Journal of
Risk and Financial Management

MDPI

Article

Practical Improvements to Mean-Variance Optimization for Multi-Asset Class Portfolios

Marin Lolic

Independent Researcher, Baltimore, MD 21210, USA; marin.lolic@gmail.com

Abstract: In the more than 70 years since Markowitz introduced mean-variance optimization for portfolio construction, academics and practitioners have documented numerous weaknesses in the approach. In this paper, we propose two easily understandable improvements to mean-variance optimization in the context of multi-asset class portfolios, each of which provides less extreme and more stable portfolio weights. The first method sacrifices a small amount of expected optimality for reduced weight concentration, while the second method randomly resamples the available assets. Additionally, we develop a process for testing the performance of portfolio construction approaches on simulated data assuming variable degrees of forecasting skill. Finally, we show that the improved methods achieve better out-of-sample risk-adjusted returns than standard mean-variance optimization for realistic investor skill levels.

Keywords: portfolio optimization; mean-variance optimization; portfolio theory; asset allocation

check for updates

Citation: Lolic, Marin. 2024. Practical Improvements to Mean-Variance Optimization for Multi-Asset Class Portfolios. *Journal of Risk and Financial Management* 17: 183. https://doi.org/10.3390/jrfm17050183

Academic Editors: Nalin Chanaka Edirisinghe and Jaehwan Jeong

Received: 8 April 2024
Revised: 24 April 2024
Accepted: 26 April 2024
Published: 29 April 2024

1. Introduction

Modern portfolio theory originated with the introduction of mean-variance optimization (MVO) by Harry Markowitz (1952). By formulating the question of optimal portfolio selection as a quadratic optimization exercise, Markowitz reduced a complex problem to the estimation of expected returns and covariances for investable assets. While his solution was conceptually and mathematically elegant, practical problems in its implementation became clear over time. Michaud (1989) and others noted that MVO often results in portfolios with dramatically unequal position sizes, with many assets reduced to zero weight, while a few assets take on very large weights. Green and Hollifield (1992) analyzed the sensitivity of MVO solutions to even modest changes in expected returns or the covariance matrix. This lack of robustness is especially troubling given the known difficulty of estimating forward-looking measures, particularly forward-looking returns (Merton 1980). Finally, research has established that MVO portfolios often perform very poorly in practice, even underperforming naive portfolio construction methods like equal weighting (Kan and Zhou 2007). Spurred by these realizations, many institutional investors have sought alternatives to portfolio optimization (Hurst et al. 2010) or have attempted to correct it in various ways.

Early proposals for MVO improvement, such as those of Frost and Savarino (1988), added constraints to portfolio weights beyond the standard requirement that all weights sum to 100% and the common avoidance of negative weights (short-selling). While added constraints can improve on the weaknesses of MVO, asset-level or asset class-level constraints are inherently arbitrary and portfolio-specific, making them difficult to generalize. Black and Litterman (1992) applied a Bayesian approach to the problem, seeking to merge an investor's expectations with market-implied figures. However, this requires the existence of a market portfolio, and especially the ability to identify such a portfolio. Jorion (1986) showed the value of "shrinking" estimates of the expected return vector to increase stability; other shrinkage methods have focused on the covariance matrix, imposing structure beyond the sample covariances taken from historical returns (Ledoit and Wolf 2003). The main complication for any shrinkage technique is the need to specify a target and a shrinkage intensity a priori. Michaud (1998) suggested the use of resampled portfolios, where the final weights would be an average of numerous portfolio optimizations, each drawn from a distribution defined by expected returns and covariances. While this approach has merit, it produces a single output portfolio without the ability to adjust the balance between optimality and robustness.

The field of robust optimization has also taken an interest in improving MVO by employing uncertainty sets (Tutuncu and Koenig 2004). In this approach, the analyst solves for the weights that produce the best outcomes under the worst realizations of returns and covariances. More recently, statistical and machine learning methods have introduced the idea of regularization to portfolio optimization (Brodie et al. 2009; Carrasco and Noumon 2010). Simply put, regularization imposes an explicit or implicit penalty on large target weights, often driving optimization toward more general solutions that should perform better across a range of outcomes (for an overview, see Tian and Zhang 2022). In the utility maximization context, this can mean adding an extra term to the objective function to discourage large weights in any single asset, or it can involve more complex reformulations of the optimization problem (Ban et al. 2018). While all of the aforementioned MVO improvements have their merits, few have achieved the combination of simplicity and results necessary to become popular among practicing portfolio managers.

We propose two intuitive methods of regularized portfolio optimization in the context of multi-asset portfolios. In both, the key idea is to approximate mean-variance optimality while reducing the concentration of the portfolio—simply defined as the sum of squared weights. The first method starts by running standard mean-variance optimization to obtain the maximum possible portfolio utility. It then performs another optimization, minimizing concentration while keeping utility at or above a preset percentage of the maximum possible utility. The second method resamples the assets across iterations, analogous to a scenario where only a certain number of assets are available for investment in a given iteration. Each method allows for adjustable levels of regularization, and each noticeably reduces portfolio concentration. Finally, we test the performance of both methods using a novel approach. We generate thousands of hypothetical return series for all available assets using a multivariate Gaussian distribution, then set expected returns and covariances as a weighted average of past values and "future" values. By toggling the weight assigned to past versus future values, we simulate different forecasting skill levels, allowing us to demonstrate the potential benefits of the two methods. We show that for all but the most implausibly skilled investors, some degree of regularization in portfolio optimization will improve risk-adjusted returns. We also show that both methods reduce the volatility in portfolio weights as optimization inputs change.

2. Background

The canonical MVO problem is given in matrix form by

$$\begin{aligned} \max \quad & \boldsymbol{\mu}^{\mathbf{T}}\mathbf{x} - \tfrac{1}{2}\gamma\mathbf{x}^{\mathbf{T}}\mathbf{V}\mathbf{x} \\ \text{s.t.} \quad & \mathbf{1}^{\mathbf{T}}\mathbf{x} = 1 \\ & \mathbf{x} \geq \mathbf{0} \end{aligned} \tag{1}$$

where $\boldsymbol{\mu}$ is the column vector of expected returns, $\mathbf{x}$ is the column vector of weights, γ is the risk aversion parameter, and $\mathbf{V}$ is the covariance matrix (Cornuejols et al. 2018). The first constraint ensures that the portfolio weights sum to one. For most investors, the second constraint applies and mandates that there are no negative weights (i.e., no short selling). MVO requires estimating expected returns and covariances among all assets, as well as knowing the investor's level of risk aversion. Following Paravisini et al. (2010), we will assume an average level of investor risk aversion throughout this paper, with γ set to 3.

To illustrate the workings of MVO and our proposed regularization methods, we will utilize a set of self-created capital markets assumptions (CMAs) detailed in Appendix A. Using standard MVO with no short-selling, we obtain the weights shown in Table 1.

Table 1. Optimal MVO Weights.

Asset	**Weight**
US LC Eq	0.0%
US SC Eq	0.0%
Non-US Dev Eq	43.8%
EM Eq	0.0%
Cash	0.0%
US Agg	15.8%
Non-US Agg	0.0%
US HY	0.0%
US FR	40.4%
REITs	0.0%

Notes: Portfolio expected return is 8.28%, expected standard deviation is 9.84%, and expected Sharpe ratio is 0.44.

One feature becomes apparent immediately: the portfolio is highly concentrated in only three assets, with the other seven assets receiving no allocations at all. This is despite relatively modest differences in expected returns and covariances among the ten total assets. Less obvious is how sensitive these estimates are to small changes. Table 2 demonstrates that adjusting the expected return of Non-US Developed Equities up or down by only 50 bps can lead to swings of over 15% in total portfolio weights. Many asset owners rightly question the wisdom of a process that produces such extreme weights and is so reactive to minor shifts in assumptions.

Table 2. Optimal MVO weights as expected return of Non-US Developed Equities changes.

Asset	**−50 bps**	**Initial**	**+50 bps**
US LC Eq	0.0%	0.0%	0.0%
US SC Eq	0.0%	0.0%	0.0%
Non-US Dev Eq	35.4%	43.8%	50.6%
EM Eq	2.4%	0.0%	0.0%
Cash	0.0%	0.0%	0.0%
US Agg	18.1%	15.8%	14.0%
Non-US Agg	0.0%	0.0%	0.0%
US HY	0.0%	0.0%	0.0%
US FR	44.1%	40.4%	35.3%
REITs	0.0%	0.0%	0.0%

Notes: For −50 bps, portfolio expected return is 7.89%, expected standard deviation is 9.19%, and expected Sharpe ratio is 0.42. For initial, portfolio expected return is 8.28%, expected standard deviation is 9.84%, and expected Sharpe ratio is 0.44. For +50 bps, portfolio expected return is 8.74%, expected standard deviation is 10.56%, and expected Sharpe ratio is 0.45.

Besides the inherent difficulty of predicting the future, there is an additional reason the weights are so unstable: very different weights can often produce similar utility. In Figure 1 below, we have constructed a simple two-asset portfolio. Asset A has an expected return of 8% with an expected standard deviation of 16%, while Asset B has an expected return of 4% with an expected standard deviation of 8%. The expected correlation between the assets is 0.2, and we assume an investor risk aversion of 3. While the maximum utility of 4.81% is achieved when Asset A has a weight of 62% (and Asset B's weight is thus 38%), the portfolio can achieve at least 98% of the maximum utility with Asset A weights of between 47% and 77%. Though the exact results depend on the expectations, most realistic capital markets assumptions will induce similar outcomes. This insight leads directly to the first method we propose.

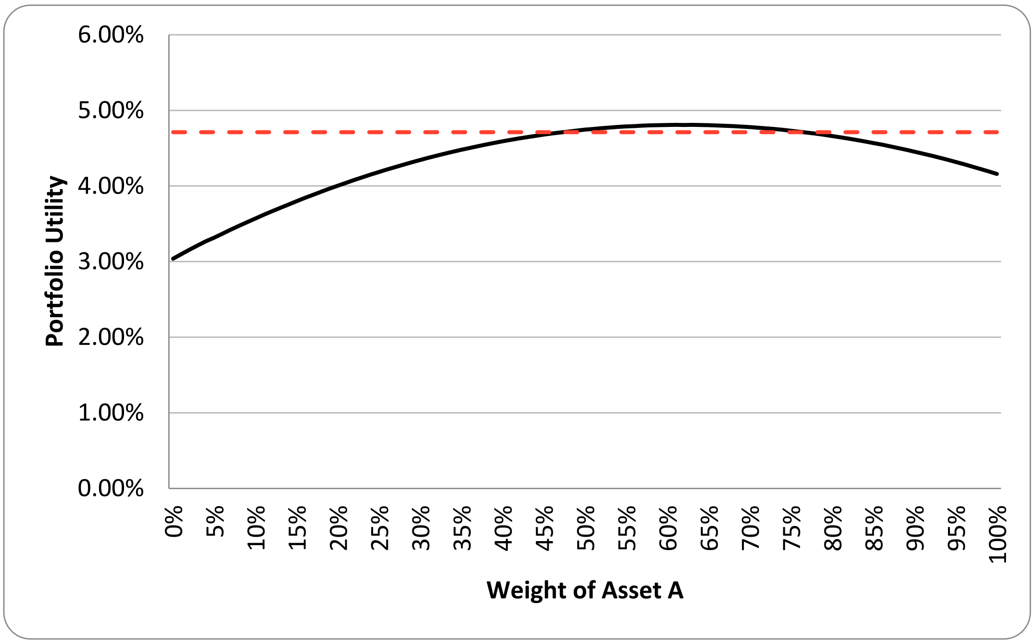

Figure 1. Portfolio utility as a function of Asset A weight. The black line plots portfolio utility, while the dashed red line is set at 98% of maximum utility.

3. Materials and Methods

3.1. Method One: Near-Optimality

If the utility function changes gradually over a wide range of weights, then it is possible to construct portfolios that are nearly optimal while minimizing concentration (the sum of squared weights). Corvalan (2005) considered a related concept, though the method we propose requires setting fewer parameters. Method One thus has two steps. In the first step, we optimize a portfolio using standard MVO as seen in (1) above; this establishes the maximum possible ex ante utility. In the second step, we minimize concentration while constraining portfolio utility to be greater than or equal to a preset percentage of the maximum possible ex ante figure. Other constraints can apply as well, and the full system is shown in (2) below.

$$\begin{aligned} \min \quad & \mathbf{x}^{\mathbf{T}}\mathbf{x} \\ \text{s.t.} \quad & \boldsymbol{\mu}^{\mathbf{T}}\mathbf{x} - \tfrac{1}{2}\gamma\mathbf{x}^{\mathbf{T}}\mathbf{V}\mathbf{x} \geq \theta\varepsilon \\ & \mathbf{1}^{\mathbf{T}}\mathbf{x} = 1 \\ & \mathbf{x} \geq \mathbf{0} \end{aligned} \tag{2}$$

For this second step, ε is the optimal ex ante utility from step one, while θ is the floor for total portfolio utility expressed as a decimal. Much like the risk aversion parameter, θ must be set by the analyst performing the optimization. A value of 1.00 for θ will result in the standard MVO portfolio, while values below 1.00 will progressively converge to an equal-weight portfolio. The parameter θ thus controls the regularization of the optimization weights, with values less than 1.00 indicating a preference for reduced portfolio concentration. We can see this method in action in Table 3, where we compare the weights generated by Method One to those of MVO. Clearly, the near-optimal portfolio is far more diversified. Later, we will demonstrate that Method One also reduces sensitivity to estimation error, thus making portfolios more robust. It can also improve risk-adjusted performance compared to MVO.

Table 3. MVO weights compared to Method One weights (θ = 0.95).

Asset	MVO	Method One
US LC Eq	0.0%	4.6%
US SC Eq	0.0%	4.6%
Non-US Dev Eq	43.8%	24.3%
EM Eq	0.0%	5.9%
Cash	0.0%	4.7%
US Agg	15.8%	15.5%
Non-US Agg	0.0%	5.2%
US HY	0.0%	8.3%
US FR	40.4%	22.4%
REITs	0.0%	4.5%

Notes: For MVO, portfolio expected return is 8.28%, expected standard deviation is 9.84%, and expected Sharpe ratio is 0.44. For Method One, portfolio expected return is 7.69%, expected standard deviation is 8.97%, and expected Sharpe ratio is 0.41.

Readers may note that the optimization in (2) is no longer strictly quadratic due to the first constraint being non-linear. Fortunately, such a quadratically constrained quadratic program (QCQP) is still convex because the covariance matrix **V** is positive semi-definite. Most optimization software supports QCQP, and possible solution techniques include interior-point methods and local search algorithms.

3.2. Method Two: Resampling Assets

One of the underlying assumptions of MVO is that there is an investable universe of assets from which to create a portfolio, and that this universe is constant. Method Two relaxes this assumption and applies a resampling approach to the universe of assets. More concretely, the method consists of a number of iterations chosen by the analyst, such as 1000 or 10,000. In each iteration, the method randomly chooses m assets out of the total n assets, and performs MVO on the subset m. Once all the iterations are complete, this method averages weights across all iterations to obtain the final weights. Figure 2 illustrates this process. Much like the number of iterations, m must be set by the analyst; this parameter controls the degree of regularization, with lower values of m corresponding to more regularization. If m is equal to n, the result is a standard MVO; if m is equal to one, the result is the equal-weight portfolio. This asset resampling approach draws inspiration from the random forest model in machine learning, where data features are resampled in the tree-building process. By limiting the optimization process to a subset of assets at each iteration, asset resampling forces the optimizer to consider investments that it would otherwise shun.

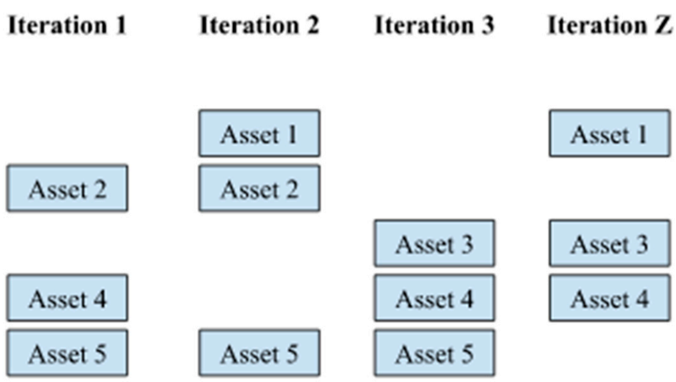

Figure 2. Schematic of Method Two with $n = 5$ and $m = 3$.

Table 4 presents the outcome of an intermediate value of *m* for the portfolio example we have used thus far. Much like Method One, resampling assets leads to lower concentration of asset weights compared to MVO. We will show later that it also leads to less turnover and better risk-adjusted performance compared to MVO.

Table 4. MVO weights compared to Method One weights (θ = 0.95) and Method Two weights (*m* = 5).

Asset	MVO	Method One	Method Two
US LC Eq	0.0%	4.6%	3.1%
US SC Eq	0.0%	4.6%	4.7%
Non-US Dev Eq	43.8%	24.3%	24.4%
EM Eq	0.0%	5.9%	7.1%
Cash	0.0%	4.7%	0.6%
US Agg	15.8%	15.5%	18.7%
Non-US Agg	0.0%	5.2%	5.5%
US HY	0.0%	8.3%	8.6%
US FR	40.4%	22.4%	27.0%
REITs	0.0%	4.5%	0.1%

Notes: For MVO, portfolio expected return is 8.28%, expected standard deviation is 9.84%, and expected Sharpe ratio is 0.44. For Method One, portfolio expected return is 7.69%, expected standard deviation is 8.97%, and expected Sharpe ratio is 0.41. For Method Two, portfolio expected return is 7.78%, expected standard deviation is 8.76%, and expected Sharpe ratio is 0.43.

4. Results

Arguably the best test of any portfolio construction method is how its created portfolios perform in terms of returns and volatility in an out-of-sample time period. To that end, we offer a way of testing the portfolio construction methods we have described, as well as almost any other approach to portfolio construction. The challenge of back-testing portfolio construction methodologies is that many require some ex ante assumptions, which can be difficult to formulate after the fact. In the case of MVO and the two methods we have proposed, tests require expected returns, expected volatilities and expected correlations.

We begin by simulating daily returns for all ten available assets using the annualized capital markets assumptions found in Appendix A. We draw the returns from a multivariate Gaussian distribution, though more complex and realistic methods could be used instead. We simulate 10,000 runs, each lasting 504 trading days (2 years) to obtain a large sample size of potential market activity. In each simulation run, we assume the position of an investor standing at the end of day 252, denoted as time *t*, that is the midpoint of the time series. While the investor knows the realized returns and covariances for days 1 through 252 (in-sample period), they are investing for days 253 through 504 (out-of-sample period), and, thus, must forecast those unseen values. For each of the 10,000 runs, the investor's return expectations at time *t* are a weighted average of the previous 252 days (R_P) and the forward 252 days (R_F), as seen in (3) below. The same relation holds for expected volatilities and expected correlations.

$$E(R)_t = (1 - \eta)R_P + \eta R_F \quad (3)$$

The parameter that controls the relative weights is η, which can be thought of as a measure of investor forecasting skill. In the case of an investor with no forecasting skill, η equals zero, and expectations are naively set to the realized values of the past 252 days; any minimal correspondence between past and future returns will be due to autocorrelation. If an investor is perfectly clairvoyant, η equals one, and expectations are the same as future 252-day realized values. The η value is thus conceptually similar to the information coefficient (IC) developed by Grinold and Kahn (2000), as each measures the correlation between predictions and outcomes.

For a given degree of regularization and level of forecasting skill, we create 10,000 portfolios (one for each simulation run) and store the weights. We also compute the realized return and volatility for each portfolio over the forward 252 days, as well as the Sharpe ratio. Finally, we record the out-of-sample median Sharpe ratios for all combinations of regularization and skill, which are shown in Tables 5 and 6 for Method One and Method Two, respectively. Each row of each table corresponds to a skill level, while the columns are degrees of regularization; note that the first column of each table is equivalent to MVO because it assumes no regularization. The median Sharpe ratios in each row are colored in green if they exceed the median MVO Sharpe for their row, while they are colored in red if they fall below.

Table 5. Method One median Sharpe ratios by regularization and skill. Left-most column corresponds to MVO.

		Minimum Utility (θ)					
		100%	**98%**	**96%**	**94%**	**92%**	**90%**
Skill (η)	**0.00**	0.26	0.27	0.29	0.30	0.31	0.31
	0.10	0.38	0.39	0.41	0.42	0.42	0.42
	0.20	0.53	0.53	0.54	0.54	0.54	0.54
	0.30	0.69	0.69	0.69	0.68	0.68	0.67
	0.40	0.86	0.86	0.85	0.84	0.83	0.81
	0.50	1.01	1.00	0.99	0.97	0.96	0.94

Notes: Numbers in each row are colored in green if they exceed the median MVO Sharpe for their row, while they are colored in red if they fall below.

Table 6. Method Two median Sharpe ratios by regularization and skill. Left-most column corresponds to MVO.

		Resampled Assets (m)					
		10	**9**	**8**	**7**	**6**	**5**
Skill (η)	**0.00**	0.26	0.30	0.30	0.31	0.32	0.32
	0.10	0.38	0.42	0.42	0.43	0.43	0.43
	0.20	0.53	0.55	0.55	0.55	0.55	0.54
	0.30	0.69	0.70	0.71	0.70	0.69	0.67
	0.40	0.86	0.87	0.86	0.86	0.85	0.82
	0.50	1.01	1.01	1.01	1.00	0.98	0.95

Notes: Numbers in each row are colored in green if they exceed the median MVO Sharpe for their row, while they are colored in red if they fall below.

Two results stand out in this analysis. First, regularization does appear to improve ex-post portfolio Sharpe ratios, especially at lower levels of forecasting skill. Only once forecasting skill (η) rises to or above 0.4 does MVO routinely outperform the regularized methods. Considering that an IC of even 0.2 corresponds to an impressive hit rate of 60%, it seems that the methods described could have wide applicability. Second, the two methods exhibit somewhat different patterns of Sharpe ratio change. Method One achieves improvements more gradually at lower skill levels than does Method 2, while Method 2 appears to be somewhat more robust at higher levels of skill. It is worth remembering that Method One explicitly solves for minimum concentration for a given level of utility, while Method Two is more of a heuristic that achieves broadly similar outcomes. Method One can also accommodate any constraints, while Method Two cannot guarantee the same due to its final step of averaging weights across iterations.

The weights calculated by each method also contain valuable information, showing the reduction in weight volatility as the degree of regularization increases. Figure 3a,b illustrate this for a subset of assets, while Tables 7 and 8 contain the results for the full asset lineup. All else equal, less volatility in weights leads to less portfolio turnover and thus reduced transaction costs. Less movement in portfolio weights over time should also mitigate harm from estimation error in forecasting asset returns and covariances.

(**a**) Bar colors denote values of θ

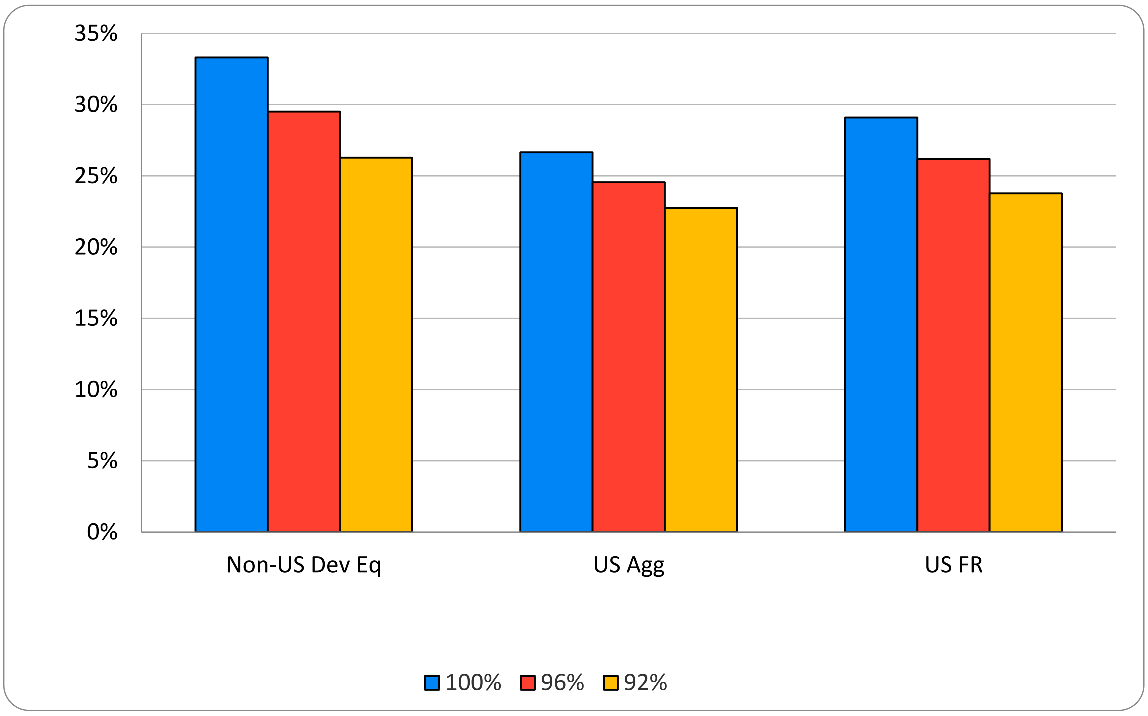

(**b**) Bar colors denote values of *m*

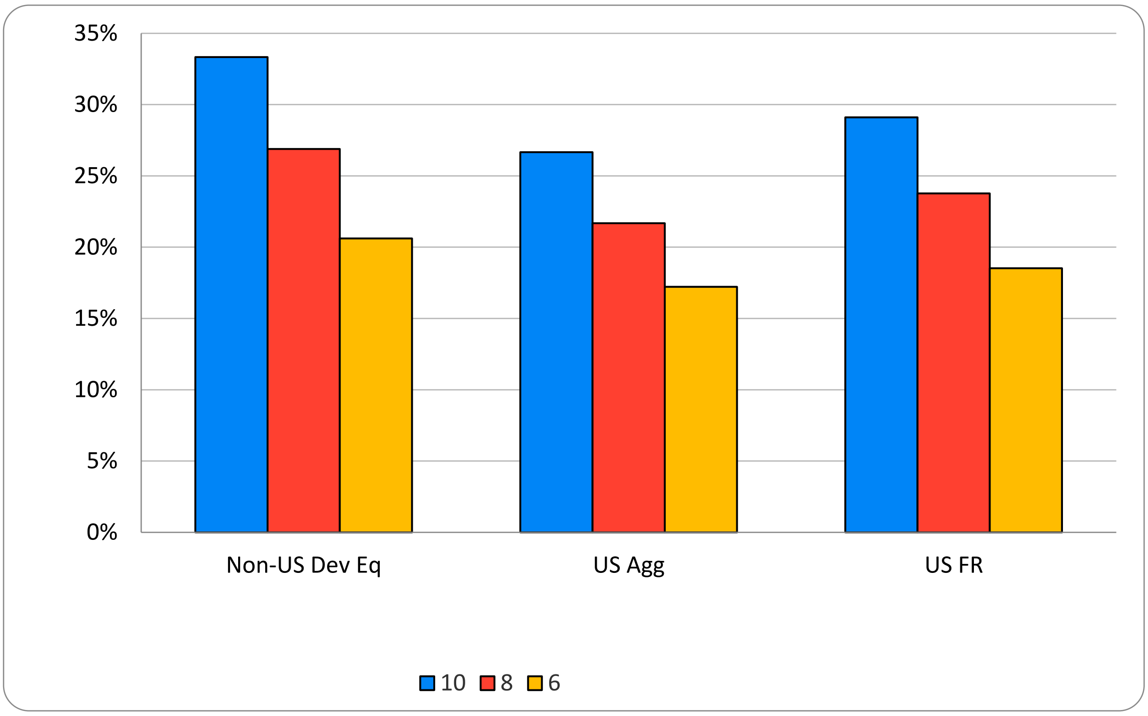

Figure 3. (**a**) Method One weight volatility by asset and regularization, $\eta = 0.15$. (**b**) Method Two weight volatility by asset and regularization, $\eta = 0.15$.

Table 7. Method One weight volatility by asset and regularization, η = 0.15. Left-most column corresponds to MVO.

	Minimum Utility (θ)					
	100%	**98%**	**96%**	**94%**	**92%**	**90%**
US LC Eq	22.4%	20.6%	19.1%	17.7%	16.4%	15.3%
US SC Eq	27.2%	25.0%	23.2%	21.6%	20.1%	18.7%
Non-US Dev Eq	33.3%	31.3%	29.5%	27.9%	26.3%	24.8%
EM Eq	32.0%	29.9%	28.1%	26.4%	24.7%	23.2%
Cash	18.0%	17.4%	16.9%	16.5%	16.0%	15.6%
US Agg	26.6%	25.5%	24.6%	23.7%	22.8%	21.9%
Non-US Agg	26.2%	25.1%	24.2%	23.3%	22.4%	21.4%
US HY	29.2%	27.6%	26.2%	24.9%	23.6%	22.3%
US FR	29.1%	27.5%	26.2%	24.9%	23.8%	22.6%
REITs	22.8%	21.1%	19.7%	18.4%	17.1%	16.0%

Table 8. Method Two weight volatility by asset and regularization, η = 0.15. Left-most column corresponds to MVO.

	Resampled Assets (m)					
	10	**9**	**8**	**7**	**6**	**5**
US LC Eq	22.4%	20.4%	18.6%	16.9%	15.2%	13.4%
US SC Eq	27.2%	24.6%	22.2%	19.8%	17.5%	15.1%
Non-US Dev Eq	33.3%	30.0%	26.9%	23.7%	20.6%	17.4%
EM Eq	32.0%	28.9%	26.0%	23.0%	20.0%	17.0%
Cash	18.0%	16.5%	15.1%	13.9%	12.7%	11.6%
US Agg	26.7%	24.1%	21.7%	19.4%	17.2%	14.9%
Non-US Agg	26.2%	23.8%	21.4%	19.2%	16.9%	14.6%
US HY	29.2%	26.4%	23.8%	21.2%	18.6%	15.8%
US FR	29.1%	26.4%	23.8%	21.1%	18.5%	15.8%
REITs	22.8%	20.8%	18.9%	17.1%	15.3%	13.4%

5. Discussion and Concluding Remarks

There are obvious practical implications for the ideas we have presented. While the results of a study can never guarantee future outcomes, our work suggests that professional investors managing institutional multi-asset class portfolios can make use of either regularization method to build portfolios that are more diversified, more robust to changing conditions, and possibly more efficient than those produced by classic MVO. The choice of which method to employ will likely depend on how many constraints are needed for the portfolio construction process, with Method One generally being more accommodating of constraints beyond those in Equation (1). Both methods allow portfolio managers to tailor the degree of regularization to their assumed skill level, as humbling as such an exercise might be. Just as importantly, portfolio managers can explain the logic of these methods to their clients without resorting to complex mathematics. Asset owners may also be drawn to the aesthetic aspects of regularized optimization, as portfolios will contain fewer extreme weights than in MVO. But practitioners should not be the only beneficiaries of this work. Finance is both an applied and an academic discipline, and the latter community can continue to develop these ideas. As always, more research is needed and welcomed.

One direction is to adapt the regularized optimization methods to single-asset class portfolios, possibly in combination with benchmark-relative performance metrics. This paper has argued for the benefits of regularized optimization in a common multi-asset setting, but we cannot guarantee that our methods produce comparable results in all investment universes (e.g., within emerging markets or among individual small-cap equities). A second option is to extend the two methods presented to the case of multi-period portfolio optimization. A more computationally intensive study could measure the effects of different volatility and correlation regimes on the results we have obtained. This might be combined with more realistic simulated returns by drawing from non-normal distributions and conditional correlation assumptions. One could also separately model skill in forecasting future returns from skill in forecasting future covariances, as the former task has generally proven much more difficult than the latter. Yet another extension may involve applying regularization to non-MVO methods of portfolio construction such as risk budgeting. Whatever the course of future research, we hope this paper has contributed both in terms of promising improvements to MVO and a method of testing various portfolio construction approaches.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are openly available at https://github.com/MarinLolic/MVO-Improvement (accessed on 7 April 2024).

Conflicts of Interest: The author declares no conflicts of interest.

Appendix A. Capital Markets Assumptions

Table A1. Assets, expected returns and expected standard deviations.

	Short Name	Long Name	Representative Index	Expected Return	Expected Standard Deviation
A	US LC Eq	United States Large Cap Equity	Russell 1000	8.5%	17.0%
B	US SC Eq	United States Small Cap Equity	Russell 2000	9.0%	19.0%
C	Non-US Dev Eq	Non-United States Developed Equity	MSCI World ex-US	10.0%	17.0%
D	EM Eq	Emerging Markets Equity	MSCI Emerging Markets	9.5%	20.0%
E	Cash	United States Cash	Merrill Lynch 3-Month US Treasury Bill	4.0%	0.0%
F	US Agg	United States Aggregate Bond	Bloomberg US Aggregate Bond	5.5%	4.0%
G	Non-US Agg	Non-United States Aggregate Bond	Bloomberg Global Aggregate ex-USD	5.0%	4.5%
H	US HY	United States High Yield Bond	Bloomberg US Corporate High Yield	7.0%	11.0%
I	US FR	United States Floating Rate Bond	Credit Suisse Leveraged Loan Index	7.5%	10.0%
J	REITs	United States Real Estate Investment Trusts	NAREIT US Real Estate	7.5%	17.0%

Table A2. Expected correlations.

	A	B	C	D	E	F	G	H	I	J
A	1.00	0.90	0.70	0.80	0.00	0.00	0.00	0.40	0.50	0.80
B	0.90	1.00	0.70	0.75	0.00	0.00	0.00	0.40	0.50	0.85
C	0.70	0.70	1.00	0.70	0.00	0.00	0.00	0.50	0.40	0.70
D	0.80	0.75	0.70	1.00	0.00	0.00	0.00	0.40	0.50	0.70
E	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00
F	0.00	0.00	0.00	0.00	0.00	1.00	0.40	0.40	0.10	0.20
G	0.00	0.00	0.00	0.00	0.00	0.40	1.00	0.30	0.10	0.20
H	0.40	0.40	0.50	0.40	0.00	0.40	0.30	1.00	0.50	0.40
I	0.50	0.50	0.40	0.50	0.00	0.10	0.10	0.50	1.00	0.50
J	0.80	0.85	0.70	0.70	0.00	0.20	0.20	0.40	0.50	1.00

References

Ban, Gah-Yi, Noureddine El Karoui, and Andrew E. B. Lim. 2018. Machine learning and portfolio optimization. *Management Science* 64: 1136–54. [CrossRef]

Black, Fischer, and Robert Litterman. 1992. Global portfolio optimization. *Financial Analysts Journal* 48: 28–43. [CrossRef]

Brodie, Joshua, Ingrid Daubechies, Christine De Mol, Domenico Giannone, and Ignace Loris. 2009. Sparse and stable Markowitz portfolios. *Proceedings of the National Academy of Sciences of the United States of America* 106: 12267–72. [CrossRef] [PubMed]

Carrasco, Marine, and Neree Noumon. 2010. Optimal Portfolio Selection Using Regularization. Working Paper. Available online: https://www.eco.uc3m.es/temp/port8.pdf (accessed on 5 January 2024).

Cornuejols, Gerard, Javier Pena, and Reha Tutuncu. 2018. *Optimization Methods in Finance*, 2nd ed. Cambridge: Cambridge University Press, pp. 91–95.

Corvalan, Alejandro. 2005. *Well Diversified Efficient Portfolios*. Working Papers No. 336. Santiago: Central Bank of Chile.

Frost, Peter A., and James E. Savarino. 1988. For better performance: Constrain portfolio weights. *Journal of Portfolio Management* 15: 29–34. [CrossRef]

Green, Richard C., and Burton Hollifield. 1992. When will mean-variance efficient portfolios be well diversified? *Journal of Finance* 47: 1785–809.

Grinold, Richard C., and Ronald N. Kahn. 2000. *Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk*, 2nd ed. New York: McGraw-Hill, pp. 147–62.

Hurst, Brian, Bryan W. Johnson, and Yao Hua Ooi. 2010. Understanding Risk Parity. White Paper. Available online: https://www.aqr.com/White-Papers/Understanding-Risk-Parity.pdf (accessed on 7 April 2024).

Jorion, Philippe. 1986. Bayes-Stein estimation for portfolio analysis. *Journal of Financial and Quantitative Analysis* 21: 279–92. [CrossRef]

Kan, Raymond, and Guofu Zhou. 2007. Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis* 42: 621–56. [CrossRef]

Ledoit, Olivier, and Michael Wolf. 2003. Improved estimation of the covariance matrix of stock returns with an application of portfolio selection. *Journal of Empirical Finance* 10: 603–21. [CrossRef]

Markowitz, Harry M. 1952. Portfolio selection. *Journal of Finance* 7: 77–91.

Merton, Robert C. 1980. On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics* 8: 323–61. [CrossRef]

Michaud, Richard O. 1989. The Markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal* 45: 31–42.

Michaud, Richard O. 1998. *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation*. Boston: Harvard Business School Press, pp. 49–62.

Paravisini, Daniel, Veronica Rappoport, and Enrichetta Ravina. 2010. *Risk Aversion and Wealth: Evidence from Person-to-Person Lending Portfolios*. NBER Working Paper No. 16063. Cambridge, MA: National Bureau of Economic Research.

Tian, Yingjie, and Yuqi Zhang. 2022. A comprehensive survey on regularization strategies in machine learning. *Information Fusion* 80: 146–66. [CrossRef]

Tutuncu, Reha H., and Mark Koenig. 2004. Robust asset allocation. *Annals of Operations Research* 132: 157–87. [CrossRef]