crisis expansion, weight rotated to Consumer Discretionary and Health Care. In the recent period, Communication Services rose to 23.9%. The pattern is the same in each window. The optimizer concentrates in whichever sectors had the lowest realized variance at that point. So the portfolio rotates with the regime instead of holding a fixed defensive tilt, and that rotation is what produces the consistent drawdown advantage.

In this sample, constrained MVO is strongest on downside protection, but less consistent on risk-adjusted return. If an investor wants to limit the worst loss, constrained MVO is the better option in this sample. An investor who cares more about raw return, or about capturing upside during growth-driven rallies, may find 1/N simpler and stronger in some regimes. Since the analysis is done in-sample, it reduces the role of estimation error compared to an out-of-sample test. Because of that, the results should be interpreted as the best-case version of what minimum-variance optimization can achieve. Whether this would still work out-of-sample is something this thesis cannot answer. This would also be the recommended direction for future research.

Table of Contents

1. Introduction

In today's investing environment, portfolio theory and what investors do in practice have never fully lined up. The Markowitz (1952) framework can help investors provide a clear theoretical answer on how to allocate capital across risky assets. It helps investors to minimize variance for a given level of return by analyzing how assets move together and then holding a portfolio that sits on the efficient frontier. However, what investors do in practice most of the time looks very different, and the gap between that has grown bigger over time. The recent market environment has especially highlighted this. Between 2020 and 2025, we saw a COVID-driven equity crash, a sharp monetary response, a supply-chain inflation shock, and an AI-driven rally. Interest rates also moved from near zero to close to 5%. In that kind of environment, it becomes more important to ask whether optimization improves portfolio outcomes or whether a simple passive strategy is enough.

At the center of the literature is the comparison between mean-variance optimization and naive diversification. On one side is Markowitz (1952). On the other side is 1/N, which is a rule that allocates equal weight to every stock and ignores other metrics. The main reference for the 1/N position is DeMiguel, Garlappi and Uppal (2009). They compared 1/N against 14 optimized strategies across seven datasets and concluded that none of them consistently outperformed the naive rule out-of-sample. This finding creates a tension when compared with the theoretical prediction from Markowitz. If mean-variance optimization is supposed to improve risk-adjusted returns, why does the simple 1/N rule perform just as well or better in practice? Michaud (1989) explains this through estimation error. Small errors in the covariance matrix or expected return estimates can cause large swings in the assigned weights. Because of that, the optimizer may end up reacting more to noise in the data than to real underlying patterns.

DeMiguel, Garlappi and Uppal (2009) show that 1/N is competitive out-of-sample, which is where estimation error has the biggest impact. Michaud (1989) explains why. But neither paper isolates what minimum-variance optimization looks like when it is evaluated in-sample rather than out-of-sample. In other words, what does MVO look like at its theoretical best? If MVO fails even in-sample, the problem likely goes beyond estimation error. It may come down to the chosen assets, the imposed constraints, or the objective function itself. But if MVO outperforms in-sample, then the result helps illustrate how much of MVO's theoretical advantage may be weakened once estimation error becomes more important. This is the gap that this thesis addresses.

The analysis tests whether the constrained minimum-variance optimized portfolio can outperform naive 1/N diversification and the S&P 500. The thesis assesses their in-sample risk-adjusted returns alongside other metrics across 33 large-cap US stocks from January 2005 to December 2025. The sample covers a 21-year window that includes two complete business cycles, the 2008 global financial crisis, the longest US expansion on record, and many regime shifts of the COVID period. The sample is constructed from the S&P 500, with three stocks from each of the 11 GICS sectors. This links the 33-stock choice to the diversification threshold from Statman (1987) while also covering the full sector structure of the US equity market. The analysis is then run over the full period and across four economically defined sub-periods. This is important because it allows the ranking to be tested across different market regimes, rather than being assessed on a single full-sample number.

The rest of the thesis covers the problem statement, literature, data, methodology, full-period results, sub-period analysis, discussion, and conclusion, in that order.

2. Problem Statement

The Markowitz (1952) framework implies that a mean-variance optimized portfolio should outperform a naive 1/N portfolio on a risk-adjusted basis, since 1/N ignores the covariance structure of returns and will generally not lie on the efficient frontier. The DeMiguel, Garlappi and Uppal (2009) result shows that this implication fails out-of-sample across a wide range of datasets and optimized strategies. But that result says nothing about whether MVO's theoretical advantage shows up when estimation error is less of a factor, as it is in an in-sample setting. This thesis addresses that gap with a direct in-sample comparison.

The research question is:

> *Does constrained minimum-variance optimization outperform naive 1/N diversification and the S&P 500 benchmark on risk-adjusted returns in-sample across a 21-year period from January 2005 to December 2025?*

Three specific comparisons help answer this question. The first is constrained MVO versus 1/N. This directly tests whether the Markowitz framework produces a measurable advantage over naive diversification. The second is constrained MVO versus the S&P 500, which shows whether any optimized construction beats the most common passive benchmark on a risk-adjusted basis. The third comparison breaks the 21-year sample into four sub-periods: pre-crisis, crisis, post-crisis, and recent. The point of this is to check whether the ranking holds across different market regimes or whether the full-period result is mainly driven by specific conditions.

The sample is 33 large-cap US stocks selected from the S&P 500, with three stocks per GICS sector across all 11 sectors, covering monthly returns from January 2005 to December 2025. The sample construction, data sources, and sub-period definitions are set out in Chapter 4 and Section 5.4. The minimum-variance optimization procedure is described in Chapter 5.

The analysis is done in-sample. This makes it possible to see what minimum-variance optimization delivers when it is evaluated on the same data on which it is estimated. Because of that, the result should be read as a best-case version of what the framework can produce. The next step would be to run an out-of-sample test in the style of DeMiguel, Garlappi and Uppal (2009). That is outside the scope of this thesis, but Section 8.3 discusses it as the main direction for future research.

3. Literature Review

3.1 Modern Portfolio Theory and Mean-Variance Optimization

Modern portfolio theory starts with Markowitz (1952). Before Markowitz, the standard way of thinking about investment was to pick individual stocks expected to deliver high returns. Markowitz reframed the problem. Instead of looking at a single stock in isolation, the idea is that what matters is how the full portfolio behaves. How the portfolio behaves depends on two things. The expected return of each asset and the covariance between them. This means the problem is no longer just about picking good individual stocks, but about how those stocks work together in a portfolio.

What Markowitz showed is that portfolio variance is not just the average of individual stock variances. It depends heavily on how stocks move together. Two stocks with identical standard deviations can produce very different portfolio risk depending on their correlation. If they move in the same direction, combining them does little to reduce risk. If they move in opposite directions, combining them can reduce risk significantly, even when each stock is individually volatile. So, diversification is not just about owning many stocks. It is more about owning stocks whose returns do not all move together.

From this came the efficient frontier. For any given level of expected return, there is one portfolio that achieves that return with the minimum variance. If you plot these portfolios across different return levels, you get a curve. Portfolios on that curve are efficient. Anything below it is dominated, because another portfolio exists with the same risk but a higher return, or the same return but lower risk. So a rational investor would hold a portfolio on the frontier. Which one depends on how much risk the investor is willing to accept.

The minimum-variance portfolio sits at the tip of the efficient frontier. It carries the lowest variance achievable across all possible weight combinations, regardless of expected return. It can also be constructed using only the covariance matrix, without needing expected return estimates. The reason this is useful is that expected returns are hard to estimate from historical data. Covariances are much more stable. So, the minimum-variance portfolio only relies on the covariance matrix, which is estimated from historical data. Since it does not require expected return estimates, it is less sensitive to estimation error, which is why it is used in this thesis. Both the constrained and unconstrained MVO portfolios in Chapter 5 target this solution, with the constrained version adding the non-negativity restriction that rules out short positions.

3.2 The Capital Asset Pricing Model and the Sharpe Ratio

Sharpe (1964) extended the mean-variance theory into the Capital Asset Pricing Model. CAPM builds on Markowitz (1952) but adds a risk-free asset. The idea is that investors combine the market portfolio with the risk-free asset rather than holding only risky assets. From that, the model says that the expected return of a stock depends on its systematic risk relative to the market. This risk is measured by beta. Lintner (1965) and Mossin (1966) derived the same result independently.

CAPM itself is not used directly in this thesis since the research question is about portfolio construction, not asset pricing. It is included here because it leads to the Sharpe ratio, which is the main performance measure used throughout. The same author, Sharpe (1966), introduced the ratio as a practical tool for comparing portfolios on a risk-adjusted basis. It is defined as excess return divided by portfolio volatility, as shown in Section 5.3. That is why it is used to compare MVO, 1/N, and the S&P 500 in Chapters 6 and 7.

3.3 Diversification: How Many Stocks Are Enough?

After Markowitz established the diversification principle, the next question was clear. How many stocks does an investor need? In practice, the gains from adding more stocks decline quickly. So, beyond some point, adding more holdings mostly adds complexity without reducing much risk.

Evans and Archer (1968) were the first to analyze this. Using US stock data, they argued that most of the diversification benefit comes by the time a portfolio reaches around 10 to 15 randomly selected stocks. Beyond that, adding more stocks produces only small reductions in variance. This then became the common view in the literature and finance textbooks.

Statman (1987) challenged that conclusion. He argued that the Evans and Archer threshold is too low because it only considers variance reduction and ignores the cost-benefit trade-off an investor faces. Statman's result is that the minimum well-diversified portfolio is around 30 stocks. The 10-stock answer from the earlier literature comes from ignoring the marginal cost of holding more stocks. Once both the marginal benefit and the marginal cost are accounted for, the answer shifts to around 30. Bodie, Kane and Marcus (2023) use the same number in their textbook treatment of diversification, which reinforces this conclusion until now.

The 33-stock sample used in this thesis is based on Statman's (1987) threshold of 30 stocks. It is slightly above the minimum, and the way it is structured also connects to the broader diversification logic from Markowitz (1952). Three stocks are selected from each of the 11 GICS sectors, so the portfolio is diversified both in the number of holdings and across industries whose returns do not move in perfect correlation. So, the sample size is based on theory rather than chosen randomly. The number comes from one of the core papers on portfolio size, and the sector structure follows directly from Markowitz. The specific stock selection procedure is described in Section 4.2.

3.4 Critiques of Mean-Variance Optimization

In practice, the Markowitz framework has been criticized for being highly sensitive to estimation error. The optimizer uses inputs such as expected returns and covariances as if they were exact, but they are only estimates. Michaud (1989) called this the Markowitz optimization enigma. The enigma is that mean-variance optimization is mathematically the right answer to the portfolio problem, but when it is applied to real data, small errors in the inputs can produce extreme and economically unreasonable weights. So the optimizer does not only find the minimum-variance portfolio. It can also amplify errors in the inputs. In that sense, it may treat noise as if it were a signal.

If estimation error is the problem, then one way to deal with it is to restrict how strongly the optimizer can respond to noisy inputs. The non-negativity constraint does exactly this. By ruling out short positions, the constrained MVO portfolio limits the optimizer's ability to take extreme leveraged positions based on unstable estimates. Jagannathan and Ma (2003) use this point and show that, in practice, a non-negativity constraint can have an effect like shrinkage on the covariance matrix. This is why the constrained version is used as the primary MVO portfolio in this thesis, and why the unconstrained version becomes unusable when the estimation window is too short.

Chan, Karceski and Lakonishok (1999) look at the same issue from the empirical side. They examine how reliable covariance matrix estimates are in practice and argue that at least 60 monthly observations are needed for the sample covariance matrix to be stable enough for optimization. Below that point, the matrix becomes poorly conditioned, and the optimized weights become unreliable. Their 60-month threshold is one of the reasons the unconstrained MVO is excluded from the pre-crisis sub-period with 36 observations and the crisis sub-period with 24 observations. Together with Michaud (1989), their work explains why optimization can still fail when the sample is too short. Jagannathan and Ma (2003) explain why the non-negativity constraint helps limit that failure.

3.5 Naive Diversification and the 1/N Debate

Michaud (1989) and Chan, Karceski and Lakonishok (1999) show that estimation error is a serious problem for optimization. So the next question is whether there is something simpler that works just as well and is not sensitive to noisy inputs. The simplest answer to that is the 1/N portfolio. Each stock gets an equal weight of 1/N, with no optimization, no covariance matrix, and no expected return estimates. The only input is the list of stocks in the sample.

DeMiguel, Garlappi and Uppal (2009) are the main reference that brought this question into the center of the literature. They compare 1/N against 14 optimized strategies across seven different datasets, including US industry portfolios, international portfolios, and Fama-French factor portfolios. The optimized strategies include the mean-variance portfolio, the minimum-variance portfolio, Bayesian approaches, and different shrinkage-based variants. What they find is that no optimized strategy consistently beats 1/N out-of-sample on the Sharpe ratio, certainty equivalent return, or turnover. So, a rule that ignores covariances completely ends up outperforming strategies that spend a great deal of effort trying to estimate them.

The explanation for this goes back to the same problem Michaud discussed earlier. The optimizer depends on estimated inputs, especially expected returns and the covariance matrix. When those estimates contain noise, the optimized portfolio can end up being based on that noise. The 1/N portfolio avoids this problem because it does not estimate anything. It simply gives each stock the same weight. DeMiguel, Garlappi and Uppal (2009) show how serious this problem can be. To reliably beat 1/N, the optimized strategy would need around 3,000 months of data with 25 assets and around 6,000 months with 50 assets. Those sample sizes are not realistic. So, in practice, 1/N is very hard to beat.

This paper has become the standard counterargument to mean-variance optimization in the academic literature. Because of that, this thesis takes 1/N seriously as a benchmark rather than as a simple baseline. The 1/N portfolio is included alongside the S&P 500 because any claim that MVO delivers useful performance needs to hold up against both, not just against the passive index.

The DeMiguel, Garlappi and Uppal (2009) result is strictly an out-of-sample finding. The optimizer estimates weights using one window, and the portfolio is evaluated on the next one, which is exactly where estimation error matters most. This thesis asks a different question. What does MVO look like when it is evaluated on the same data it is estimated on? That reduces the role of estimation error but does not remove it completely. What it can show is whether MVO's theoretical advantage appears at all when estimation error is the least of a problem.

4. Data

4.1 Data Source and Sample Period

The source of the monthly stock and return data is from the Center for Research in Security Prices (CRSP). Access to this database was provided via the Wharton Research Data Services (WRDS) platform. CRSP is a provider of one of the most comprehensive and widely used databases of historical US equity returns. This is why it was a reliable option for this thesis. It allowed access to the stock's monthly returns, which were constructed from dividend-adjusted closing prices. Each of the 33 sample stocks needed a variable MthRet to be extracted over the full sample period. In the end, it resulted in a dataset of 8,316 observations. In other words, 33 stocks multiplied by 252 monthly returns.

The data is limited to US stocks for a few reasons. First, the United States represents the largest and most liquid equity market. This means that the stock prices are more likely to reflect all available information more efficiently. Second, much of the literature from Chapter 3 was grounded in US data. By using the same market, the findings can be positioned within that literature. Third, the S&P 500 index, which is used as a natural and accessible benchmark, is listed in the United States. It represents a realistic alternative to any US equity investor, and limiting the analyzed stocks to the same market allows for the findings to be more comparable.

The chosen sample period runs from January 2005 to December 2025. This gives a clean 21-year window that spans two complete US business cycles, the 2008 global financial crisis, the long post-crisis expansion, and the COVID-19 shock of 2020. The start date is also chosen to sit comfortably after the 2002 Sarbanes-Oxley Act, so the sample is drawn entirely from the post-SOX reporting regime, where disclosure quality and accounting standards are consistent (U.S. Congress, 2002). Portfolio optimization literature uses monthly frequency as a standard practice, and for that, it was chosen as well. Monthly data also provides enough observations for reliable covariance matrix estimation, and it avoids the microstructure noise associated with daily returns.

As mentioned briefly above, the S&P 500 Total Return Index is used as the passive benchmark throughout the analysis. The S&P 500 represents the most widely recognized measure of broad US equity market performance. On top of that, this index can be purchased by anyone at negligible cost. This is what makes it a natural opportunity cost for any active strategy, as discussed by Bodie, Kane and Marcus (2023). It is used in the form of total return data, which accounts for dividend reinvestment in addition to the price increase. The dataset is sourced from Investing.com, as CRSP coverage did not extend to December 2025 at the time of data collection. It covers the full sample period from January 2005 to December 2025.

4.2 Stock Selection Procedure

The stock selection follows a rule-based procedure to keep it objective. It is important to ensure that the data is available across the full sample period. The 11 sectors of the Global Industry Classification Standard (GICS) form the basis for the portfolios. Sectors are Energy, Materials, Industrials, Consumer Discretionary, Consumer Staples, Health Care, Financials, Information Technology, Communication Services, Real Estate, and Utilities. GICS is an industry-developed tool by MSCI and S&P Dow Jones Indices in 1999. This classification provides a global framework for the classification of different companies based on their business activity (MSCI and S&P Dow Jones Indices, 1999). Currently, there are 11 sectors that include Communication Services, which was introduced in September 2018. This was due to an expansion of existing Telecommunication Services (S&P Dow Jones Indices, 2018). Three stocks were selected from each of the 11 sectors. Three stocks per sector gives 33 in total, which sits just above Statman's (1987) minimum of 30. It also balances two practical considerations. Two per sector leaves each sector under-diversified, so a single stock shock would hit the sector allocation too hard. Five per sector would produce 55 assets, which makes the covariance matrix harder to estimate reliably, especially in the shorter sub-periods. Three keeps the portfolio diversified within each sector while staying tractable for Solver across all sub-periods.

Each of the 11 sectors was assigned the three largest companies by market capitalization. This was done using the S&P 500 holdings ranking published by Slickcharts (2026), accessed in March 2026. The reason for choosing stocks with the largest market capitalization is that they are generally more liquid and investable. This is what helps with the reduction of noisy price data (Bodie, Kane and Marcus, 2023). On top of that, large-cap stocks are more likely to provide continuous historical returns across the full sample period. This is important because it helps to eliminate missing observations, which would complicate the construction of the covariance matrix. Lastly, this selection process provides a transparent and objective rule. Slickcharts ranks S&P 500 holdings by index weight, but it does not include each stock's GICS sector. To resolve that, each company's sector classification was verified manually by comparing it with the official MSCI and S&P Dow Jones Industries sector definitions (MSCI and S&P Dow Jones Indices, 1999). Because the ranking was pulled in March 2026 and the sample ends in December 2025, there is a three-month gap. Given the stability of large-cap stocks, this is unlikely to have a significant effect on the sector composition.

A data completeness requirement is applied to each stock in the dataset. Each selected stock was required to have continuous monthly return data. Any stock whose return did not extend continuously from January 2005 to December 2025 was excluded. Excluded stocks were replaced by the next largest company within the same sector. This approach ensured that the final sample consisted of stocks with 252 consecutive monthly returns.

Due to this requirement, the data resulted in several substitutions. Meta Platforms (META) was excluded as it did not go public until May 2012 and was replaced by Netflix (NFLX). Netflix was the next largest company in the Communication Services sector. Tesla (TSLA) was excluded due to its June 2010 IPO and replaced by McDonald's (MCD). This was the only change in Consumer Discretionary. Within the Financials sector, both Visa (V) and Mastercard (MA) were excluded due to their March 2008 and May 2006 IPOs. Because of that, Bank of America (BAC) was chosen as the 5th largest qualifying company to fill the third position in the sector. In the Communication Services sector, T-Mobile (TMUS) was excluded as the corporate entity only dates to 2013 and was replaced by Verizon (VZ). Lastly, in the Materials sector, Linde (LIN) was excluded as the current entity was created through a merger in 2018 and replaced by Sherwin-Williams (SHW). In all the above cases, the replacement stock was always the next largest qualifying company within the same sector with a continuous return history from January 2005. The full 33-stock sample with all substitution notes is shown in Table A1 in the Appendix.

This selection procedure introduces survivorship bias. All 33 selected companies are large enough in March 2026 to rank near the top of their sector. This implies that they survived and grew over the full sample period. Brown et al. (1992) explain survivorship bias, and this limitation applies here. The results should be interpreted as describing the performance of large, established, surviving US firms. Most importantly, it may not generalize to the wider US equity market.

5. Methodology

5.1 Portfolio Construction

This thesis constructs three portfolios, with the fourth being the S&P 500 used as a benchmark. Constructed portfolios are used to evaluate the performance of minimum-variance optimization against simpler alternatives. The three constructed portfolios are a constrained MVO portfolio, an unconstrained MVO portfolio, and a naive 1/N portfolio. These four portfolios (including the S&P 500) then help to cover the spectrum from mathematically optimized to completely passive. In other words, the four portfolios allow comparison across different levels of complexity.

The constrained MVO portfolio is constructed using the Markowitz (1952) mean-variance optimization framework. The portfolio minimizes total portfolio variance subject to two constraints. All weights must sum to one, and no individual weight can be negative. The non-negativity constraint means that short selling is not allowed, which reflects the practical limitations faced by most investors. The unconstrained MVO portfolio follows the same optimization procedure but allows individual weights to become negative, which implies short positions. The third portfolio is the 1/N equally weighted portfolio, where each of the 33 stocks receives an identical weight of 1/33. This equals approximately 3.03%. This portfolio is the primary naive benchmark. Lastly, the S&P 500 Total Return index is the passive market benchmark, as introduced in Section 4.1.

The two MVO portfolios test whether mathematical optimization delivers superior risk-adjusted performance compared to simple alternatives. Comparing the constrained and unconstrained MVO portfolios directly shows what difference the no-short-selling constraint makes on performance.

5.2 Mean-Variance Optimization Framework

Both constrained and unconstrained MVO portfolios are focused on the minimum variance. The objective is to find the set of portfolio weights that produces the lowest possible variance. It is calculated using the minimum-variance framework introduced by Markowitz (1952). The idea is that a rational investor cares not only about return but also about variance.

The expected return of a portfolio is the weighted sum of the individual asset returns:

$$R_p = \sum_{i=1}^{N} w_i R_i \quad (1)$$

Portfolio return is used to evaluate performance after optimization, not as an input to the optimizer.

Formally, the constrained MVO problem is written as:

$$\min_{w} w^{\mathrm{T}} \Sigma w \ (2)$$

subject to

$$w^{\mathrm{T}} 1 = 1$$

$$w_i \geq 0$$

The portfolio standard deviation, used as the measure of risk, is the square root of the portfolio variance:

$$\sigma_p = \sqrt{w^{\mathrm{T}} \Sigma w} (3)$$

Here, w is the 33-by-1 vector of portfolio weights, and Σ is the 33-by-33 annualized covariance matrix of monthly returns. The objective function $w^T \Sigma w$ is the portfolio variance. The first constraint requires the weights to sum to one, meaning the portfolio is fully invested. The second constraint rules out short selling by requiring each weight to be non-negative. The unconstrained MVO portfolio uses the same objective and the same full investment constraint but drops the non-negativity requirement. This means individual weights can take negative values. This thesis implements the global minimum-variance version of MVO, following Jagannathan and Ma (2003), because expected return estimates introduce additional noise without improving the optimization outcome.

The optimization requires a covariance matrix that captures how the returns of each stock move together. For that, a 33x33 covariance matrix is constructed from monthly return data. In Excel, this is calculated using the COVARIANCE.S function, producing monthly covariances that are multiplied by 12 to annualize them, since variance scales linearly with time under the standard assumption that monthly returns are independently and identically distributed. With 252 monthly observations, the covariance matrix exceeds the 60-observation minimum recommended by Chan, Karceski and Lakonishok (1999) for stable matrix estimation. When the number of observations falls below the number of assets, the matrix becomes rank-deficient. That means the usual inverse-based solution does not work, and the unconstrained weights become unreliable. When observations only slightly exceed the number of assets, the matrix is technically invertible but near-singular, which means small estimation errors in the inputs get amplified into large swings in the optimized weights, as Michaud (1989) describes. Both cases are relevant for the shorter sub-periods analyzed in Section 5.4.

Excel's Solver add-in is then used to find the optimal portfolio weights. Solver uses the Generalized Reduced Gradient (GRG) Nonlinear method to numerically solve this quadratic program. Since the objective function is quadratic and the constraints are linear, the problem is convex when the covariance matrix is positive semi-definite. In this setting, Solver's GRG Nonlinear method is expected to converge to the global minimum. It searches across all possible weight combinations to find the one that minimizes total portfolio variance. The objective cell is set to the portfolio variance formula, and Solver then adjusts the 33 individual stock weights to minimize this value. This process is repeated separately for each period analyzed.

The two MVO portfolios differ only in their weight constraints. The constrained MVO portfolio requires all weights to be non-negative, while the unconstrained version allows weights to take any value. This distinction becomes relevant in the sub-period analysis, where shorter estimation windows make the unconstrained MVO portfolio more sensitive to noise in the data.

The analysis assumes a risk-free rate of zero. DeMiguel, Garlappi and Uppal (2009) apply the same assumption in their comparison of optimized and naive portfolios. As this thesis directly tests their findings in-sample, using the same assumption keeps the results consistent with the existing literature. Because the same risk-free rate assumption is applied to all portfolios, the comparison remains internally consistent.

Using a time-varying rate, such as the 3-month Treasury bill rate, would be a valid robustness check, but it is unlikely to change the main interpretation.

It is worth acknowledging that the actual US short-term interest rate moved significantly over the 2005 to 2025 sample. It reached close to 5% in 2007, fell to near zero between 2010 and 2015, and returned close to 5% in 2023 and 2024. So, a constant-zero assumption is a simplification, and the reported Sharpe ratios should be read as directional rather than as absolute risk-adjusted returns an investor would have earned in real time.

5.3 Performance Metrics

Six metrics are used to evaluate each portfolio across all time periods. Annualized return, annualized volatility, Sharpe ratio, maximum drawdown, best month, and worst month. Together, these cover both average performance and the more extreme ends of the return distribution.

Annualized return is the primary measure used for measuring the portfolio's profitability. It is calculated as the average monthly return multiplied by 12. This is the arithmetic annualization method used by DeMiguel, Garlappi and Uppal (2009). Geometric annualization would produce lower return estimates since it accounts for compounding, but the arithmetic method is used here to keep the results directly comparable with the existing literature. Annualized volatility measures the dispersion of monthly returns around the mean. It is calculated by using the standard deviation of monthly returns multiplied by the square root of 12. Bodie, Kane and Marcus (2023) describe annualized volatility as the standard measure of portfolio risk.

The Sharpe ratio measures risk-adjusted return by dividing excess return over the risk-free rate by portfolio volatility. With a risk-free rate of zero, as established in Section 5.2, this simplifies to annualized return divided by annualized volatility. The ratio is included because two portfolios with identical returns can carry very different levels of risk, and raw return alone does not capture that difference. The formula, following Sharpe (1966), is:

$$SR = \frac{R_p - R_f}{\sigma_p} \quad (4)$$

where R_p is the annualized portfolio return, R_f is the risk-free rate, and σ_p is the annualized portfolio volatility. In this thesis, $R_f = 0$, following DeMiguel, Garlappi and Uppal (2009).

Maximum drawdown measures the largest peak-to-trough decline in portfolio value over the sample period. It is defined as:

$$\mathrm{MDD} = \min_t \left(\frac{V_t - \max_{\tau \le t} V_\tau}{\max_{\tau \le t} V_\tau} \right) \quad (5)$$

It is calculated from the cumulative drawdown sequence, which records the decline from the previous portfolio peak at each point in time. Maximum drawdown shows the worst sustained loss, which is a standard measure of downside risk in portfolio evaluation (Bodie, Kane and Marcus, 2023).

The best month and worst month are the maximum and minimum single monthly returns recorded by each portfolio. They complete the picture that the Sharpe ratio and maximum drawdown leave open. These metrics show the single sharpest movements in either direction, which gives a sense of how each portfolio behaves in extreme monthly conditions.

5.4 Sub-Period Structure

The full-period analysis alone cannot show how each portfolio behaves under different market conditions. Sub-period analysis addresses this by breaking the full sample into four economically distinct periods. This is what makes it possible to test a portfolio's performance across different market regimes. Each sub-period is defined using official US business cycle dates published by the National Bureau of Economic Research Business Cycle Dating Committee (NBER BCDC).

The first sub-period is the pre-crisis period, running from January 2005 to December 2007. It covers 36 monthly observations. NBER BCDC (2008) officially determined that the US business cycle reached its peak in December 2007, so the full window falls within the expansion. With only 36 observations, the covariance matrix is near singular and would produce unreliable weights for the unconstrained MVO portfolio. It is therefore excluded from this period. Even though the 60-month threshold from Chan, Karceski and Lakonishok (1999) is not met, the constrained MVO remains usable since the non-negativity constraint limits the optimizer's sensitivity to the poorly conditioned matrix. Results are reported for the constrained MVO, 1/N, and S&P 500 only.

The second period runs from January 2008 to December 2009, described as a crisis. It covers 24 monthly observations. NBER BCDC (2008) officially determined that the US recession began in January 2008, and NBER BCDC (2010) confirmed that it ended in June 2009. The period end is set at December 2009 rather than June 2009 to maintain a clean calendar year cut. With only 24 observations for 33 assets, the covariance matrix is rank-deficient, since there are fewer observations than assets. For this reason, the unconstrained MVO portfolio is excluded from the crisis period analysis. The constrained MVO remains usable even at 24 observations because the non-negativity constraint prevents the kind of extreme leveraged positions that make unconstrained optimization unstable at short horizons (Jagannathan and Ma, 2003). As in the pre-crisis period, constrained optimization is reported despite not meeting the 60-month threshold. The constrained MVO, 1/N, and S&P 500 are evaluated as normal.

The third period, post-crisis, runs from January 2010 to December 2019. It covers 120 monthly observations. NBER BCDC (2010) confirmed that the recession ended in June 2009, marking the start of the post-crisis expansion. NBER BCDC (2020) confirmed that this expansion lasted until February 2020. While global events such as the 2011 Eurozone sovereign debt crisis and the 2015 to 2016 China slowdown caused volatility spikes in US equity markets, neither triggered a US recession, and both are captured in the return data within this sub-period. The post-crisis period is set to end in December 2019 rather than February 2020 to maintain a clean calendar year boundary and to keep it economically distinct from the period that follows. With 120 observations, this is the longest and best-conditioned sub-period in the analysis, comfortably exceeding the number of assets. It also exceeds the 60-month minimum recommended by Chan, Karceski and Lakonishok (1999). All four portfolios are analyzed in this period.

The recent period runs from January 2020 to December 2025. It covers 72 monthly observations. NBER BCDC (2020) determined that the US business cycle peaked in February 2020, marking the start of the COVID-19 recession. January 2020 is used as the period start to maintain a clean calendar year, consistent with the approach applied to the other sub-periods. This timeframe contains multiple distinct economic shocks. It begins with the COVID-19 market crash in early 2020, followed by a sharp recovery. Next is the inflation shock in 2021 and 2022, and an AI-driven equity rally at the end of the period. With 72 observations, all four portfolios are analyzed in this period.

6. Full Period Analysis (Jan 2005 - Dec 2025)

Table 1: Portfolio Performance Metrics, January 2005 – December 2025.

Metric	Constrained MVO	Unconstrained MVO	1/N Equal Weight	S&P 500
Annualised Return	12.67%	13.04%	16.80%	11.32%
Annualised Volatility	10.43%	9.68%	14.27%	14.83%
Sharpe Ratio	1.215	1.347	1.177	0.764
Max Drawdown	-21.22%	-14.16%	-41.31%	-50.95%
Best Month	+9.31%	+8.29%	+12.29%	+12.82%
Worst Month	-9.05%	-7.70%	-15.53%	-16.80%

Note: The risk-free rate is set to zero. Sharpe ratio is calculated as annualised return divided by annualised volatility. Max drawdown is the largest peak-to-trough decline over the sample period.

Figure 3: Efficient Frontier and Portfolio Positions, January 2005 – December 2025

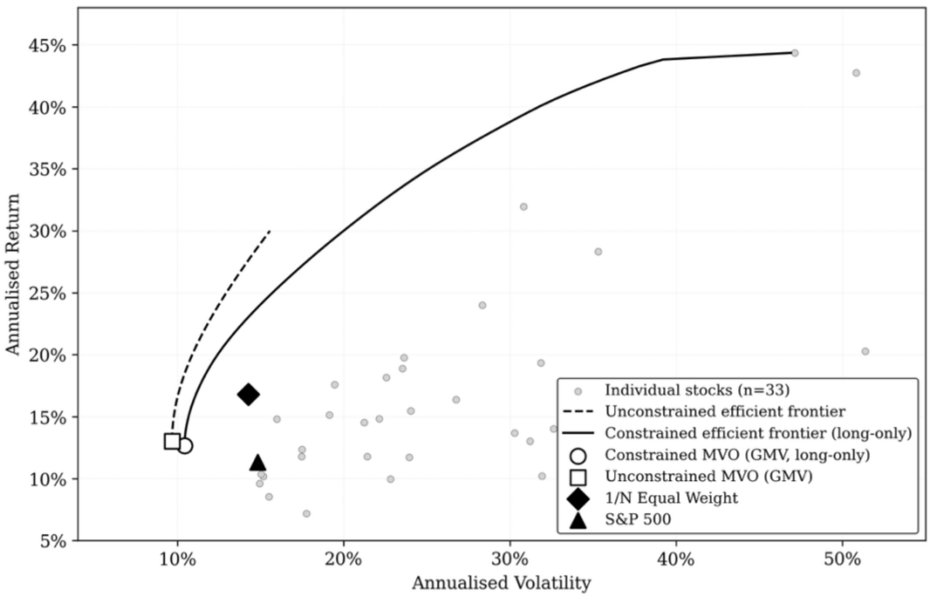

Figure 3 plots all four portfolios and 33 individual stocks against the efficient frontier. Both MVO portfolios sit at the left edge of the frontier, confirming that the optimizer finds the minimum-variance solution. The 1/N portfolio and the S&P 500 both sit well inside the frontier, with higher volatility for their level of return.

Table 1 shows a clear comparison of the different performance metrics over the full sample period. The main finding is that the two MVO portfolios produce the highest risk-adjusted return. The 1/N portfolio generates the highest raw return of the four at high volatility. The S&P 500 has the lowest return and the highest volatility of the four portfolios.

The 1/N portfolio generates the highest annualized return, which amounts to 16.80%, and annualized volatility reaching 14.27%. The constrained MVO returns 12.67% at a volatility of 10.43%, and the unconstrained MVO returns 13.04% at 9.68%. To put that into perspective, MVO portfolios give up some of the return compared to 1/N, but the benefit of that is significantly lower volatility. This is what some investors might prefer, since it delivers lower risk. The S&P 500 sits at the worst position of the four. It returns 11.32% with 14.83% volatility. It is the worst performer because it reaches the lowest return, while taking on the most risk out of them all.

The reason 1/N produces the highest raw return is structural to how it is built. 1/N allocates the same weight to every sector regardless of volatility or historical performance. So, high-return sectors like Information Technology and Consumer Discretionary get the same 3.03% in each stock as defensive sectors like Utilities. Over a 21-year period that includes the post-2009 recovery and the AI-driven rally of 2023 to 2025, those high-growth sectors contributed most of the upside. The 1/N captures that upside in full. The optimizer does not. It deliberately underweights high-variance sectors even when they are performing well, which is why constrained MVO returns 12.67% against 16.80% for 1/N. That 4.13 percentage point gap is the return cost of variance minimization. Whether that cost is worth paying depends on how the volatility side of the trade-off plays out, which is where the Sharpe ratio and drawdown results come in.

The Sharpe ratio makes this difference more visible. Despite producing the highest raw return, 1/N finishes third on risk-adjusted performance at 1.177, behind the unconstrained MVO at 1.347 and constrained MVO at 1.215. The extra return 1/N earns does not compensate for the additional volatility it carries. The S&P 500 trails all three constructed portfolios at 0.764. This in-sample result contrasts with DeMiguel, Garlappi and Uppal (2009), who found that 1/N often outperformed optimized strategies out-of-sample on risk-adjusted measures.

The volatility results are not accidental. Minimum-variance optimization, by construction, minimizes portfolio variance subject to the weights summing to one. So, MVO having the lowest volatility is the direct outcome of the objective function defined in Section 5.2. The constrained MVO records 10.43% annualized volatility, and the unconstrained 9.68%. Both are below 1/N at 14.27% and the S&P 500 at 14.83%. This also explains the pattern in Table 8, where the constrained MVO concentrates heavily in low-volatility sectors like Consumer Staples and Utilities. The full sector-weight analysis is presented in Section 7.6, but the link between the objective function and sector tilt is what drives the whole volatility and drawdown result.

Maximum drawdown points in the same direction. The constrained MVO has a maximum drawdown of -21.22% over the full period, and the unconstrained MVO is even better at -14.16%. The 1/N portfolio drops to -41.31% and the S&P 500 to -50.95%. Compared to the

S&P 500, the MVO portfolios lost roughly half or less at their worst point. The 1/N portfolio again shows the same pattern as in the return and Sharpe ratio results. It takes on more downside risk than the MVO portfolios.

The best and worst months confirm the same pattern at the individual month level. Both MVO portfolios stay within a closer range than 1/N and the S&P 500, whose worst single months exceed -15%. The optimizer limits not just cumulative drawdown but also the sharpest monthly swings.

Figure 1. Cumulative Growth of $1 Invested, January 2005 to December 2025

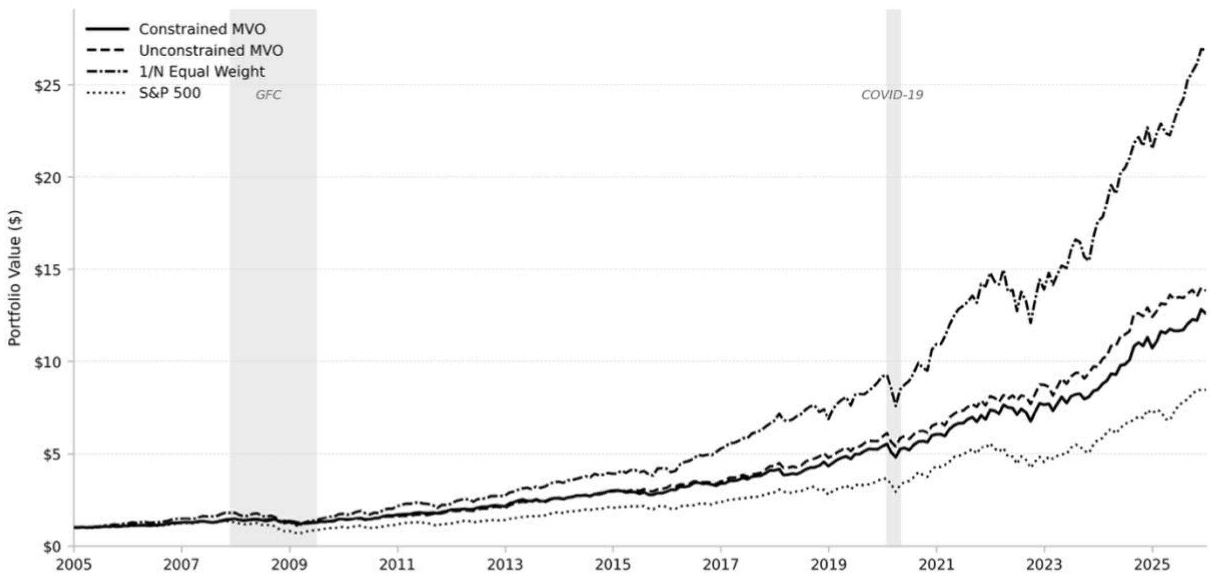

Note: Growth of $1 invested at the start of January 2005. Returns are arithmetic monthly returns compounded over time. Shaded areas indicate the Global Financial Crisis (December 2007–June 2009) and the COVID-19 market drawdown (February–April 2020).

The cumulative chart in Figure 1 helps to unveil something that summary statistics alone do not show. It visualizes the different paths each portfolio would end up in if it started at $1 in January 2005. It also shows the different routes portfolios went through depending on the period.

The 1/N portfolio finishes with the highest ending value. All stocks have equal weight, so when high-growth sectors perform well, 1/N can capture that upside. This happens in two parts of the sample. First, from 2009 to 2014, the post-crisis recovery rewarded broad exposures across all sectors. Second, from 2023 onwards, the AI-driven rally pushed the Information Technology and Communication Services to new highs. The MVO portfolios do not capture that upside in the same way, because they underweight sectors with higher variance, even if they perform well.

The MVO portfolios do not finish with the highest dollar value, but they protect much better when markets fall. This can be spotted in 2008. The S&P 500 dropped sharply, while the constrained MVO stayed much closer to its previous level. Since most of its weight was concentrated in defensive sectors such as Utilities and Consumer Staples, it was able to protect itself. By 2010, the MVO portfolios were already close to their previous highs, while the S&P 500 needed much longer to recover. A similar pattern also appeared during the COVID crash in 2020. The S&P 500 and 1/N fell sharply, while the MVO portfolios absorbed the shock better.

The S&P 500 finishes last. It follows the general market direction as 1/N, but the gap that started during the 2008 crisis was never fully closed and only widened over time.

So, the summary statistics give the average picture, but the cumulative path shows when each portfolio earned or lost its position in the ranking, as shown in Figure 1. The full-period averages hide what is happening underneath. Chapter 7 tests whether the ranking holds across all four sub-periods or comes from specific periods.

7. Sub-Period Analysis

7.1 Pre-Crisis Period (Jan 2005 to Dec 2007)

Table 2: Portfolio Performance Metrics, January 2005 – December 2007.

Metric	Constrained MVO	Unconstrained MVO	1/N Equal Weight	S&P 500
Annualised Return	8.31%	N/A	20.13%	8.59%
Annualised Volatility	4.25%	N/A	9.57%	7.79%
Sharpe Ratio	1.957	N/A	2.104	1.102
Max Drawdown	-2.75%	N/A	-4.21%	-4.84%
Best Month	+3.08%	N/A	+5.97%	+4.43%
Worst Month	-1.82%	N/A	-4.21%	-4.18%

Note: The risk-free rate is set to zero. Unconstrained MVO is excluded from this period — 36 observations for 33 assets produces a near-singular covariance matrix, making the unconstrained results unreliable.

Table 2 shows that the pre-crisis period is one of the samples where 1/N leads on risk-adjusted return. Its Sharpe ratio of 2.104 is above constrained MVO at 1.957, reversing the ranking from the full period. All three portfolios perform well, with drawdowns remaining below -5% and the S&P 500 trailing at 1.102. The unconstrained MVO is excluded from this period as established in Section 5.4.

This period falls within an economic expansion. Board of Governors of the Federal Reserve System (2005) describes the period as one of steady growth, low core inflation, and low long-term interest rates. Overall, the economic market conditions favored business activity. Because of that, the pre-crisis is classified as a calm period with no major shocks to equity markets. The results justify this description as well. All three portfolios keep annualized volatility below 10% with a small drawdown.

The Sharpe ratio results show that this is the one period where 1/N is most competitive on a risk-adjusted basis. The 1/N portfolio achieves a Sharpe ratio of 2.104, which is above the constrained MVO at 1.957. The finding here is the exact opposite of the Chapter 6 results. This means that in a favorable market environment, 1/N can outperform the optimized portfolio. When all sectors are rising together, spreading weight equally across 33 stocks captures broad market gains without needing to identify the minimum-variance combination. Lastly, the S&P 500 achieves a Sharpe ratio of 1.102 here, which is higher than its full-period Sharpe ratio. The pre-crisis result is best explained by market conditions, although the short estimation window of 36 observations for 33 assets also fits the pattern DeMiguel, Garlappi and Uppal (2009) describe, where optimization underperforms when the number of assets is close to the number of observations. This is discussed further in Section 8.1.

Drawdown differences are minimal in this period. All three portfolios stay above -5%, and the worst single months are small compared to what Chapter 6 reports for the full sample. In a stable expansion, downside risk is not as big of a threat as losses remain small regardless of strategy. This is consistent with Bodie, Kane and Marcus (2023), who argue that in expanding markets, diversification alone is often sufficient to limit drawdowns.

7.2 Crisis Period (Jan 2008 to Dec 2009)

Table 3: Portfolio Performance Metrics, January 2008 – December 2009.

Metric	Constrained MVO	Unconstrained MVO	1/N Equal Weight	S&P 500
Annualised Return	1.39%	N/A	1.14%	-8.58%
Annualised Volatility	12.14%	N/A	23.43%	23.51%
Sharpe Ratio	0.114	N/A	0.049	-0.365
Max Drawdown	-17.94%	N/A	-39.24%	-46.42%
Best Month	+7.12%	N/A	+12.29%	+9.57%
Worst Month	-6.77%	N/A	-15.53%	-16.80%

Note: The risk-free rate is set to zero. Unconstrained MVO is excluded from this period — 24 observations for 33 assets produces a rank-deficient covariance matrix. The unconstrained solution converges to near-zero variance, which is not a genuine optimum.

As seen from the table, the crisis period produces the largest differences between portfolios in the entire analysis. The constrained MVO achieves a Sharpe ratio of 0.114 and the 1/N portfolio 0.049, both near zero. S&P 500 drops to -0.365. All metrics fall sharply relative to both the full period and the pre-crisis period. The unconstrained MVO is excluded from this period.

The crisis period runs from January 2008 to December 2009, as defined in Section 5.4. It all started when Lehman Brothers collapsed in September 2008. Following that, a severe shock hit global financial markets. Stock prices fell significantly, volatility spiked, and the correlation between assets rose dramatically (Federal Reserve History, no date). That triggered a chain reaction among investors. Investors panicked, and it led to selling their equity across all sectors at once. Instead of stable, predictable return patterns, the crisis produced sharp and sudden moves in both directions. Table 3 shows those effects for the best and worst months.

What shows the scale of the financial crisis's impact is the maximum drawdown. S&P 500 recorded a severe -46.42% drop. The 1/N portfolio reached -39.24%, and the constrained MVO reached -17.94%. The constrained MVO lost roughly half of what the S&P 500 lost at its worst point. This is a key result that highlights the downside protection from optimization. The constrained MVO also performs significantly better than 1/N on drawdown, which shows that optimization adds clear value over 1/N.

Sharpe ratios and returns point in the same direction. As all portfolios produce low or negative returns, the constrained MVO is the only portfolio that records a positive Sharpe ratio, while 1/N is near zero, and the S&P 500 is negative. The constrained MVO generates an annualized return of 1.39%. The 1/N portfolio returns 1.14%, and the S&P 500 produces a negative annualized return of -8.58%. The optimizer finds a weight combination that stays positive on a risk-adjusted basis even through the toughest period in the sample.

The key reason for constrained MVO's strong performance is tied to its sector allocation. It systematically shifted toward the most defensive sectors. Section 7.6 presents the full sector weight analysis, but one number is worth mentioning here. During the crisis, the optimizer concentrated 71.6% of the portfolio in Utilities and Consumer Staples combined. These are the two sectors with the lowest sensitivity to the economic cycle, which means their returns held up while cyclical sectors collapsed. However, the in-sample nature of this result is discussed as a limitation in Section 8.2.

7.3 Post-Crisis Period (Jan 2010 to Dec 2019)

Table 4: Portfolio Performance Metrics, January 2010 – December 2019.

Metric	Constrained MVO	Unconstrained MVO	1/*N* Equal Weight	S&P 500
Annualised Return	15.30%	15.34%	17.39%	13.56%
Annualised Volatility	7.87%	6.95%	11.42%	12.46%
Sharpe Ratio	1.944	2.207	1.522	1.089
Max Drawdown	-5.06%	-4.13%	-12.24%	-16.26%
Best Month	+6.99%	+5.66%	+9.93%	+10.93%
Worst Month	-5.06%	-4.13%	-7.11%	-9.03%

Note: The risk-free rate is set to zero. All four portfolios are included. With 120 observations this is the best-conditioned sub-period in the analysis.

Table 4 reports the post-crisis results. This is the longest sub-period and the best-conditioned for the estimation window in the entire analysis. The results are most robust for all portfolios, especially MVO. The constrained MVO achieves a Sharpe ratio of 1.944, and the unconstrained reaches 2.207. Both significantly exceed results from Chapter 6. The 1/N portfolio reaches 1.522, below both MVO portfolios but above its full period result. The S&P 500 trails at 1.089, but it is still higher compared to the full period. So even the worst performer in this sub-period does better than over the full 252 months.

The post-crisis period lies within the longest economic expansion in US history, as explained in section 5.4. Throughout this decade, the Federal Reserve maintained historically low interest rates and introduced large-scale quantitative easing programs after the financial crisis (Federal Reserve History, no date). Equity markets recovered steadily, and volatility remained relatively low. This created the most stable estimation environment of any sub-period in the analysis. On top of the economic conditions, there are 120 monthly observations, which means that the covariance matrix is well-conditioned.

MVO performs best during this period, as confirmed by the Sharpe ratio. The unconstrained MVO also performs better than the constrained in this period. With a longer, more stable sample, it allows short positions to improve the portfolio's efficiency rather than introducing instability. That explains why both are above 1/N and the S&P 500. Unlike the pre-crisis period, where 1/N was competitive, the stable estimation environment here gives the optimizer more room to find a genuine advantage. The gap between constrained MVO at 1.944 and 1/N at 1.522 is 0.422 on the Sharpe ratio. This is the largest gap and clearest case of MVO winning out over 1/N. With 120 months of stable conditions, the in-sample optimizer has the strongest advantage of any sub-period.

The drawdown results show the same pattern even more clearly. The constrained MVO records a drawdown of -5.06% and the unconstrained -4.13%, compared to -12.24% for 1/N and -16.26% for the S&P 500. Both MVO portfolios cut drawdown to roughly a third of what the S&P 500 experienced. Even in a decade with no major recession, the optimizer systematically avoids the worst monthly losses that hit the passive benchmark and the 1/N portfolio.

The return-volatility pattern is the same as in the full period. 1/N leads on raw return at 17.39% but at higher volatility, while both MVO portfolios sit below 8% volatility. The result is the same as Chapter 6, as 1/N earns more but takes on proportionally more risk, which is why it finishes behind on Sharpe.

Because this is the longest sub-period and the one where MVO performs best, its weight in the full-period results is discussed in Section 8.1.

7.4 Recent Period (Jan 2020 to Dec 2025)

Table 5: Portfolio Performance Metrics, January 2020 – December 2025.

Metric	Constrained MVO	Unconstrained MVO	1/N Equal Weight	S&P 500
Annualised Return	13.32%	12.31%	19.36%	15.59%
Annualised Volatility	12.15%	9.70%	16.43%	17.15%
Sharpe Ratio	1.096	1.269	1.178	0.909
Max Drawdown	-14.25%	-8.10%	-19.28%	-23.86%
Best Month	+10.11%	+7.25%	+11.93%	+12.82%
Worst Month	-7.73%	-6.43%	-11.34%	-12.35%

Note: The risk-free rate is set to zero. All four portfolios are included. This period covers the COVID-19 crash, the 2021-2022 inflation shock, and the subsequent AI-driven equity rally.

The numbers for this period are in Table 5. The recent period is the hardest to read because it contains three distinct market regimes over 72 months, as defined in Section 5.4. Optimization remains competitive on a risk-adjusted basis, but constrained MVO falls slightly behind 1/N on Sharpe, while the unconstrained MVO records the highest Sharpe ratio of the period. The constrained MVO reaches a Sharpe ratio of 1.096, which is below the full period benchmark of 1.215. This is the sub-period where constrained MVO's advantage over 1/N shrinks the most, with the Sharpe gap narrowing to just 0.082 compared to 0.422 in the post-crisis period.

The COVID-19 crash in March 2020 was the sharpest shock of the period. Baker, Bloom, and Davis (2020) showed that US stock market volatility during the first half of 2020 exceeded even the 2008 financial crisis. That was followed by a rebound driven by fiscal stimulus and monetary easing, then a surge in inflation in 2021 and 2022, and finally an AI-driven equity rally. Each regime rewarded different portfolio characteristics, which is why this period is the hardest to interpret.

On raw return, 1/N leads at 19.36%, with the S&P 500 at 15.59% and constrained MVO at 13.32%. The unconstrained MVO finishes last on return at 12.31%, despite achieving the highest Sharpe ratio of the period. This is because its short positions cut volatility so aggressively, down to 9.70%, that a lower return still produces the best risk-adjusted result. The AI-driven rally explains why 1/N leads on return. It held full equal weight in every sector, including Information Technology and Communication Services, which were the strongest performers from 2023 onwards. That return advantage is what brought 1/N's Sharpe closer to constrained MVO in this period.

Sharpe ratios still favor MVO, but only the unconstrained. The constrained MVO concentrates on defensive sectors, which limits its upside when growth stocks drive the market. As a result, 1/N benefits from the broad exposure that the AI rally rewarded. Even the S&P 500 closes the gap to 0.909, the highest it achieves in any sub-period.

The drawdown results are more consistent with the full period pattern. The constrained MVO drops -14.25% and the unconstrained only -8.10%, while 1/N reaches -19.28% and the S&P 500 falls to -23.86%. MVO holds up better on the downside, even though 1/N moves ahead on Sharpe in this period. The unconstrained MVO's result is worth pointing out since the -8.10% is the smallest drawdown of this period, and the reason for that is its short positions. They provided some protection during the COVID crash in early 2020. So even when the Sharpe ratio gap shrinks, the optimization still cuts the worst loss roughly in half compared to the S&P 500.

The conditions within this sub-period shift fast. Within this in-sample framework, MVO adapts to each regime through its covariance-based weights, but it cannot anticipate regime changes before they happen. The tech-driven rally is a good example of that. The portfolio that captured most of the growth exposure was rewarded over variance minimization, which is also what pulls the full period Sharpe ratios down from the post-crisis highs.

7.5 Comparative Summary

Table 6: Sharpe Ratio Across Sub-Periods and Full Period.

Period	Constrained MVO	Unconstrained MVO	1/N Equal Weight	S&P 500
Pre-Crisis (2005–2007)	1.957	N/A	2.104	1.102
Crisis (2008–2009)	0.114	N/A	0.049	-0.365
Post-Crisis (2010–2019)	1.944	2.207	1.522	1.089
Recent (2020–2025)	1.096	1.269	1.178	0.909
Full Period (2005–2025)	1.215	1.347	1.177	0.764

Note: The risk-free rate is set to zero. Unconstrained MVO is excluded from the Pre-Crisis and Crisis periods due to insufficient observations for reliable covariance matrix estimation. Full period results from Table 1 are included as the benchmark.

Table 7: Maximum Drawdown Across Sub-Periods and Full Period.

Period	Constrained MVO	Unconstrained MVO	1/N Equal Weight	S&P 500
Pre-Crisis (2005–2007)	-2.75%	N/A	-4.21%	-4.84%
Crisis (2008–2009)	-17.94%	N/A	-39.24%	-46.42%
Post-Crisis (2010–2019)	-5.06%	-4.13%	-12.24%	-16.26%
Recent (2020–2025)	-14.25%	-8.10%	-19.28%	-23.86%
Full Period (2005–2025)	-21.22%	-14.16%	-41.31%	-50.95%

Note: Maximum drawdown is the largest peak-to-trough decline over each period. Unconstrained MVO is excluded from the Pre-Crisis and Crisis periods due to insufficient observations for reliable covariance matrix estimation. Full period results from Table 1 are included as the benchmark.

Tables 6 and 7 put Sharpe ratios and maximum drawdowns across all four sub-periods and the full period side by side. Together, they show whether a portfolio performs well on average and whether it protects the investor when things go wrong.

The Sharpe pattern in Table 6 is not one-sided. Constrained MVO leads in the crisis period at 0.114 and in the post-crisis period at 1.944. But 1/N is ahead in the pre-crisis at 2.104 versus 1.957, and again in the recent period at 1.178 versus 1.096. So, on the Sharpe ratio, the result is split. Constrained MVO wins in the two sub-periods where market conditions are either extreme or very stable for estimation. 1/N wins in the calm bull market and in the recent period, where broad exposure was rewarded. Over the full period, constrained MVO leads 1.215 to 1.177, a gap of just 0.038. In practice, the higher turnover required by MVO rebalancing could generate enough transaction costs to close the difference, which means the two portfolios are close to equal on a risk-adjusted basis. The S&P 500 trails all three constructed portfolios across all sub-periods.

Table 7 is where the case for optimization is clearest. Constrained MVO records the lowest maximum drawdown in every single sub-period. In the crisis, the gap is largest, at -17.94% versus -39.24% for 1/N and -46.42% for the S&P 500. In the post-crisis and recent periods, the absolute drawdowns are smaller, but the ranking holds. Across all sub-periods, the constrained MVO maximum drawdown is roughly half that of 1/N and less than half that of the S&P 500. Within the in-sample design, the optimizer finds the weights that get the portfolio closest to the lowest maximum drawdown.

Returning to the research question from Chapter 2, within this in-sample setting, the evidence is split depending on which metric is used. On Sharpe, constrained MVO leads in two of four sub-periods and over the full period, but the advantage is modest, and 1/N is competitive when conditions are favorable. On drawdown, the result is consistent across every sub-period. Constrained MVO limits the worst loss better than both 1/N and the S&P 500, regardless of the market regime. The clearest evidence for optimization is the drawdown results, not the Sharpe ratio results.

7.6 Sector Weight Analysis

Table 8: Constrained MVO Portfolio Sector Allocation Across Periods.

GICS Sector	Full Period	Pre-Crisis (2005-2007)	Crisis (2008-2009)	Post-Crisis (2010-2019)	Recent (2020-2025)	1/N Equal Weight
Consumer Staples	28.2%	22.8%	26.1%	13.0%	26.9%	9.1%
Utilities	25.6%	26.3%	45.5%	21.1%	18.9%	9.1%
Communication Services	12.6%	0.0%	5.3%	2.7%	23.9%	9.1%
Health Care	9.1%	10.9%	0.0%	13.6%	16.4%	9.1%
Consumer Discretionary	8.8%	0.0%	0.0%	17.1%	5.1%	9.1%
Energy	6.5%	8.9%	23.1%	3.0%	4.3%	9.1%
Financials	5.2%	11.1%	0.0%	16.4%	0.0%	9.1%
Information Technology	1.0%	0.0%	0.0%	3.0%	4.6%	9.1%
Materials	2.6%	4.0%	0.0%	6.0%	0.0%	9.1%
Real Estate	0.4%	0.0%	0.0%	4.2%	0.0%	9.1%
Industrials	0.0%	16.0%	0.0%	0.0%	0.0%	9.1%
Total	**100%**	**100%**	**100%**	**100%**	**100%**	**100%**

Note: Sector weights shown for the constrained MVO portfolio only. The 1/N portfolio allocates equally across all 11 sectors at 9.1% each. The crisis period shows a heavy concentration in Utilities (45.5%) and Consumer Staples (26.1%), totalling 71.6% in two defensive sectors.

Table 8 shows how the constrained MVO portfolio allocates across the 11 GICS sectors in each period. The 1/N equal weight is included as a reference. The performance tables in the previous sections show what happened. This table shows why.

Even over the full period, the optimizer leans heavily toward defensive sectors. Consumer Staples takes 28.2% and Utilities 25.6%, making up over half the portfolio between just two sectors. Information Technology, one of the best-performing sectors of the past decade, receives only 1.0%. The optimizer is not chasing returns. It is finding the lowest variance combination, and over the full 252 months, that combination is dominated by low-volatility sectors.

The crisis period makes this most visible. During January 2008 to December 2009, the constrained MVO put 45.5% in Utilities and 26.1% in Consumer Staples. That is 71.6% of the entire portfolio in two defensive sectors. Energy picked up another 23.1%. Everything else received close to zero. This is the direct explanation for the -17.94% maximum drawdown from Section 7.2. When the financial crisis hit and cyclical sectors collapsed, the optimizer was already positioned in the parts of the market that held up.

The weights shift across periods depending on the market environment. In the pre-crisis period, the portfolio is more spread out. Industrials take 16.0% and Financials 11.1%, sectors that performed well during the stable expansion. Post-crisis, the allocation moves toward Consumer Discretionary at 17.1% and Health Care at 13.6%. This was a long recovery period where volatility was low, and more sectors contributed to stability. In the recent period, Communication Services jumped to 23.9%. It captures lower-volatility exposure within the tech-adjacent space, while Information Technology itself only receives 4.6%. So, the optimizer is not static. It reads the covariance structure of each period and adjusts.

Figure 2. Constrained MVO Sector Allocation Across Periods

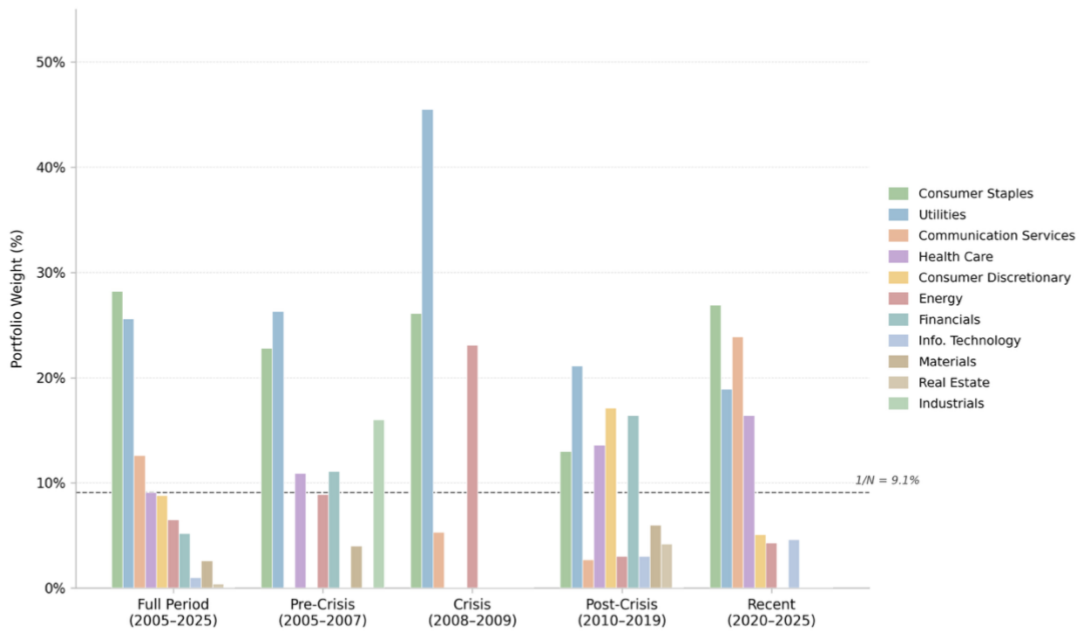

Note: Sector weights for the constrained MVO portfolio only. Dashed line = 9.1% equal weight under 1/N. Unconstrained MVO not shown.

The transitions between regimes are where the economic story is clearest. Moving from pre-crisis to crisis, the optimizer dropped cyclical sectors entirely. Financials fell from 11.1% to 0.0% and Industrials from 16.0% to 0.0%, while Utilities jumped from 26.3% to 45.5%. It loaded into the one sector whose returns held up when everything else was falling. Moving from crisis to post-crisis, the allocation broadened out. Defensive concentration came down, and new weight appeared in Consumer Discretionary, Health Care, and Financials. The covariance structure no longer rewarded loading into just two sectors. Moving from post-crisis to recent, Communication Services jumped from 2.7% to 23.9%, and Consumer Staples returned to 26.9%, while Financials dropped back to 0.0%. Each transition follows the same logic. The optimizer re-reads the covariance structure and shifts weight toward whichever sectors carry the least variance in that regime.

That is what makes it active in a specific technical sense, without any return forecasting. It simply responds to changes in how stocks co-move. The 1/N portfolio never moves. It holds 9.1% in every sector regardless of what is happening in the market. In a calm period like post-crisis, that works well enough since broad diversification is sufficient. But in a crisis, holding 9.1% in Financials, Industrials, and Consumer Discretionary while those sectors collapse is exactly the kind of exposure the constrained MVO avoids entirely. The optimizer tilts toward whatever is low variance, given the data. 1/N does not. That is why constrained MVO leads on drawdown in every single sub-period, even when it falls behind on raw return or the Sharpe ratio (Figure 2).

8. Discussion

8.1 Interpretation of Main Findings

Constrained MVO records the lowest maximum drawdown in every sub-period. This directly supports the central claim of Markowitz (1952) that variance minimization produces a better risk outcome than naive weighting. However, on the Sharpe ratio, the advantage is smaller and dependent on the market regime. So, the answer depends on what the investor cares about. To avoid the worst losses, constrained MVO is the better option. For maximizing return, 1/N wins.

The reason the two metrics diverge is that they measure different things. Drawdown captures the worst-case loss, which is exactly what the optimizer targets. Sharpe rewards broad exposure when growth sectors rally, which favors 1/N. That is why the ranking switches between sub-periods. No single answer covers both metrics.

As discussed in Section 3.5, DeMiguel, Garlappi and Uppal (2009) show that 1/N outperforms optimized strategies out-of-sample when the number of assets is large relative to the estimation window. The pre-crisis period, with 36 observations for 33 assets, is where that concern is most relevant. At that ratio, the covariance matrix is near-singular, and the optimizer is working with noisy estimates. That is exactly the setting where their result would predict 1/N to do well. However, since this thesis evaluates portfolios in-sample, the estimation error effect cannot be separated from the fact that the pre-crisis period is also a broad expansion where equal weighting already does well on its own. So, the pre-crisis result is consistent with the prediction, but it cannot confirm it. The recent period is the cleaner comparison. With 72 observations for 33 assets, estimation noise is not the issue, yet 1/N still leads on raw return. That is driven by the AI-driven rally, not by a poorly conditioned matrix.

The objective function explains why MVO beats 1/N on Sharpe but not on raw return. Over the full period, constrained MVO returns 12.67% against 16.80% for 1/N, which is a gap of 4.13%. But at a volatility of 10.43% versus 14.27%, the Sharpe ratio favors MVO. The unconstrained MVO performs even better on Sharpe when included, but its exclusion from two sub-periods and its reliance on short selling make the constrained version the more relevant comparison. The 0.038 Sharpe gap between constrained MVO and 1/N is modest. The difference is not in the Sharpe number itself but in how each portfolio gets there. MVO gets there through lower volatility and drawdown protection. 1/N gets there through higher raw returns.

The sector weight analysis in Section 7.6 explains the pattern mechanically. The non-negativity constraint stops extreme short positions that would make the solution unstable. That stability is part of why constrained MVO holds up across regimes.

But that level of concentration carries a risk that the in-sample design hides. If the optimizer had identified the wrong defensive sectors, the result could have been worse than 1/N. Out-of-sample, concentration at that scale could become a vulnerability rather than an advantage. This is one reason why the drawdown results, while the most consistent finding, should be interpreted with caution.

The full-period Sharpe ratio should also be read considering how the sample is weighted. The post-crisis period accounts for nearly half of the total sample. It covers 120 out of 252 months. It is also the sub-period where MVO performs best, with a Sharpe ratio of 1.944 for the constrained MVO portfolio. So, the full-period Sharpe ratio of 1.215 is partly a product of this extended stretch of stable, low-volatility conditions where MVO thrives. If the sample had been weighted more toward crisis or recent-type conditions, the full-period advantage would likely be smaller. A different sample split might have produced a smaller gap.

8.2 Limitations

The biggest limitation is the in-sample design. Portfolio weights are estimated and evaluated on the same data within each period. The optimizer sees the full crisis period when building the crisis weights, which is why it concentrates so heavily in defensive sectors. A real investor in January 2008 had none of that information. This inflates the crisis results and means the findings should be interpreted as describing what optimization can achieve given the data, not what it would have achieved in practice. This limitation was first flagged in Section 7.2 and is the primary reason out-of-sample testing would be the natural next step.

The selection procedure introduces survivorship bias, as discussed in Section 4.2. All 33 stocks were selected based on their March 2026 market capitalization ranking. This means they are companies that survived and grew over the full sample period. Following Brown et al. (1992), this implies the results describe the performance of large, established, surviving US firms rather than the broader equity market. Stocks that failed or declined over the sample period are excluded by construction. The results likely overstate what would be achievable in a portfolio that did not have the benefit of this future outlook selection.

Between sub-periods, MVO requires significant rebalancing as the optimal weights can shift substantially across regimes. The 1/N portfolio always returns to equal weights, so its turnover is minimal. In practice, the transaction costs from MVO rebalancing would narrow or eliminate the Sharpe advantage. The sector weight shifts between sub-periods are substantial. For example, Utilities drops from 45.5% during the crisis to 21.1% in the post-crisis period, while Financials swings from 0.0% to 16.4%. These swings in weights imply high portfolio turnover at regime transitions. After accounting for those costs, the Sharpe advantage of constrained MVO over 1/N could narrow or disappear, so the two strategies may end up close to equal on risk-adjusted return.

The unconstrained MVO is excluded from the pre-crisis and crisis periods due to insufficient observations for reliable covariance matrix estimation, as established in Section 5.4. This limits the cross-period comparability of the unconstrained portfolio and means the constrained version carries the full weight of the MVO case in two of four sub-periods.

The S&P 500 is sourced from Investing.com rather than CRSP because CRSP coverage did not extend to December 2025 at the time of data collection, as noted in Section 4.1. The figures match published index returns, and the data is consistent across the full sample period. A single source for all series would be preferable and remains a direction for future data collection.

The zero risk-free rate assumption is discussed in Section 5.2 and is the main caveat for reading the reported Sharpe ratios as directional rather than absolute.

8.3 Out-of-Sample Considerations and Future Research

The in-sample design means the results cannot be generalized to new data. A rolling window out-of-sample approach, as applied by DeMiguel, Garlappi and Uppal (2009), would be the next step. That would mean estimating weights on a training window and evaluating performance on the following period. Their finding that 1/N outperforms optimized strategies out-of-sample remains the open question this thesis raises but does not answer. The key empirical question is whether the drawdown advantage survives once the optimizer loses access to the data it is being tested on. How long the estimation window should be is part of that question. The post-crisis sub-period, with 120 monthly observations, is where MVO performs best. In a rolling design, an estimation window that long may not be realistic, while the 60-month baseline from Chan, Karceski and Lakonishok (1999) might be too short for 33 assets.

It is also unclear whether these results generalize beyond this sample. This thesis uses 33 large-cap US stocks from the S&P 500. Whether the same patterns would hold with mid-cap stocks, international equities, or a different number of sectors is something the data here cannot answer. DeMiguel, Garlappi and Uppal (2009) tested across seven datasets, which is part of what makes their finding robust. This thesis tests one. Expanding the assets would help show whether the drawdown advantage of constrained MVO generalizes beyond large-cap US equities.

9. Conclusion

This thesis examines whether constrained minimum-variance optimization outperforms naive 1/N diversification and the S&P 500 in-sample across January 2005 to December 2025. The answer depends on the metric. On drawdown, constrained MVO wins in every sub-period without exception. On Sharpe, it leads over the full period and in two of four sub-periods, but the advantage is modest and conditional on the market regime.

Drawdown is where the results are most consistent. Constrained MVO records the lowest maximum drawdown in all four sub-periods. The gap is widest during the crisis, where the optimized portfolio suffered only around 40% of the loss recorded by the passive benchmark. In the other sub-periods, the absolute drawdowns are smaller, but the ranking does not change.

On Sharpe, constrained MVO leads in the crisis and post-crisis periods but 1/N edges ahead in the pre-crisis and recent periods. Over the full period, constrained MVO leads but the gap is small. The unconstrained MVO produces the strongest Sharpe result when included, but its exclusion from two sub-periods and reliance on short selling make the constrained version the more relevant comparison for the research question.

The S&P 500 trails all three constructed portfolios on both metrics in every period. As a benchmark, it shows that any form of active allocation, whether optimized or naive, produced better results than simply holding the market index over this sample.

Even in-sample, 1/N is competitive on Sharpe in two sub-periods, which fits with what DeMiguel, Garlappi and Uppal (2009) showed out-of-sample. This thesis addresses an in-sample question that their study was not designed to answer. MVO's theoretical advantage does appear when estimation error is reduced by in-sample evaluation, but it is conditional on the metric and the market regime. The main limitation is the in-sample design. The optimizer sees the full data before assigning weights, which inflates the results, particularly during the crisis. Survivorship bias is also built into the sample since all 33 stocks were selected based on their March 2026 market capitalization. After accounting for higher MVO turnover, the two strategies are effectively equivalent on the Sharpe ratio. The case for constrained MVO therefore becomes drawdown protection, not an advantage in the Sharpe ratio.

What this sample cannot answer is whether the drawdown advantage survives once the optimizer loses access to the data on which it is tested. The in-sample result sets the ceiling, and an out-of-sample test would show how much of it holds in practice.

Bibliography

Baker, S.R., Bloom, N. and Davis, S.J. (2020) 'COVID-induced economic uncertainty', *NBER Working Paper No. 26983*. Cambridge, MA: National Bureau of Economic Research.

Board of Governors of the Federal Reserve System (2005) *Annual Report 2005: Monetary Policy and the Economic Outlook*. Washington, DC: Federal Reserve. Available at: https://www.federalreserve.gov/boarddocs/rptcongress/annual05/sec1/c2.htm (Accessed: 10 April 2026).

Bodie, Z., Kane, A. and Marcus, A.J. (2023) *Investments*. 13th edn. New York: McGraw-Hill.

Brown, S.J., Goetzmann, W., Ibbotson, R.G. and Ross, S.A. (1992) 'Survivorship bias in performance studies', *Review of Financial Studies*, 5(4), pp. 553-580.

Chan, L.K.C., Karceski, J. and Lakonishok, J. (1999) 'On portfolio optimization: Forecasting covariances and choosing the risk model', *Review of Financial Studies*, 12(5), pp. 937-974.

DeMiguel, V., Garlappi, L. and Uppal, R. (2009) 'Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?', *Review of Financial Studies*, 22(5), pp. 1915-1953.

Evans, J.L. and Archer, S.H. (1968) 'Diversification and the reduction of dispersion: An empirical analysis', *Journal of Finance*, 23(5), pp. 761-767.

Federal Reserve History (no date) *The Great Recession and its Aftermath*. Available at: https://www.federalreservehistory.org/essays/great-recession-and-its-aftermath (Accessed: 2 April 2026).

Investing.com (no date) *S&P 500 Historical Data*. Available at: https://www.investing.com/indices/us-spx-500-historical-data (Accessed: 15 March 2026).

Jagannathan, R. and Ma, T. (2003) 'Risk reduction in large portfolios: Why imposing the wrong constraints helps', *Journal of Finance*, 58(4), pp. 1651-1683.

Lintner, J. (1965) 'The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets', *Review of Economics and Statistics*, 47(1), pp. 13-37.

Markowitz, H. (1952) 'Portfolio selection', *Journal of Finance*, 7(1), pp. 77-91.

Michaud, R.O. (1989) 'The Markowitz optimization enigma: Is "optimized" optimal?', *Financial Analysts Journal*, 45(1), pp. 31-42.

Mossin, J. (1966) 'Equilibrium in a capital asset market', *Econometrica*, 34(4), pp. 768-783.

MSCI and S&P Dow Jones Indices (1999) *Global Industry Classification Standard (GICS) Methodology*. New York: MSCI Inc. and S&P Dow Jones Indices LLC.

NBER Business Cycle Dating Committee (NBER BCDC) (2008) *Business Cycle Dating Committee Announcement December 1, 2008*. Cambridge, MA: NBER. Available at: https://www.nber.org/news/business-cycle-dating-committee-announcement-december-1-2008 (Accessed: 3 April 2026).

NBER Business Cycle Dating Committee (NBER BCDC) (2010) *Business Cycle Dating Committee Announcement September 20, 2010*. Cambridge, MA: NBER. Available at: https://www.nber.org/news/business-cycle-dating-committee-announcement-september-20-2010 (Accessed: 5 April 2026).

NBER Business Cycle Dating Committee (NBER BCDC) (2020) *Business Cycle Dating Committee Announcement June 8, 2020*. Cambridge, MA: NBER. Available at: https://www.nber.org/news/business-cycle-dating-committee-announcement-june-8-2020 (Accessed: 4 April 2026).

S&P Dow Jones Indices (2018) *Communication Services Sector to Replace Telecommunication Services in the GICS Structure*. New York: S&P Dow Jones Indices LLC.

Sharpe, W.F. (1964) 'Capital asset prices: A theory of market equilibrium under conditions of risk', *Journal of Finance*, 19(3), pp. 425-442.

Sharpe, W.F. (1966) 'Mutual fund performance', *Journal of Business*, 39(1), pp. 119-138.

Slickcharts (2026) *S&P 500 Companies by Weight*. Available at: https://www.slickcharts.com/sp500 (Accessed: 15 March 2026).

Statman, M. (1987) 'How many stocks make a diversified portfolio?', *Journal of Financial and Quantitative Analysis*, 22(3), pp. 353-363.

U.S. Congress (2002) *Sarbanes-Oxley Act of 2002*, Public Law 107-204. Washington, DC: U.S. Government. Available at: https://www.congress.gov/bill/107th-congress/house-bill/3763 (Accessed: 7 April 2026).

Appendix

This appendix contains the full 33-stock sample with substitution notes (Table A1), the constrained MVO portfolio weights for the full period and all four sub-periods (Tables A2 to A6), the unconstrained MVO portfolio weights for the three periods where the optimizer is included (Tables A7 to A9), and the Excel Solver setup and convergence documentation (Table A10). The appendix does not count toward the page or character limit.

Table A1: Full sample of 33 stocks

Three stocks per GICS sector across all 11 sectors. Tickers are listed in the order used throughout the analysis. The substitution notes column records changes from the top three by S&P 500 market capitalization, where the original stock did not have continuous monthly returns from January 2005.

Ticker	Company Name	GICS Sector	Substitution Notes
NVDA	NVIDIA Corp.	Information Technology	
MSFT	Microsoft Corp.	Information Technology	
AAPL	Apple Inc.	Information Technology	
GOOGL	Alphabet Inc.	Communication Services	
NFLX	Netflix Inc.	Communication Services	replaced META (May 2012 IPO)
VZ	Verizon Communications	Communication Services	replaced TMUS (corporate entity dates to 2013)
LLY	Eli Lilly and Co.	Health Care	
JNJ	Johnson & Johnson	Health Care	
MRK	Merck & Co.	Health Care	
BRK.B	Berkshire Hathaway	Financials	
JPM	JPMorgan Chase & Co.	Financials	
BAC	Bank of America Corp.	Financials	replaced V (March 2008 IPO) and MA (May 2006 IPO); chosen as 5th largest qualifying
AMZN	Amazon.com Inc.	Consumer Discretionary	
HD	Home Depot Inc.	Consumer Discretionary	
MCD	McDonald's Corp.	Consumer Discretionary	replaced TSLA (June 2010 IPO)
WMT	Walmart Inc.	Consumer Staples	
COST	Costco Wholesale	Consumer Staples	
PG	Procter & Gamble	Consumer Staples	
XOM	Exxon Mobil Corp.	Energy	
CVX	Chevron Corp.	Energy	
COP	ConocoPhillips	Energy	
GE	GE Aerospace	Industrials	
CAT	Caterpillar Inc.	Industrials	
BA	Boeing Co.	Industrials	

Ticker	Company Name	GICS Sector	Substitution Notes
PLD	Prologis Inc.	Real Estate	
AMT	American Tower	Real Estate	
WELL	Welltower Inc.	Real Estate	
SHW	Sherwin-Williams Co.	Materials	replaced LIN (current entity created through 2018 merger)
NEM	Newmont Corp.	Materials	
FCX	Freeport-McMoRan	Materials	
NEE	NextEra Energy	Utilities	
SO	Southern Company	Utilities	
DUK	Duke Energy Corp.	Utilities	

Table A2: Constrained MVO weights, full period (Jan 2005 to Dec 2025)

Stock-level and sector-level weights for the constrained MVO portfolio over the full 252-month period. The non-negativity constraint forces all weights to be non-negative, and the full-investment constraint forces them to sum to one.

Ticker	GICS Sector	Weight
NVDA	Information Technology	0.00%
MSFT	Information Technology	0.99%
AAPL	Information Technology	0.00%
GOOGL	Communication Services	2.64%
NFLX	Communication Services	2.75%
VZ	Communication Services	7.21%
LLY	Health Care	5.91%
JNJ	Health Care	1.83%
MRK	Health Care	1.41%
BRK.B	Financials	5.17%
JPM	Financials	0.00%
BAC	Financials	0.00%
AMZN	Consumer Discretionary	0.00%
HD	Consumer Discretionary	0.96%
MCD	Consumer Discretionary	7.79%
WMT	Consumer Staples	16.45%
COST	Consumer Staples	0.05%
PG	Consumer Staples	11.68%
XOM	Energy	6.54%
CVX	Energy	0.00%
COP	Energy	0.00%
GE	Industrials	0.00%
CAT	Industrials	0.00%
BA	Industrials	0.00%
PLD	Real Estate	0.00%
AMT	Real Estate	0.39%

Ticker	GICS Sector	Weight
WELL	Real Estate	0.00%
SHW	Materials	0.00%
NEM	Materials	2.58%
FCX	Materials	0.00%
NEE	Utilities	0.93%
SO	Utilities	10.64%
DUK	Utilities	14.08%
Total		100.00%

Sector-level totals (constrained MVO, full period):

GICS Sector	Constrained MVO Weight
Information Technology	0.99%
Communication Services	12.60%
Health Care	9.14%
Financials	5.17%
Consumer Discretionary	8.75%
Consumer Staples	28.18%
Energy	6.54%
Industrials	0.00%
Real Estate	0.39%
Materials	2.58%
Utilities	25.64%
Total	100.00%

Table A3: Constrained MVO weights, pre-crisis period (Jan 2005 to Dec 2007)

Constrained MVO weights for the pre-crisis period. With 36 monthly observations, this period falls below the 60-month threshold from Chan, Karceski and Lakonishok (1999) but contains more observations than assets. The unconstrained MVO portfolio is excluded from this period.

Ticker	GICS Sector	Weight
NVDA	Information Technology	0.00%
MSFT	Information Technology	0.00%
AAPL	Information Technology	0.00%
GOOGL	Communication Services	0.00%
NFLX	Communication Services	0.00%
VZ	Communication Services	0.00%
LLY	Health Care	3.22%
JNJ	Health Care	7.67%
MRK	Health Care	0.00%
BRK.B	Financials	1.41%
JPM	Financials	0.00%
BAC	Financials	9.74%
AMZN	Consumer Discretionary	0.00%

Ticker	GICS Sector	Weight
HD	Consumer Discretionary	0.00%
MCD	Consumer Discretionary	0.00%
WMT	Consumer Staples	20.35%
COST	Consumer Staples	2.41%
PG	Consumer Staples	0.00%
XOM	Energy	0.00%
CVX	Energy	0.00%
COP	Energy	8.88%
GE	Industrials	9.60%
CAT	Industrials	2.00%
BA	Industrials	4.44%
PLD	Real Estate	0.00%
AMT	Real Estate	0.00%
WELL	Real Estate	0.00%
SHW	Materials	0.35%
NEM	Materials	3.66%
FCX	Materials	0.00%
NEE	Utilities	0.00%
SO	Utilities	21.70%
DUK	Utilities	4.60%
Total		100.00%

Sector-level totals (constrained MVO, pre-crisis):

GICS Sector	Constrained MVO Weight
Information Technology	0.00%
Communication Services	0.00%
Health Care	10.89%
Financials	11.14%
Consumer Discretionary	0.00%
Consumer Staples	22.75%
Energy	8.88%
Industrials	16.03%
Real Estate	0.00%
Materials	4.00%
Utilities	26.30%
Total	100.00%

Table A4: Constrained MVO weights, crisis period (Jan 2008 to Dec 2009)

Constrained MVO weights for the crisis period. With 24 observations for 33 assets, the covariance matrix is rank-deficient. The non-negativity constraint stabilizes the constrained solution. The unconstrained MVO portfolio is excluded.

Ticker	GICS Sector	Weight
NVDA	Information Technology	0.00%
MSFT	Information Technology	0.00%
AAPL	Information Technology	0.00%
GOOGL	Communication Services	0.00%
NFLX	Communication Services	5.25%
VZ	Communication Services	0.00%
LLY	Health Care	0.00%
JNJ	Health Care	0.00%
MRK	Health Care	0.00%
BRK.B	Financials	0.00%
JPM	Financials	0.00%
BAC	Financials	0.00%
AMZN	Consumer Discretionary	0.00%
HD	Consumer Discretionary	0.00%
MCD	Consumer Discretionary	0.00%
WMT	Consumer Staples	26.11%
COST	Consumer Staples	0.00%
PG	Consumer Staples	0.00%
XOM	Energy	23.12%
CVX	Energy	0.00%
COP	Energy	0.00%
GE	Industrials	0.00%
CAT	Industrials	0.00%
BA	Industrials	0.00%
PLD	Real Estate	0.00%
AMT	Real Estate	0.00%
WELL	Real Estate	0.00%
SHW	Materials	0.00%
NEM	Materials	0.00%
FCX	Materials	0.00%
NEE	Utilities	5.10%
SO	Utilities	8.78%
DUK	Utilities	31.64%
Total		100.00%

Sector-level totals (constrained MVO, crisis):

GICS Sector	Constrained MVO Weight
Information Technology	0.00%
Communication Services	5.25%
Health Care	0.00%
Financials	0.00%
Consumer Discretionary	0.00%
Consumer Staples	26.11%
Energy	23.12%
Industrials	0.00%
Real Estate	0.00%
Materials	0.00%
Utilities	45.52%
Total	100.00%

Table A5: Constrained MVO weights, post-crisis period (Jan 2010 to Dec 2019)

Constrained MVO weights for the post-crisis period. With 120 monthly observations, this is the longest and best-conditioned sub-period in the analysis. All four portfolios are included.

Ticker	GICS Sector	Weight
NVDA	Information Technology	2.83%
MSFT	Information Technology	0.19%
AAPL	Information Technology	0.00%
GOOGL	Communication Services	0.75%
NFLX	Communication Services	1.92%
VZ	Communication Services	0.00%
LLY	Health Care	13.18%
JNJ	Health Care	0.00%
MRK	Health Care	0.44%
BRK.B	Financials	15.86%
JPM	Financials	0.00%
BAC	Financials	0.49%
AMZN	Consumer Discretionary	0.00%
HD	Consumer Discretionary	0.00%
MCD	Consumer Discretionary	17.14%
WMT	Consumer Staples	5.14%
COST	Consumer Staples	0.00%
PG	Consumer Staples	7.80%
XOM	Energy	0.00%
CVX	Energy	0.00%
COP	Energy	3.02%
GE	Industrials	0.00%
CAT	Industrials	0.00%
BA	Industrials	0.00%

Ticker	GICS Sector	Weight
PLD	Real Estate	0.00%
AMT	Real Estate	4.16%
WELL	Real Estate	0.00%
SHW	Materials	2.56%
NEM	Materials	3.47%
FCX	Materials	0.00%
NEE	Utilities	0.00%
SO	Utilities	14.65%
DUK	Utilities	6.40%
Total		100.00%

Sector-level totals (constrained MVO, post-crisis):

GICS Sector	Constrained MVO Weight
Information Technology	3.02%
Communication Services	2.67%
Health Care	13.62%
Financials	16.35%
Consumer Discretionary	17.14%
Consumer Staples	12.95%
Energy	3.02%
Industrials	0.00%
Real Estate	4.16%
Materials	6.03%
Utilities	21.05%
Total	100.00%

Table A6: Constrained MVO weights, recent period (Jan 2020 to Dec 2025)

Constrained MVO weights for the recent period, covering the COVID-19 crash, the 2021 to 2022 inflation shock, and the AI-driven equity rally. All four portfolios are included.

Ticker	GICS Sector	Weight
NVDA	Information Technology	0.00%
MSFT	Information Technology	4.60%
AAPL	Information Technology	0.00%
GOOGL	Communication Services	5.42%
NFLX	Communication Services	0.00%
VZ	Communication Services	18.51%
LLY	Health Care	5.17%
JNJ	Health Care	10.28%
MRK	Health Care	0.92%
BRK.B	Financials	0.00%
JPM	Financials	0.00%
BAC	Financials	0.00%
AMZN	Consumer Discretionary	4.32%

Ticker	GICS Sector	Weight
HD	Consumer Discretionary	0.00%
MCD	Consumer Discretionary	0.75%
WMT	Consumer Staples	13.75%
COST	Consumer Staples	0.00%
PG	Consumer Staples	13.17%
XOM	Energy	4.27%
CVX	Energy	0.00%
COP	Energy	0.00%
GE	Industrials	0.00%
CAT	Industrials	0.00%
BA	Industrials	0.00%
PLD	Real Estate	0.00%
AMT	Real Estate	0.00%
WELL	Real Estate	0.00%
SHW	Materials	0.00%
NEM	Materials	0.00%
FCX	Materials	0.00%
NEE	Utilities	0.87%
SO	Utilities	0.00%
DUK	Utilities	17.98%
Total		100.00%

Sector-level totals (constrained MVO, recent):

GICS Sector	Constrained MVO Weight
Information Technology	4.60%
Communication Services	23.93%
Health Care	16.37%
Financials	0.00%
Consumer Discretionary	5.07%
Consumer Staples	26.92%
Energy	4.27%
Industrials	0.00%
Real Estate	0.00%
Materials	0.00%
Utilities	18.85%
Total	100.00%

Table A7: Unconstrained MVO weights, full period (Jan 2005 to Dec 2025)

Stock-level and sector-level weights for the unconstrained MVO portfolio over the full period. Negative weights indicate short positions.

Ticker	GICS Sector	Weight
NVDA	Information Technology	-1.07%
MSFT	Information Technology	+3.72%

Ticker	GICS Sector	Weight
AAPL	Information Technology	+0.08%
GOOGL	Communication Services	+3.22%
NFLX	Communication Services	+2.62%
VZ	Communication Services	+4.82%
LLY	Health Care	+6.94%
JNJ	Health Care	+0.61%
MRK	Health Care	+1.56%
BRK.B	Financials	+9.24%
JPM	Financials	+9.51%
BAC	Financials	-4.98%
AMZN	Consumer Discretionary	+0.50%
HD	Consumer Discretionary	+5.21%
MCD	Consumer Discretionary	+6.17%
WMT	Consumer Staples	+13.39%
COST	Consumer Staples	+4.13%
PG	Consumer Staples	+12.67%
XOM	Energy	+10.52%
CVX	Energy	-9.52%
COP	Energy	+4.97%
GE	Industrials	-1.48%
CAT	Industrials	-1.01%
BA	Industrials	+1.49%
PLD	Real Estate	-9.17%
AMT	Real Estate	+3.51%
WELL	Real Estate	-8.53%
SHW	Materials	-1.84%
NEM	Materials	+3.45%
FCX	Materials	-1.83%
NEE	Utilities	+3.35%
SO	Utilities	+12.02%
DUK	Utilities	+15.74%
Total		+100.00%

Sector-level totals (unconstrained MVO, full period):

GICS Sector	Unconstrained MVO Weight
Information Technology	2.73%
Communication Services	10.66%
Health Care	9.11%
Financials	13.76%
Consumer Discretionary	11.88%
Consumer Staples	30.19%
Energy	5.97%
Industrials	-0.99%

GICS Sector	Unconstrained MVO Weight
Real Estate	-14.19%
Materials	-0.23%
Utilities	31.11%
Total	100.00%

Table A8: Unconstrained MVO weights, post-crisis period (Jan 2010 to Dec 2019)

Unconstrained MVO weights for the post-crisis period. With 120 stable observations, short positions improve portfolio efficiency rather than introducing instability.

Ticker	GICS Sector	Weight
NVDA	Information Technology	+1.29%
MSFT	Information Technology	+3.43%
AAPL	Information Technology	+0.14%
GOOGL	Communication Services	+3.20%
NFLX	Communication Services	+1.99%
VZ	Communication Services	+0.18%
LLY	Health Care	+15.43%
JNJ	Health Care	-29.31%
MRK	Health Care	+1.11%
BRK.B	Financials	+20.29%
JPM	Financials	+6.10%
BAC	Financials	+2.30%
AMZN	Consumer Discretionary	-1.54%
HD	Consumer Discretionary	-2.75%
MCD	Consumer Discretionary	+22.88%
WMT	Consumer Staples	+4.93%
COST	Consumer Staples	+8.51%
PG	Consumer Staples	+5.51%
XOM	Energy	+12.67%
CVX	Energy	-23.95%
COP	Energy	+12.09%
GE	Industrials	+0.93%
CAT	Industrials	-4.10%
BA	Industrials	-2.38%
PLD	Real Estate	-9.47%
AMT	Real Estate	+6.26%
WELL	Real Estate	-0.01%
SHW	Materials	+3.61%
NEM	Materials	+5.11%
FCX	Materials	+0.93%
NEE	Utilities	+6.18%
SO	Utilities	+14.99%
DUK	Utilities	+13.42%

Ticker	GICS Sector	Weight
Total		+100.00%

Sector-level totals (unconstrained MVO, post-crisis):

GICS Sector	Unconstrained MVO Weight
Information Technology	4.86%
Communication Services	5.37%
Health Care	-12.76%
Financials	28.69%
Consumer Discretionary	18.60%
Consumer Staples	18.95%
Energy	0.82%
Industrials	-5.55%
Real Estate	-3.21%
Materials	9.65%
Utilities	34.59%
Total	100.00%

Table A9: Unconstrained MVO weights, recent period (Jan 2020 to Dec 2025)

Unconstrained MVO weights for the recent period. Short positions taken during the COVID-19 crash in early 2020 contribute to the smallest drawdown of any portfolio in this sub-period at -8.10%.

Ticker	GICS Sector	Weight
NVDA	Information Technology	-7.10%
MSFT	Information Technology	-3.99%
AAPL	Information Technology	+4.90%
GOOGL	Communication Services	+10.93%
NFLX	Communication Services	-5.55%
VZ	Communication Services	+7.86%
LLY	Health Care	-5.07%
JNJ	Health Care	+36.81%
MRK	Health Care	+0.53%
BRK.B	Financials	-5.17%
JPM	Financials	+21.91%
BAC	Financials	-21.35%
AMZN	Consumer Discretionary	+16.51%
HD	Consumer Discretionary	-0.10%
MCD	Consumer Discretionary	+7.68%
WMT	Consumer Staples	+11.61%
COST	Consumer Staples	+18.46%
PG	Consumer Staples	+20.74%
XOM	Energy	+30.56%
CVX	Energy	-22.25%
COP	Energy	+1.00%

Ticker	GICS Sector	Weight
GE	Industrials	+0.80%
CAT	Industrials	+0.32%
BA	Industrials	+4.63%
PLD	Real Estate	-16.44%
AMT	Real Estate	+10.11%
WELL	Real Estate	-22.40%
SHW	Materials	-13.81%
NEM	Materials	-4.13%
FCX	Materials	-1.22%
NEE	Utilities	+3.66%
SO	Utilities	-14.77%
DUK	Utilities	+34.31%
Total		+100.00%

Sector-level totals (unconstrained MVO, recent):

GICS Sector	Unconstrained MVO Weight
Information Technology	-6.19%
Communication Services	13.24%
Health Care	32.26%
Financials	-4.60%
Consumer Discretionary	24.09%
Consumer Staples	50.81%
Energy	9.31%
Industrials	5.76%
Real Estate	-28.72%
Materials	-19.17%
Utilities	23.20%
Total	100.00%

Table A10: Excel Solver setup and convergence

Eight optimizations are reported: five constrained MVO optimizations and three unconstrained MVO optimizations. All optimizations were solved using Microsoft Excel's Solver add-in with the Generalized Reduced Gradient (GRG) Nonlinear method. The objective in every case is to minimize portfolio variance, $w'\Sigma w$. The full-investment constraint $w'1 = 1$, is imposed in all optimizations. For constrained MVO, the non-negativity restriction, $w_i \geq 0$, is also imposed for all 33 assets.

Optimization	Engine	Convergence	Final Variance	Status
Constrained MVO, full period	GRG Nonlinear	0.0001	0.0109	Solver found a solution
Constrained MVO, pre-crisis	GRG Nonlinear	0.0001	0.0018	Solver found a solution
Constrained MVO, crisis	GRG Nonlinear	0.0001	0.0147	Solver found a solution

Optimization	Engine	Convergence	Final Variance	Status
Constrained MVO, post-crisis	GRG Nonlinear	0.0001	0.0062	Solver found a solution
Constrained MVO, recent	GRG Nonlinear	0.0001	0.0148	Solver found a solution
Unconstrained MVO, full period	GRG Nonlinear	0.0001	0.0094	Solver found a solution
Unconstrained MVO, post-crisis	GRG Nonlinear	0.0001	0.0048	Solver found a solution
Unconstrained MVO, recent	GRG Nonlinear	0.0001	0.0094	Solver found a solution

Solver options used in every case were: maximum time unlimited, convergence tolerance 0.0001, and maximum subproblems unlimited. The full-investment constraint is imposed in all optimizations. For constrained MVO, the non-negativity restriction binds only for assets assigned zero weight.

The unconstrained MVO is excluded from the pre-crisis and crisis periods because the covariance matrix is estimated from only 36 and 24 monthly observations, respectively. In the crisis period, the 33-by-33 covariance matrix is estimated from only 24 monthly observations, which makes the unconstrained solution unstable and economically unreliable. Excel produces extreme leveraged weights, including a Utilities weight of 146.14%, so the result is excluded rather than interpreted as a valid portfolio. The full-period unconstrained result is estimated over a window that includes those same 24 crisis months. A larger sample reduces the problem but does not eliminate it.

Table A11: Sub-period boundary sensitivity for the crisis period

The crisis sub-period ends in December 2009 rather than at the NBER cycle trough of June 2009 to maintain a clean calendar year. This table reports the constrained MVO maximum drawdown under both definitions.

Period Definition	Observations	Constrained MVO MaxDD	1/N MaxDD	S&P 500 MaxDD
Jan 2008 to Dec 2009 (used in thesis)	24	-17.94%	-39.24%	-46.42%
Jan 2008 to Jun 2009 (NBER trough)	18	-17.94%	-39.24%	-46.42%

The maximum drawdown is unchanged under both definitions because the worst cumulative peak-to-trough losses occur before June 2009, which falls inside both windows. Extending the crisis period to December 2009 changes the number of observations, but it does not change the maximum drawdown comparison.