Revisiting methods for assessing parameter estimation error in empirical datasets

Benedikt Himbert, Julia Kapraun*

April 2017

Abstract

In this paper, we propose a robust methodology for assessing estimation error in empirical datasets. In contrast to recent findings, we show that established portfolio strategies outperform the equally weighted portfolio rule even in larger portfolios, despite increasing absolute estimation error. We find that, when short-sale constraints are introduced and in consideration of portfolios that do not rely on expected return estimates the relationship between portfolio size and measurement error, as previously hypothesised in literature, does generally hold. Measurement error in these strategies, however, is significantly reduced and their application to larger asset universes is yet beneficial. Finally, we discuss the usefulness of statistics to assess estimation error and propose an intuitive measure of return-loss due to "unfavourable estimation error" based on the downside deviation of the return distribution that quantifies the share of measurement error investors should ultimately be concerned about.

Keywords: Mean-Variance Portfolios, Parameter Estimation Error, Measurement Error, Downside Risk

*The authors work at WHU - Otto Beisheim School of Management, Burgplatz 2, 56179 Vallendar, Germany.

Corresponding author is Benedikt Himbert, benedikt.himbert@whu.edu, +49-(0)261-6509-393

1 Introduction

Markowitz's (1952) mean-variance optimal framework has become one of the, if not the most important benchmark model in financial research. But since it requires the estimation of both, expected return and the covariance matrix, portfolio strategies targeting mean-variance optimality are prone to parameter estimation error. This error occurs when the sample moments are poor estimates of the true parameters ('sampling error') or the distribution of returns is non-stationary over time. Misspecification in the historic (ex-post) estimates ultimately leads to poor out-of-sample performance as compared to the performance under ex-ante parameter knowledge (Chopra and Ziemba (1993); Broadie (1993)).

It is generally noticeable that the conditions affecting the magnitude of parameter estimation error have not been sufficiently exploited. Only recently, DeMiguel et al. (2009b), Kan and Zhou (2007) and Duchin and Levy (2009) also attempted to better understand the conditions that affect the degree of measurement error in Markowitz's mean-variance optimal strategy and its extensions. Based on their findings, we want to further examine the hypothesis that the degree of parameter estimation error in mean-variance optimal investment strategies depends on the number of assets in the portfolio they are applied to as well as on investors' risk aversion. More specifically, Kan and Zhou (2007) provide an analytical, closed-form solution for a utility-loss function of using the parameter estimates in finding the mean-variance optimal weights, rather than knowing the true parameters. They conclude that estimation error is higher and mean-variance optimal strategies perform relatively worse in larger portfolios as a result of estimation error. This is empirically supported by findings of DeMiguel et al. (2009b), who claim that no mean-variance optimal strategy will on average outperform the equally weighted diversification that is free from estimation error. However, such solution is not available for short-sale constrained portfolio problems as it involves solving a non-linear quadratic function. In fact, Duchin and Levy (2009) earlier find the short-sale constrained mean-variance optimal strategy to perform relatively better in larger portfolios. To arrive at these highly contradictory findings, researchers have so far relied on methods in which the portfolios of different size are of more or less entirely dissimilar composition. It may seem quite obvious that no meaningful relationship between portfolio size and strategy performance can and should be inferred from the comparison of entirely different portfolio constituents. DeMiguel et al. (2009b), for example, compare the performance of investment strategies applied to a universe that consist of 20 Fama French portfolios with a universe of 8 MSCI country indices to assess the impact of increasing the size of the opportunity set (portfolio size).[1] Duchin and Levy (2009) fix their asset universe to the 30 Fama French industry portfolios and assess the impact of increasing portfolio size within the universe of the N=30 assets. Their methodology trades one flaw for another as "when $n < 30$, they consider the first n assets of the 30 Fama French industry portfolios". In other words, the choice of the "additional" asset is completely arbitrary and a robust estimation can neither be achieved. To remove this sampling bias and to resolve the prevailing conflict in literature we introduce an improved methodology to assess the impact of parameter estimation risk in empirical datasets.

In the first part of this paper we review the most relevant literature and provide an introduction to the methods used to model estimation error. We then present our more robust methodology and apply it to the most commonly employed mean-variance investment strategies and their extensions. In doing so, we confirm that with increasing portfolio size the measurement error in strategies that rely on parameter estimates increases, but show that higher absolute estimation error does not necessarily lead to worse ex-ante performance relative to other investment strategies and the equal weighting investment rule. In this context, we point out that out-of-sample outpeformance is dependent on an investor's peformance evaluation method. Finally, we resolve the apparent conflict in results between DeMiguel et al. (2009b) and Duchin and Levy (2009) and propose an intutive measure of return-loss due to "unfavourable estimation error" based on the Sortino ratio.

[1] Where the portfolios are taken as assets.

2 A review of relevant literature and methods used to study estimation error

Harry Markowitz's (1952) single-period theory on the optimal portfolio weights that provide the best trade-off between the mean return and risk (as measured by variance) of a portfolio, is one of the most fundamental ideas of portfolio management. To describe it in line with our improved methodology, let $\mathbb{R}^N$ be the space of N excess return vectors of length T, where N is number of assets in the asset universe under consideration. Then $\mathbb{R}^n \subset \mathbb{R}^N$ is a space of n excess return vectors, where n is the size of a portfolio that consists of assets drawn from a universe of size N. Denote by r_t and r_{ft} the rates of return on n risky assets and the risk-free asset at time t (t=1,2,...,T). Then the excess returns (from here on "returns") at time t are defined as $R_t = r_t - r_{ft}\mathbf{1_n}$, where $\mathbf{1_n}$ is an nx1 vector of ones.

In a portfolio of n assets, let $\mu = (\mu_1, \mu_2, ...\mu_n)'$ be the vector of their respective mean returns and $w = (w_1, w_2, ..., w_n)'$ the vector of portfolio weights. It follows that any portfolio will have an expected return $w'\mu$ and variance $w'\Sigma w$, where Σ is the covariance matrix of asset returns. Given a target value μ_* for the mean return of a portfolio, Markowitz (1952) defines an efficient portfolio (w^*) as one that allocates wealth to its constituents subject to the solution of the following optimization problem:

$$arg\ \underset{w}{min}\ w'\Sigma w \quad s.t.\ w'\mu = \mu_*\ and\ \mathbf{1_n}'w = 1. \tag{1}$$

The constraint $\mathbf{1_n}'w = 1$ thereby ensures that the portfolio is fully invested. For the estimated portfolio weights to actually be the true optimal investment proportions, one must assume the true parameters μ and Σ to be known. In practice, however, μ and Σ are unobservable and need to be estimated from historical data. Under the assumption that the distribution of returns is i.i.d. normal, the sample mean return $\hat{\mu}$ and the covariance matrix $\hat{\Sigma}$, estimated over T observations, are then the best estimates of μ and Σ.[2]

[2]We assume T > n so that Σ is invertable.

$$\hat{\mu} = \frac{1}{T}\sum_{t=1}^{T} R_t \tag{2}$$

$$\hat{\Sigma} = \frac{1}{T-1}\sum_{t=1}^{T}(R_t - \hat{\mu})'(R_t - \hat{\mu}) \tag{3}$$

Lai et al. (2011) point out that these are also "method-of moments estimates without the assumption of normality and when the i.i.d. assumption is replaced by weak stationarity." In the classical "plug-in" method, the obtained sample estimates are then treated as if they were the true parameters and plugged into (1). A vast amount of research has, however, proven that replacing μ and Σ by their sample estimates may lead to poor portfolio performance in the following "out-of-sample" periods (T+1,T+2,...) as result of erroneous estimation (Jobson and Korkie (1981); Michaud (1989)). Chopra and Ziemba (1993) discuss the impact of misspecification in normally distributed portfolio selection problems and find that the largest misspecification is a result of errors in the mean returns that are about ten times as important as errors in the covariance matrix estimates. Jagannathan and Ma (2003) even claim that the "estimation error in the sample mean is so large that nothing much is lost in ignoring the mean all together".

Kan and Zhou (2007) analytically study the estimation error and derive an expression for the "out-of-sample loss due to employment of estimated parameters rather than the true moments". Following their study, we consider an investor who chooses the set of portfolio weights that maximises his utility defined as:

$$U(w) = w'\mu - \frac{\gamma}{2}w'\Sigma w. \tag{4}$$

Maximizing (4) may be understood as problem (1), with the addition of a constant γ (where $\gamma > 0$) that is the slope parameter and can be used as an indicator of an investor's risk aversion[3]. There are good reasons to optimise the mean-variance utility function instead of the Markowitz mean-variance problem (1). Bodnar et al. (2013) note that only investing into the utility maximizing portfolio investors will get the true mean-variance efficient portfolio. This is the case, since all solutions to this problem will lie on the upper part of the parabola of mean-variance efficient portfolios, while this does not always hold true under (1).

[3]Though it is generally not equal to the risk aversion coefficient as defined by Pratt (1964).

Further, it is known that the set of optimal weights that maximises (4) can be expressed as:

$$w^* = \frac{1}{\gamma}\mu\Sigma^{-1}. \tag{5}$$

Based on the optimal weights, w^*, we can estimate the corresponding maximum utility an investor could obtain by plugging (5) into (4):

$$U(w^*) = \frac{1}{2\gamma}\mu'\Sigma^{-1}\mu. \tag{6}$$

In pratice, w^* is not known, since μ and Σ are unkown and, thus, needs to be estimated as a function of the data ($\hat{w} = f(R_1, ..., R_T)$). For each set of portfolio weights, the realised out-of-sample mean and variance are then given by $\tilde{\mu} = \hat{w}'\mu$ and $\tilde{\sigma}^2 = \hat{w}'\Sigma\hat{w}$ respectively. These are used to express the expected utility, conditional on the weights being chosen as $\hat{w}$ (and the strategy $\hat{w}$ played infinitely many times).

$$E[\tilde{U}(\hat{w})] = \hat{w}'\mu - \frac{\gamma}{2}\hat{w}'\Sigma\hat{w} \tag{7}$$

The expected loss from using $\hat{w}$ rather than w^* can then be defined by the following expression:

$$E[L(w^*, \hat{w})] = U(w^*) - E[\tilde{U}(\hat{w})]. \tag{8}$$

Since our focus remains on the relationship between expected performance of the investment rules as a function of the portfolio size n, we are particularly interested in Kan and Zhou's derivation of an expression that relates the expected loss in out-of-sample performance to n, T, and γ. Note that when the parameters μ and Σ are estimated by $\hat{\mu}$ and $\hat{\Sigma}$ through (2) and (3), the vector of estimated optimal portfolio weights is given as:

$$\hat{w} = \frac{1}{\gamma}\hat{\mu}\hat{\Sigma}^{-1}. \tag{9}$$

Assuming that the distributions of returns R_t are jointly normal, $\hat{\mu}$ and $\hat{\Sigma}$ are independent and distributed as $\hat{\mu} \sim N(\mu, \Sigma/T)$ and $T\hat{\Sigma} \sim W_n(T-1, \Sigma)$. Following Kan and Zhou (2007) it can then be shown that, when T $>$ n+4, the following closed-form relationship holds:

$$E[L(w^*, \hat{w})] = (1-k)\frac{(\mu'\Sigma^{-1}\mu)}{2\gamma} + \frac{nT(T-2)}{2\gamma(T-n-1)(T-n-2)(T-n-4)} \tag{10}$$

$$where \quad k = \left(\frac{T}{(T-n-2)}\right)\left[2 - \frac{T(T-2)}{(T-n-1)(T-n-4)}\right].$$

We observe that as n increases, the loss increases, and as γ increases, the loss decreases. Note that this relationship as such can only be said to hold for short-sale unconstrained mean-variance optimal portfolios since there exists a linear, closed-form solution. Under the same assumptions and using Monte Carlo methods, DeMiguel et al. (2009b) empirically show that in an unconstrained setting, the naive equally weighted investment strategy (which does not rely on any parameter estimates) is "more likely to outperform (out-of-sample) when the portfolio size is large, because this improves the potential for diversification while at the same time increases the number of parameters to be estimated in mean-variance optimal strategies". They find evidence that naively diversified portfolios consistently have higher out-of-sample Sharpe ratios and certainty equivalent return than those constructed based on estimated moments of the return series. From simulated data, they further find that the estimation window needed for the sample-based mean-variance strategy to outperform the naive investment rule is around 3000 months for a portfolio with 25 assets. Results of Tu and Zhou (2011) are in line with these findings, as both studies show that increasing the estimation period favourably affects the performance of mean-variance based strategies. Given a practitioners approach to estimating the moments of the returns series over significantly shorter horizons, this leads one to the assumption that mean-variance optimal strategies should not outperform their heuristic counterparts on average. Our research is aimed to show whether this also holds for the minimum-variance and short-sale constrained derivations of the Markowitz mean-variance optimal problem that are obtained by solving a quadratic optimisation problem. The hypotheses to be tested may, thus, be summarized as follows:

Hypothesis 1.1: Estimation error in shortsale constrained mean-variance optimal (minimum-variance) portfolios increases with the size n of the asset universe they are applied to.

Hypothesis 1.2: Shortsale constrained mean-variance optimal (minium-variance) portfolio strategies perform relatively worse when applied to larger asset universes and consistently underperform the 1/n rule due to estimation error.

Hypothesis 2.1: Estimation error in shortsale constrained mean-variance optimal portfolios decreases with higher investor risk aversion γ.

Hypothesis 2.2: Assuming 2.1 holds, then estimation error in mean-variance optimal portfolios can be reduced by constraining the volatility of the ex-post optimal portfolio.

3 Towards a more robust estimation method

Our novel methodology is an extension of the "rolling-sample" approach that has become the standard model for assessing parameter estimation risk and considers a large number of samples over time. The algorithm can be described as follows:

1. Given a datasets of length M trading days of asset prices for n assets, choose an estimation window of length T where M > T and a rolling interval P.

2. Starting in T+1, estimate $\hat{\mu}$ and $\hat{\Sigma}$ from the sample return series over the previous T days for n assets. Based on the estimated parameters find the vector of estimated optimal portfolio weights ($\hat{w}$), subject to the chosen portfolio strategy.

3. "Hold" the n assets in proportion $\hat{w}$ over an observation ("out-of-sample") period of length O-trading days and estimate the daily returns over period T+1 to T+1+O.

4. Drop the earliest P days of returns and add the next P days of returns of the M-days period. Repeat steps 2-4.

Following through with this algorithm when P=O, it becomes a portfolio strategy with rebalacing frequency P and will ultimately yield a series of (M − T) daily out-of-sample returns that can be used for performance assessment.

In our study, we choose an estimation window of length T=900 days (that is 30 months of daily returns). While research has shown that longer estimation windows can significantly improve the performance of investment strategies when relying on parameter estimates, estimation windows of more than 60 months are rarely used in practice.[4] Following DeMiguel et al. (2009b) we employ a O=30 days observation window and P=30 days rolling interval. Other studies like Duchin and Levy (2009) only assess the performance over the following year. Whether extending the estimation and observation period is ultimately favourable, depends largely on the statisitical properties of the data. Broadie (1993) and Scherer (2007) point out that increasing the estimation period when the data is non-stationary may have an unfavourable impact on out-of-sample performance.

Definition 1: *We propose an extension of the rolling window estimation to all possible combinations of a portfolio of n assets, drawn from a sample of N assets. Hence, for a given portfolio size n there are $c_{n|N}$ ways to sample it from the universe of N assets, where $c_{n|N}$ is given by:*

$$c_{n|N} = \begin{pmatrix} N \\ n \end{pmatrix} = \frac{N!}{(N-n)!\,n!}. \tag{11}$$

This yields the $c_{n|N} \times n$ matrix of possible return series combinations that defines the space $\mathbb{R}^N$. To achieve a more robust estimate of the out-of-sample performance, we find the mean ex-ante return and variance across all $c_{n|N} \times n$ portfolio combinations. In the same way, we estimate all other performance statistics. Employing this methodology allows for a very robust estimation of the effect of increasing n on parameter estimation error as it fixes the sample universe and models all combinations an investor could choose within for a chosen portfolio size. Achieving this desirable degree of robustness, unfortunately, comes at the cost of high computational requirements. Our proposed methodology can be considered suitable for testing datasets, where N is smaller than 25. When the ratio of N/n becomes larger, $c_{n|N}$ grows exponentially and optimising portfolios across all combinations becomes difficult. Previous research in this field, however, rarely considers any N larger than 30. This makes our approach suitable for assessing the most commonly employed sample sizes.

In addition, we consider the case where the joint distribution of returns over M observations is known and employ the investment strategies on the latter. There is then no estimation error and the mean-variance optimal portfolio will outperform. Employing any other rule would be irrational. We, thus, label the corresponding returns as the performance under "true parameter knowledge".

[4] DeMiguel et al. (2009b)

Table 4.1: List of tested asset-allocation strategies

This table presents the employed asset-allocation strategies and their abbreviations that will be used in this paper. In our notation, a strategy can either be identified by its indicator variable z or its abbreviation.

z	**Asset-allocation strategy**	**Abbr.**
1	(Naive) equal weighting of all portfolio constituents	1/n or ew
2	Shortsale unconstrained mean-variance optimal portfolio	mvo
3	Shortsale constrained mean-variance optimal portfolio	mvo-c
4	Shortsale unconstrained minimum-variance portfolio	mv
5	Shortsale constrained minimum-variance portfolio	mv-c

4 Description of investment rules employed

In our study, we test the most commonly employed mean-variance investment strategies in their respective short-sale unconstrained and constrained forms against the naive diversification that does not rely on parameter estimates. A comprehensive summary of all tested strategies and their corresponding abbreviations is given in Table 4.1. To find the optimal portfolio weights, all mean-variance strategies rely on the estimated covariance matrix $\hat{\Sigma}$ of portfolio asset returns and, subject to the chosen strategy, directly on the estimated mean of the return series, $\hat{\mu}$, as well. These parameters are estimated over a rolling estimation period ("ex-post") of length T, from which the optimal weight of each asset is derived. The optimised investment proportions are expected to also be optimal over the following out-of-sample ("ex-ante") period. For this assumption to hold, the joint distribution of returns must be stable over time. Let $\mathring{\mu}$ and $\mathring{\Sigma}$ be the "true", unknown in-sample parameters estimated over the entire series of M-days. Then the ex-post estimated optimal portfolio weights $\hat{w}$ will only correspond to the ex-ante optimal portfolio weights $\mathring{w}$, when $\hat{\mu}$ and $\hat{\Sigma}$ are perfect estimators of $\mathring{\mu}$ and $\mathring{\Sigma}$.[5]

[5] To assess the effect of estimation error we assume that $\mathring{\mu}$ and $\mathring{\Sigma}$ constitute the "population parameters", i.e. $\mathring{\mu} = \mu$ and $\mathring{\Sigma} = \Sigma$. While these are estimated over a sufficiently long dataset, these are generally not perfect estimates of the unkown population moments and in themselves subject to estimation error.

4.1 Equally weighted (naively diversified) portfolios

When assigning an equal weight to each portfolio asset i (i=1, 2, ..., n), the vector w_{ew} is simply constituted by:

$$w_{ew} = (\frac{1}{n}, ..., \frac{1}{n})' \, \epsilon \, \mathbb{R}^n. \tag{12}$$

The mean return of the portfolio is then an equally weighted average of the returns of all n portfolio constituents. In estimating the portfolio weights, the equally weighted portfolio does not take the dependencies between portfolio assets into account and, therefore, does not rely on the estimated moments of the return series. When investors have no knowledge about the future distribution of returns or cannot rely on historic return distributions, the 1/n rule will be optimal. Despite its long history, the 1/n rule is still one of the most commonly used investment rules, especially by private investors who tend to hold a relatively smaller number of assets (Benartzi and Thaler (2001)).

4.2 (Unconstrained) sample based mean-variance optimal portfolios

The sample based mean-variance optimal portfolio that maximises expected investor utility can be understood as problem (4) when the investor relies on moment estimates. The optimization problem becomes to find the vector of optimal portfolio weights $\hat{w}_{mvo}$ so that:

$$\underset{w}{max} \; w'\hat{\mu} - \frac{\gamma}{2} w'\hat{\Sigma} w \quad s.t. \; \mathbf{1_n'} w = 1. \tag{13}$$

By assuming an explicit degree of risk aversion (γ) we can find a single mean-variance optimal portfolio that all investors with the same risk appetite should hold. Most of the models considered throughout this paper can be expressed as a derivation of (13), subject to their own constraints. Hence, it is important that the chosen utility function is valid and that results are not sensitive to misspecification of the utility function. In this context, Kallberg and Ziemba (1984) find that “utility functions with similar levels of Arrow-Pratt absolute risk aversion result in similar optimal portfolio weights, irrespective of the functional form of the utility.” We, therefore, deem the results of our analysis sufficiently robust against any misspecification.

As shown in (9) there is a linear solution to the optimisation problem available, yielding the tangency portfolio:

$$\hat{w}_{mvo} = \frac{1}{\gamma}\hat{\mu}\hat{\Sigma}^{-1}. \tag{14}$$

4.3 (Unconstrained) sample based minimum-variance portfolio

Chopra and Ziemba (1993) suggest that because forecasting returns is so difficult and estimation error so pronounced, portfolio outcomes (i.e. out-of-sample performance) could be improved by assuming that all stocks possess the same expected returns. This idea forms the basis for optimisation towards a minimum-variance portfolio objective. In the absence of short-sale constraints, solving the following optimisation problem yields the vector of optimal weights for the minimum-variance portfolio $\hat{w}_{mv}$:

$$arg\ \underset{w}{min}\ w'\hat{\Sigma}w \quad s.t.\ \mathbf{1_n'}w = 1. \tag{15}$$

It can be shown that there also exists a linear solution for the short-sale unconstrained minimum-variance portfolio (e.g. Kempf and Memmel (2006)) that is given by:

$$\hat{w}_{mv} = \frac{\hat{\Sigma}^{-1}\mathbf{1_n}}{\mathbf{1_n'}\hat{\Sigma}^{-1}\mathbf{1_n}}. \tag{16}$$

The underlying assumption of equal expected returns across all assets is stark but Jagannathan and Ma (2003), among others, provide empirical evidence that the minimum-variance portfolio often performs better out-of-sample than any other mean-variance optimal portfolio strategy. This suggests that estimating an additional moment of the return series, in this case the mean, adds significant parameter estimation risk. Kallberg and Ziemba (1984) discuss the impact of misspecification in normally distributed portfolio selection problems and find that the largest misspecification is a result of the errors in sample means that are about ten times as important as errors in the sample covariance matrix. This holds true for small ratios of n/T, where the estimation error in mean returns does account for most of the loss of out-of sample performance, as pointed out by Kan and Zhou (2007). For large n/T, however, the error in the covariance matrix becomes increasingly important.

4.4 Constrained sample based mean-variance optimal portfolio

When estimation error is not a concern, imposing constrains should hurt. Green and Hollifield (1992) show that mean-variance portfolios contain both extreme positive and negative weights. Introducing a short-sale constraint should then reduce the in-sample performance of the investment strategy. On the other hand, when estimation error is present, research has shown that by constraining the strategies we may be able to mitigate possible extreme differences between the true and the observed out-of-sample distribution of returns (Frost and Savarino (1988)).

The optimisation problem for finding the sample based mean-variance optimal portfolio corresponds to (13) with the additional constraint $w_i \geq 0$ (i=1,2,...,n), i.e. that the weight of each asset i must be larger (or equal to) zero. While the problem appears not much different from its short-sale unconstrained form, the solution of optimal weights can only be attained by solving a quadratic problem. Imposing the short-sale constraint yields the following Lagrangian function:

$$\mathcal{L} = w'\hat{\mu} - \frac{\gamma}{2}w'\hat{\Sigma}w + w'\lambda \quad s.t.\ \mathbf{1_n'}w = 1, \tag{17}$$

where $\lambda = (\lambda_1, ..., \lambda_n)' \epsilon \mathbb{R}^n$ is the vector of Lagrange multipliers for the shortsale constraint and thus $\lambda \geq 0$.[6] Hence, there is no linear, closed-form solution available. As a result, it is neither possible to derive an expression for the expected loss in out-of-sample performance due to parameter estimation error as in Kan and Zhou (2007) for this problem. The relationship of portfolio size, investor risk aversion and the observed out-of-sample loss can, therefore, not be deduced to an analytical solution but requires an empirical assessment.

4.5 Constrained sample based minimum-variance portfolio

Jagannathan and Ma (2003) provide an extensive assessment of the impact of imposing short-sale constraints on the performance of the minimum-variance portfolio. They show that short-sale constraining a minimum-variance portfolio can be understood as shrinking the elements of the covariance matrix of the unconstrained portfolio, where the sample covariance matrix is replaced by:

$$\hat{\Sigma}_{mv-c} = \hat{\Sigma}_{mv} - \lambda \mathbf{1}_n' - \mathbf{1}_n \lambda', \tag{18}$$

where λ, again, is the vector of Lagrange multipliers for the short-sale constraint. If the short-sale constraint for the i-th asset is binding (that is $w_i = 0$, (i=1,2,...,n)), then the covariance of the i-th

[6] $\lambda_i = 0\, if\, w_i > 0,\ i = 1, 2, ...n$

asset with another asset is shrunk by λ_i, the magnitude of the Lagrange multiplier associated with the short-sale constraint. If an estimated large covariance is due to estimation error, the shrinkage implied by the short-sale constraint leads to a better estimate of the true (population) covariance.[7] Jagannathan and Ma's finding, that when the non-negativity constraints are in place, there is not much further benefit from applying other shrinkage estimators, leads us to the decision to not incorporate shrinking methods in our work (though it could certainly be done in future research).

[7] For a more detailed discussion we refer the reader to Jagannathan and Ma (2003).

5 Methods for evaluating portfolio performance and estimation error

5.1 Sharpe ratios[8]

Practitioners mostly prefer to use the Sharpe ratio as a measure of a portfolio's performance. The out-of-sample Sharpe ratio of a strategy z (where z encompasses the space of strategies defined in Table 4.1) is defined as the ratio of sample mean and the standard deviation of the out-of-sample excess returns (denoted by $\tilde{\mu}_z$ and $\tilde{\sigma}_z$ respectively):

$$\widetilde{SR}_z = \frac{\tilde{\mu}_z}{\tilde{\sigma}_z}. \tag{19}$$

Reporting the Sharpe ratio revokes practitioners' interest. But when the objective function is given by (13), that is maximising expected utility, then Sharpe ratio may not be the appropriate measure. For once, the Sharpe ratio does not rely on an assumption of risk aversion and, therefore, largely differs from the objective function. Though, it is a well-known fact that when the true parameters are known, maximizing Sharpe ratio and the utility may be seen as equivalent, once the true parameters are unknown, the two are different (Tu and Zhou (2011)).

For an optimisation-based investment strategy to be of any value, it must provide an investor with a better return for the same level of risk than the naive diversification rule.

5.2 Certainty equivalent return

To compare the value of the investment strategies to an investor, we assume that investor preferences can be fully described by the mean and variance of a portfolio. At all times, the investor aims to

[8]Note that all of the statistics presented in the following are reported as averages across $c_{n|N}$ simluations.

maximise his expected utility for an investment strategy z. Maximising the utility therefore corresponds to finding the ex-ante mean-variance optimal portfolio for an investor with risk aversion coefficient γ. (20) can, thus, be understood as the certainty-equivalent return ('CER') that an investor is willing to accept instead of adopting the risky portfolio rule z. The out-of-sample CER is given by:

$$\widetilde{CER}_z = \tilde{\mu}_z - \frac{\gamma}{2}\tilde{\sigma}_z^2. \tag{20}$$

This facilitates the comparison between strategies as the higher the CER, the higher investor utility.

5.3 (Mean) absolute deviation

To assess the estimation error of any strategy z, we measure the distance from in-sample strategy optimal mean return and volatility to the out-of-sample observed parameters.

Definition 2: *Let $\dot{\mu}_z$ and $\dot{\sigma}_z$ denote the strategy specific mean return and standard deviation subject to the vector of "true", in-sample optimal weights $\mathring{w}_z$, i.e. $\dot{\mu}_z = \mathring{w}_z'\mathring{\mu}$ and $\dot{\sigma}_z = \sqrt{\mathring{w}_z'\mathring{\Sigma}\mathring{w}_z}$. Then the distance from $\tilde{\mu}_z$ to $\dot{\mu}_z$ and $\tilde{\sigma}_z$ to $\dot{\sigma}_z$ is respectively given by:*

$$\widetilde{AD}_z^{\mathring{\mu}} = \mid \dot{\mu}_z - \tilde{\mu}_z \mid \tag{21}$$

$$\widetilde{AD}_z^{\mathring{\sigma}} = \mid \dot{\sigma}_z - \tilde{\sigma}_z \mid . \tag{22}$$

The average of absolute deviations across $c_{n|N}$ observations then yields the mean absolute deviation ($\widetilde{MAD}_z^{\mathring{\mu}}$, $\widetilde{MAD}_z^{\mathring{\sigma}}$).

5.4 Return-loss due to estimation error

In addition to the absolute deviation, we also define a measure of "return-loss" caused by not knowing the true parameters. We propose the following relationship.

Definition 3: *When $\dot{\mu}_z$ and $\dot{\sigma}_z$ are the mean return and standard deviation of strategy z subject to the vector of "true", in-sample optimal weights $\mathring{w}_z$. Then the return-loss caused by not knowing $\mathring{\mu}$ and $\mathring{\Sigma}$ can be defined as:*

$$\widetilde{RL}_z = \frac{\dot{\mu}_z}{\dot{\sigma}_z}\tilde{\sigma}_z - \tilde{\mu}_z. \quad (23)$$

It is, therefore, an expression of the additional return needed for an investment rule z to perform, out-of-sample, as well as the same strategy under true parameter knowledge. The measure relies on linear transformation of the Sharpe ratio of the in-sample strategy-optimal portfolio by the attained out-of volatility under ex-post estimated weights. The average return-loss is then an expression of the average loss against the Sharpe ratio of the in-sample, strategy optimal portfolio. Note the underlying assumption that Sharpe ratio is the objective that investment strategies are evaluated by. The measure is not necessarily proportional to the absoulte deviation (21, 22) since strategies that do not explicitly optimise weightings towards having the maximum Sharpe ratio may in fact attain ex-ante Sharpe ratios that are higher than those of the true strategy optimal portfolios. It is often not recognized that estimation error can, thus, positively affect the performance in terms of some objective other than the explicit strategy target. In other words, while no minimum-variance portfolio can have a lower variance than its ex-ante optimal counterpart, it can have a higher Sharpe ratio (and/or certainty equivalent return).[9] This also explains why different strategies may achieve the best out-of-sample performance when evaluated by the two different performance measures and why (10) may intuitively not hold for minimum-variance portfolios. Nevertheless, we deem (23) the approriate measure since it is the expected loss due to estimation error in the Sharpe ratio objective that fund managers (for whom the Sharpe ratio is a natural benchmark) are ultimately concerned about.

5.5 Return-loss due to unfavourable estimation error

In the following we introduce a measure of "unfavourable" loss due to estimation error that also accounts for non-normality of the out-of-sample return series in performance evaluation.[10] Let the (M-T)x1 vector of out-of-sample daily returns of each strategy be denoted by $\tilde{R}_z$. Then this return series can be decomposed into a threshold return τ minus a downside measure, denoted by $max[\tilde{R}_z - \tau, 0]$, plus an upside measure, denoted by $max[\tau - \tilde{R}_z, 0]$. A measure of downside performance, labeled the lower partial moment of the return distribution, can then be expressed through:

[9]Similarly, no mean-variance optimal portfolio will have a higher certainty equivalent return than the ex-ante optimal counterpart but its Sharpe ratio could still be higher.

[10]A commonly addressed criticism is the non-normality of the returns data. A quick Bera-Jarque test (of which the results are not reported here) reveals that none of the return series are normally distributed. Most financial data suffers from this shortcoming but Tu and Zhou (2004) confirm that, although the assumption is heavily criticised, it can be assumed to hold for a mean-variance investor.

$$lpm_z(\tau, h) = Emax[(\tilde{R}_z - \tau)^h, 0] = \frac{1}{(M-T)} \sum_{t=1}^{(M-T)} (\tilde{R}_{z,t} - \tau)^h d_t, \tag{24}$$

$$d_t = \begin{cases} 0 & \tilde{R}_{z,t} > \tau \\ 1 & \tilde{R}_{z,t} \leq \tau \end{cases}$$

where the h-th degree of the lower partial moment defines how the "penalty function" is shaped. By assuming $h = 2$ and setting $\tau = 0$ we obtain $lpm(0, 2)$, the semi-variance of the return distribution. This measure of dispersion of negative returns corresponds to the measure of downside variance introduced by Sortino and Price (1994) in their construction of the Sortino ratio.

Definition 4: *Let* $\dot{\mu}_z$ *and* $l\dot{pm}_z(0, 2)$ *denote the strategy specific mean return and downside variance subject to the vector of "true", in-sample optimal weights* $\mathring{w}_z$. *Assume that the mean of the out-of-sample returns and the corresponding downside variance are given by* $\tilde{\mu}_z$ *and* $l\tilde{pm}_z(0, 2)$ *respectively. Then the return-loss due to estimation error in the downside risk measure may be stated as follows:*

$$\widetilde{URL}_z = \frac{\dot{\mu}_z}{l\dot{pm}_z(0, 2)} \, l\tilde{pm}_z(0, 2) - \tilde{\mu}_z. \tag{25}$$

By using this measure we account for the presence of larger upper partial moments and thus higher variance in the positive part of the return distribution. Such variance is penalised in consideration of the Sharpe ratio (19) and the corresponding measure of return-loss in (23). Thus, the value of our measure increases the more the out-of-sample distribution of returns deviates from the assumption of i.i.d normality.

6 Empirical analysis

Our primary dataset consists of 4,351 trading days[11] of daily excess return data on the 17 Fama-French industry portfolios. The data is obtained from Kenneth French's website[12] and expressed as the excess return over the prevailing risk-free rate. Returns on the 17 portfolios are taken as asset returns (i.e. each of the portfolios is to be understood as an asset), which seems acceptable in the light of previous research. One must always bear in mind, though, that the equally-weighted Fama French portfolios will on average exhibit lower volatility than a single stock as a result of inherent diversification. Nonetheless, they provide a great representation of the overall market. Working with these portfolios is important, as we do not want to limit ourselves to one index that will be biased by the selection of stocks based on some characteristic. On the other hand, we believe that choosing a number of random stocks like Jagannathan and Ma (2003) or Jorion (1986) does not sufficiently serve the explanatory power of our model, since selection of random stocks may also lead to a stock universe that is tilted towards some industry or sector. Using daily returns in our analysis benefits optimisation strategies while it does not affect the 1/n rule. Jagannathan and Ma (2003) show that when using daily returns data, the true parameters and, thus, the optimal weights in mean-variance portfolios can be estimated more precisely. To ensure robustness of our findings, we also ran the analysis for a dataset of the same length that consists of the 12 Fama-French industry portfolios.[13]

The results section of our work is divided into two parts. In the first part, we examine the performance of the aforementioned investment strategies and their sensitivity to portfolio size by means of the performance indicators, introduced in section 5. In the second part we present our findings on how the portfolio size and degree of risk aversion influence the magnitude of estimation error. Which strategy performs best, out-of-sample, will depend on the stability of return distributions over time. Unstable return distributions induce erroneous ex-post parameter estimates and, thus, parameter es-

[11] Between September 1997 and December 2015.

[12] *http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/*

[13] The results of this analysis are in line with our findings for the 17 Fama-French industry portfolios and, therefore, not reported here. Data on the 12 Fama-French portfolios is available upon request.

timation error. For large degrees of parameter estimation error, we would expect the equally weighted rule to perform relatively better. If distributions are stable and parameter estimates close to the "true" moments of the return distribution, mean-variance optimisation strategies should be able to achieve their respective targets and outperform the equal weighting rule from an out-of-sample perspective.

6.1 Out-of-sample performance of mean-variance investment strategies in dependency of n and γ

The various strategies tested are given in the rows of Table 6.1, while the Sharpe ratio for a range of portfolio sizes can be found in the respective columns. Panel A considers an investor with γ=3, whereas in Panel B we report the results for γ=1. In the first line of each panel we present the Sharpe ratio of the in-sample optimal short-sale unconstrained mean-variance optimal portfolio. When the true parameters are known (and the return series is stationary over time), it would be irrational for an investor with risk aversion γ to hold any other portfolio.

It is evident that the Sharpe ratio of the (in-sample) optimal mvo* strategy increases with the portfolio size. Out-of-sample, however, the mvo strategy exhibits a negative relationship with portfolio size when γ=3, the Sharpe ratio decreases and the mvo strategy constantly underperforms the naive diversification rule. While this finding is in line with DeMiguel et al. (2009b), the opposite holds when γ=1. Here we observe an increase in Sharpe ratio up to n=15, when the set of portfolio assets is extended, but total underperformance against the 1/n is even higher. It is generally noteworthy that lower risk aversion leads to worse out-of-sample performance in terms of Sharpe ratio. The magnitude of this finding may be sample specific, but note that the implied risk-aversion of the maximum Sharpe ratio portfolios (that make no explicit assumption of γ) will be closer to 3 than 1. Hence, the Sharpe ratio as the performance measure for Markowitz mean-variance optimal strategies as specified in (1) is somewhat biased and alternative measures of performance should also be taken into consideration. Looking at certainty equivalent return ('CER') as reported in Table 6.2 we find that for all mean-variance optimal strategies, the CER decreases out-of-sample with increasing n (while at the same time, it becomes larger for the naive diversification). One observes that mvo strategies not only underperform the 1/n rule but yield significant negative CER across both levels of risk aversion. A

Table 6.1: Average out-of-sample Sharpe ratios (p.a.) of selected mean-variance portfolio strategies

This table presents the average, annualised out-of-sample Sharpe ratios of investment strategies. The values are obtained through a rolling-window analysis over M=4,351 days, with the parameters estimated over T=900 days and the portfolios rebalanced every 30 days (expressed as averages of $c_{n|17}$ *observations). The average significance of the out-of-sample returns at the 1%, 5% and 10% levels, is estimated by a student t-test with* $H_0 : \tilde{\mu} = 0$ *and* $H_0 : \tilde{\mu} \neq 0$*, is indicated by ***,** and * respectively.*

	n					
Strategy (N=17)	3	6	9	12	15	17
Panel A *(*$\gamma = 3$*)*						
mvo*	1.2907***	1.4940***	1.6270***	1.7387***	1.8422***	1.9098***
ew (1/n)	0.8512***	0.8602***	0.8636***	0.8653***	0.8664***	0.8669***
mvo	0.7761***	0.7630***	0.7554***	0.7436**	0.7299**	0.7209**
mvo-c	0.9615***	0.9998***	1.0161***	1.0235***	1.0273***	1.0282***
mv	1.1083***	1.3388***	1.4911***	1.5995***	1.6813***	1.7270***
mv-c	1.0197***	1.0932***	1.1148***	1.1173***	1.1140***	1.1128***
Panel B *(*$\gamma = 1$*)*						
mvo*	1.0051***	1.1181***	1.2638***	1.4018***	1.5306***	1.6135***
mvo	0.5015*	0.5643*	0.5919*	0.6023*	0.6041*	0.6032*
mvo-c	0.9118***	0.9265***	0.9426***	0.9669***	0.9941***	1.0158***

ew: equal weighting, **mvo***: mean-variance optimal under true parameter knowledge, **mvo**: sample-based mean-variance optimal, **mvo-c**: sample-based short-sale constrained mean-variance optimal, **mv**: sample-based minimum-variance, **mv-c**: sample-based short-sale constrained minimum-variance

Table 6.2: Average certainty equivalent return (p.a.) of selected mean-variance portfolio strategies

This table presents the average, annualised certainty equivalent return (CER) of investment strategies. The values are obtained through a rolling-window analysis over M=4,351 days, with the parameters estimated over T=900 days and the portfolios rebalanced every 30 days. CER is given in percentages and expressed as an average of $c_{n|17}$ *observations.*

	n					
Strategy (N=17)	3	6	9	12	15	17
Panel A ($\gamma = 3$)						
ew (1/n)	11.4852	12.0106	12.1858	12.2734	12.3259	12.3506
mvo	1.1681	-27.6319	-61.7835	-99.1614	-139.1289	-167.0141
mvo-c	14.4877	15.9118	16.5940	16.9068	17.0539	17.0879
mv	16.3704	19.0627	20.1558	20.6573	20.9041	21.0124
mv-c	14.9184	15.9366	16.0081	15.8200	15.6125	15.5433
Panel B ($\gamma = 1$)						
ew (1/n)	18.5515	18.7267	18.7851	18.8143	18.8318	18.8400
mvo	-23.5751	-116.2702	-221.7786	-335.8810	-457.3370	-541.9481
mvo-c	21.7106	23.2412	24.2845	25.2390	26.1242	26.79260
mv	21.2207	22.4390	22.8446	22.9499	22.9413	22.9445
mv-c	20.1501	20.2066	19.8793	19.4736	19.1418	19.0258

ew: equal weighting, **mvo**: sample-based mean-variance optimal, **mvo-c**: sample-based short-sale constrained mean-variance optimal, **mv**: sample-based minimum-variance, **mv-c**: sample-based short-sale constrained minimum-variance

rational investor would, thus, never pursue an investment in these strategies. Further, note that in contrast to results obtained in consideration of the Sharpe ratio, the CER for the mvo strategy when γ=1 also decreases in larger portfolios.

When imposing short-sale constraints on the mean-variance optimal portfolios the relationship is reversed. As it can be seen from Tables 6.1 and 6.2 the mvo-c portfolio strategy consistently outperforms the ew rule (e.g. when γ=1 and n=15, mvo-c on average attains a Sharpe ratio of 1.0273 vs. the ew rule with only 0.8664). The performance as measured by Sharpe ratio and CER is now strictly positively related to the portfolio size and so is the outperformance over the naive diversification. This supports earlier findings by Duchin and Levy (2009) and stands in stark contrast to DeMiguel et al. (2009b). Furthermore, it becomes evident that short-sale constraints are absolutely necessary when employing mean-variance optimal investment strategies. While mvo-c peform better than the 1/n strategy across both levels of risk aversion, they fall short of the peformance of minimum-variance portfolios that outpace the 1/n diversification and all other investment rules in terms of Sharpe ratio. These results confirm Jagannathan and Ma (2003) as ignoring the estimates of expected returns in portfolio construction leads to better risk-adjusted performance. Bearing this in mind, from Table 6.2 we observe that when γ=1 the short-sale constrained mean-variance optimal strategy consistently yields the highest CER. Hence, as opposed to Jagannathan and Ma (2003), we also show that a mean-variance optimal portfolio strategy can on average outperform the minium-variance portfolio rule and so there is value in exploiting estimates of expected returns.[14]

It becomes evident that which mean-variance strategy yields the best out-of-sample performance is ultimately dependent on the measure used to assess it. Finally, note that when the assets used to construct minimum-variance portfolios are themselves large portfolios, adding non-negativity constraints does not help much since Sharpe ratios and CER for mv-c compared to the mv strategies are persistently lower. With increasing sample size, the probability that a shortsale constraint becomes binding increases and the performance gap becomes larger. Thus, for minimum-variance portfolios imposing constraints not only hurts when the true parameters are known, but also when they are estimated from past returns.

6.2 Results from studying estimation error

Table 6.3 reports the average absolute deviation in terms of realised return and volatility of the various investment strategies compared to the respective true strategy optimal portfolio.[15] We find that when there is estimation error in any of the asset return series, increasing the portfolio size generally results in a higher probability that the optimal weights are based on at least one set of highly erroneous parameter estimates and, thus, higher estimation error. This finding is in line with

[14] This observation does not hold for $\gamma = 3$. To explain this, note that problem (13) corresponds to (15) when the vector $\hat{\mu}$ is a vector of zeros and the investor exhibits an infinite degree of risk aversion. The optimal minium-variance portfolios will always contain the same stocks in their same proportions regardless of the actual degree of investor risk aversion. Hence, when evaluating them in terms of certainty equivalent return they should be expected to perform relatively better for higher levels of risk aversion. It follows that investors with very high degrees of risk aversion will always be relatively better off, investing in the minimum-variance portfolio.

[15] Mean absolute deviations of the ew strategy are not reported since it is free from estimation error.

the bulk of research and validates our methodology. Furthermore, measurement error in the mvo strategy that directly relies on estimates of expected return is significantly larger than in the mv strategy that does not rely on the latter. It follows that the error in mean return estimates must be relatively larger than the error in covariance estimates. In fact, the mv is the only strategy for which the mean absolute deviation in both mean returns and volatilities is reduced in larger sample sizes. To our surprise, imposing short-sale constraints on the minimum-variance portfolio leads to higher mean absolute deviations. While Jagannathan and Ma (2003) have shown that imposing long-only constraints on the minimum-variance portfolio does generally not help much in terms of out-of-sample performance, most literature suggests that imposing short-sale constraints on the mv portfolio reduces measurement error. It appears that the shrinkage of covariance matrix unfavourably affects the out-of-sample performance of minimum-variance investment strategies. Imposing short-sale constraints on the mvo portfolio, however, sucessfully reduces estimation error.

The return-loss of the rebalanced ex-post optimal portfolios against their in-sample counterpart, reported in Table 6.4, illustrates the difference in Sharpe ratios. It can be seen that with increasing portfolio size, the loss in return from applying any strategy that relies on parameter estimates also increases. While the annual return-loss against the in-sample mean-variance optimal portfolio is 25.46% for the mvo when n=3, it grows to 157.22% when n=17. The same relationship holds for less risk-averse investors and across all investment rules. Furthermore, from Table 6.4 we conclude that the impact of estimation error is up to 55 times larger in short-sale unconstrained strategies that rely on mean return estimates than in those that are only based on the covariance estimates (e.g. compare 109.14% (mvo) vs. 2.69% (mv) when n=12 in Panel A). This result is striking and appears as a strong case against mean-variance optimal investment strategies in the light of parameter estimation error. The magnitude of this difference can, however, easily be controlled for by imposing short-sale constraints, where loss due to estimation error in mvo-c is "only" up to 3 times larger than in mv-c. Short-sale constraints, thus, become absolutely necessary in controlling estimation error and improving out-of-sample performance, when portfolio strategies directly rely on mean return estimates. It should be noted that in line with mean absoulte deviation statistics, imposing short-sale constraints on the minimum-variance portfolio also leads to higher return-loss due to estimation error. Finally, our results also show that a relatively worse out-of-sample performance of the mv-c vs. the mv strategy may directly be related to higher estimation error in the covariance estimates.

Table 6.3: Mean absolute deviations (p.a.) of mean-variance investment strategies

This table presents the average annualised out-of-sample absolute deviations (as derived in 5.3) in terms of realised portfolio return and risk of investment strategies from their in-sample optimal counterparts. The values are obtained through a rolling-window analysis over M=4,351 days, with the parameters estimated over T=900 days and the portfolios rebalanced every 30 days. Figures are expressed as averages of $c_{n|17}$ *observations.*

	n											
Strategy (N=17)	3		6		9		12		15		17	
	$\widetilde{MAD}_z^{\hat{\mu}}$	$\widetilde{MAD}_z^{\hat{\sigma}}$	$\widetilde{MAD}_z^{\hat{\mu}}$	$\widetilde{MAD}_z^{\hat{\sigma}}$	$\widetilde{MAD}_z^{\hat{\mu}}$	$\widetilde{MAD}_z^{\hat{\sigma}}$	$\widetilde{MAD}_z^{\hat{\mu}}$	$\widetilde{MAD}_z^{\hat{\sigma}}$	$\widetilde{MAD}_z^{\hat{\mu}}$	$\widetilde{MAD}_z^{\hat{\sigma}}$	$\widetilde{MAD}_z^{\hat{\mu}}$	$\widetilde{MAD}_z^{\hat{\sigma}}$
Panel A *($\gamma = 3$)*												
mvo	0.1703	0.2385	0.1803	0.4322	0.1639	0.5660	0.1397	0.6740	0.0969	0.7674	0.0169	0.8241
mvo-c	0.0449	0.0405	0.0467	0.0476	0.0551	0.0501	0.0659	0.0497	0.0784	0.0462	0.0876	0.0422
mv	0.0310	0.0114	0.0277	0.0070	0.0245	0.0045	0.0223	0.0034	0.0221	0.0030	0.0229	0.0029
mv-c	0.0306	0.0116	0.0273	0.0108	0.0281	0.0112	0.0318	0.0118	0.0369	0.0122	0.0398	0.0125
Panel B *($\gamma = 1$)*												
mvo	0.5035	0.8692	0.5568	1.3877	0.5154	1.7567	0.4427	2.0653	0.3101	2.3381	0.0510	2.5050
mvo-c	0.0501	0.0522	0.0490	0.0642	0.0517	0.0675	0.0560	0.0690	0.0609	0.0697	0.0633	0.0698

MAD: mean absolute deviation, **mvo**: sample-based mean-variance optimal, **mvo-c**: sample-based short-sale constrained mean-variance optimal, **mv**: sample-based minimum-variance, **mv-c**: sample-based short-sale constrained minimum-variance

Table 6.4: Average out-of-sample return-loss (p.a.) of mean-variance investment strategies

This table presents the average annualised out-of-sample return-loss (as derived in 5.4) of investment strategies due to estimation error. The values are obtained through a rolling-window analysis over M=4,351 days, with the parameters estimated over T=900 days and the portfolios rebalanced every 30 days. Figures are given in percentages and expressed as averages of $c_{n|17}$ *observations.*

	n					
Strategy (N=17)	3	6	9	12	15	17
Panel A *(*$\gamma = 3$*)*						
mvo	25.4580	54.6938	81.8342	109.1356	137.5019	157.2228
mvo-c	8.5538	11.2855	12.9994	14.2082	15.0253	15.3798
mv	3.9725	3.0494	2.7343	2.6891	2.7672	2.8411
mv-c	4.6235	4.2692	4.4665	4.9231	5.4536	5.7575
Panel B *(*$\gamma = 1$*)*						
mvo	65.6470	119.4459	185.7562	260.3611	341.4889	399.0968
mvo-c	9.9285	13.2602	14.8750	15.8462	16.7820	17.2841

mvo: sample-based mean-variance optimal, **mvo-c**: sample-based short-sale constrained mean-variance optimal, **mv**: sample-based minimum-variance, **mv-c**: sample-based short-sale constrained minimum-variance

In Table 6.5 we illustrate the return-loss due to estimation error in the semi-variance. Comparing the results in Table 6.5 with those reported in Table 6.4, it becomes evident that the unfavourable estimation error in short-sale unconstrained mean-variance optimal strategies must be relatively higher than the estimation error in the upper part of the sample return distribution. We also show that an investor is exposed to signifcantly more downside risk by pursuing a short-sale unconstrained mean-variance optimal strategy. Again, by imposing short-sale constraints one is able to achieve an impressive reduction in return-loss due to estimation error in the lower partial moments (particularly in larger sample sizes). For the mvo-c strategy, unfavourable estimation error is lower than the overall estimation error observed earlier. The same can be shown for minimum-variance strategies, in which we find relatively lower unfavourable estimation error and so conlcude that the out-of-sample return distribution must include more extreme positive than negative returns. Given an expected unfavourable return-loss of a mere 0.76% p.a. for investors in the long-short minimum-variance strategy, one must not be too concerned with estimation error when pursuing the mv rule.

Table 6.5: Average "unfavourable" out-of-sample return-loss (p.a.) of mean-variance investment strategies

This table presents the average annualised "unfavourable" out-of-sample return-loss (as derived in 5.5) of investment strategies due to estimation error. The values are obtained through a rolling-window analysis over M=4,351 days, with the parameters estimated over T=900 days and the portfolios rebalanced every 30 days. Figures are given in percentages and expressed as averages of $c_{n|17}$ *observations.*

	n					
Strategy (N=17)	3	6	9	12	15	17
Panel A *(*$\gamma = 3$*)*						
mvo	33.6733	72.7890	105.9240	137.3261	168.5599	189.7230
mvo-c	5.0775	7.5584	8.8291	9.2913	9.1120	8.6846
mv	0.6557	0.5467	0.6557	0.7480	0.7767	0.7601
mv-c	1.2249	1.4757	1.6797	1.7984	1.8657	1.9097
Panel B *(*$\gamma = 1$*)*						
mvo	93.6646	169.2103	248.6587	332.0202	418.7219	478.4411
mvo-c	6.2570	9.5836	11.0681	11.8393	12.4742	12.8510

mvo: sample-based mean-variance optimal, **mvo-c**: sample-based short-sale constrained mean-variance optimal, **mv**: sample-based minimum-variance, **mv-c**: sample-based short-sale constrained minimum-variance

When looking at Panels A and B of the return-loss statistics, one finds the loss to be relatively higher in mvo-c strategies that assume a lower risk aversion (γ). This is in line with Kan and Zhou's proposed closed-form solution for the short-sale unconstrained mean-variance optimal portfolio problem and confirms Hypothesis 2.1. Levy and Duchin (2003), in this context, point out that optimal portfolios that maximise expected utility with higher degrees of risk aversion tend to be more diversified and include less extreme weights relative to optimal portfolios that maximise expected utility with lower degrees of risk aversion.

Note that (13) can be generalized when γ is unknown by:[16]

$$\hat{w}_{mvo} = \frac{\hat{\Sigma}^{-1}\mathbf{1}}{\mathbf{1}'\hat{\Sigma}^{-1}\mathbf{1}} + \gamma\left(\hat{\Sigma}^{-1} - \frac{\hat{\Sigma}^{-1}\mathbf{1}\mathbf{1}'\hat{\Sigma}^{-1}}{\mathbf{1}'\hat{\Sigma}^{-1}\mathbf{1}}\right)\hat{\mu}. \tag{26}$$

It can be seen that by assuming values of γ from 0 to $+\infty$, this problem describes the solution to the mean-variance efficient frontier. From the formula it can be inferred that more risk-averse investors will choose mean-variance optimal combinations with lower levels of total risk (since $\frac{1}{\mathbf{1}'\hat{\Sigma}^{-1}\mathbf{1}}$ is the variance of the minimum-variance portfolio). Hypothesis 2.1 may thus be translated into the sensitivity of estimation error to the volatility of the ex-post optimal portfolio. From our results in part 6 and related literature we deduce that the out-of-sample volatility is on average larger than the estimated volatility due to parameter estimation error. Bearing in mind that in long-only portfolios the volatility cannot be reduced through short-selling of another asset but only through diversification, constraining the volatiliy in mean-variance efficient portfolios could be understood as an additional constraint to extreme weights in the portfolio constituents. More diversification and less extreme weights should lead to less overall estimation error and likely to better relative performance.

To assess the impact of constraining the total risk level of a mean-variance optimal portfolio on estimation error, we turn back to the optimisation problem underlying the efficient frontier (1). We employ a methodology that allows us to estimate the impact of estimation error in terms of ex-ante loss in return for a given risk level. For portfolios of size n, we estimate the ex-post optimal short-sale constrained mean-variance optimal portfolios for a set of k=5 risk levels σ_k ($\sigma_1 < \sigma_2 < \sigma_3 < \sigma_4 < \sigma_5$) over an estimation period of length T. For each n, this yields k=5 efficient portfolios ($\hat{w}_k$) in which the stocks are held proportionally to the respective mean-variance optimal weights. We then simulate the

[16] For a comprehensive discussion and derivation, we refer the reader to Bodnar et al. (2013).

performance of these portfolios over the following (ex-ante) observation period of length O and record their out-of-sample return and standard deviation:

$$arg\ \underset{w}{max}\ w'\hat{\mu} \quad s.t.\ \sqrt{w'\hat{\Sigma}w} = \sigma_k\ and\ \mathbf{1_n'}w = 1\ and\ w \geq 0 \tag{27}$$

$$\tilde{\mu}_k = \hat{w}_k \mathring{\mu} \tag{28}$$

$$\tilde{\sigma}_k = \sqrt{\hat{w}_k' \mathring{\Sigma} \hat{w}_k}. \tag{29}$$

In the next step, we find the ex-ante mean-variance optimal portfolios (w_k^*) with the additional constraint that they must have the same volatility as the observed out-of-sample variance of the ex-post optimal portfolios. The difference in return of the two portfolios at the same risk level is then an expression of the inherent return-loss due to estimation error. For each risk level σ_k, this may be illustrated as follows:

$$arg\ \underset{w}{max}\ w'\mathring{\mu} \quad s.t.\ \sqrt{w'\mathring{\Sigma}w} = \tilde{\sigma}_k\ and\ \mathbf{1_n'}w = 1\ and\ w \geq 0 \tag{30}$$

$$\mu_k^* = w_k^* \mathring{\mu} \tag{31}$$

$$RL_k(w_k^*, \hat{w}_k) = \mu_k^* - \tilde{\mu}_k. \tag{32}$$

We run this analysis with simulated data from a one-factor model since it allows us to bind the variance within a specific range (making the created portfolios comparable). Following Mackinlay and Pastor (2000), DeMiguel et al. (2009b) and Tu and Zhou (2011), we assume that the market is composed of a risk free asset and n risky assets that include F factors. A factor model of the form $R_y = \alpha + \beta R_x + \varepsilon$ can then be used to generate the excess returns of the remaining n-F risky assets. R_y here is the vector of generated excess asset returns, R_x is the vector of excess returns on the factor portfolios, β the vector of factor loadings, α the vector of mispricing coefficients and ε the vector of noise.

When simulating the data we follow Tu and Zhou (2011) and assume the excess return of the factor to be normally distributed with an annual mean and volatility of 8% and 16% respectively ($R_x \sim N(0.08, 0.16)$). The mispricing is set to 0 as in DeMiguel et al. (2009b) and the factor loadings (β) are evenly distributed between 0.5 and 1.5. Volatility is a function of the residual variance-covariance matrix ($\varepsilon \sim N(0, \Sigma_\varepsilon)$), where the variance-covariance matrix is assumed to be diagonal with its elements drawn from an uniform distribution with a support of [0.10,0.30]. The cross-sectional average annual idiosyncratic volatility is set to 20% and the generated return series will be i.i.d. normal. We generate 10,000 datasets and 50,000 portfolios to ensure robustness of our results.

Table 6.6 presents the results of our analysis. We find evidence that the degree of (loss due to) estimation error increases with the standard deviation of the ex-post mean-variance optimal portfolio. The return-loss for the high-risk portfolios is about 5 times higher than the return-loss of the low-risk asset combinations. When we allow for higher levels of risk for the estimated mean-variance optimal portfolio, it appears to take more extreme positions in stocks with potentially highly erroneous parameter estimates. This ultimately leads to higher losses due to misspecification of the estimates. To illustrate this, compare the average annual return-loss of the low risk portfolio for n=25 (30.41%) to the high risk portfolio (109.96%). We conclude that constraining the maximum level of acceptable volatility in mean-variance optimal portfolios could help to control estimation error and confirm Hypothesis 2.2. Table 6.5 also reveals how sensitive this finding is to the portfolio size n. Again, we confirm that the overall level of estimation error increases with the portfolio size. The increase is not proportional across risk levels but marginally higher for portfolios found further on the right of the ex-post efficient frontier.

These findings illustrate why minimum-variance portfolios are generally more robust against estimation error and frequently outperform mean-variance optimal portfolios (or those designed to maximise Sharpe ratio) in the out-of-sample period. But first and foremost, this implies that one can control for estimation error by constraining the total risk of the portfolio. When doing so, investors should be able to achieve better out-of-sample performance in terms of Sharpe ratios. Or in other words, choosing the riskier portfolio will then lead to relatively worse ex-ante performance. This may be confirmed by results reported in Table 6.7 that provide strong evidence of the out-of-sample Sharpe ratios for high-risk portfolios being remarkably lower than Sharpe ratios for low-risk portfolios. This finding holds across all n.

Table 6.6: Average out-of-sample return-loss (p.a.) of the volatility constrained ex-post mean-variance optimal investment strategy vs. the ex-ante optimal strategy

This table presents the average annualised return-loss of the short-sale constrained ex-post mean-variance optimal investment strategy (as derived in 5.4) under varying ex-post volatility levels. The values are obtained through Monte-Carlo methods, with the parameters estimated over T=900 days and the portfolio constituents held in their optimal proportion over the following O=90 days. The underlying return series is generated from a one-factor model, following MacKinlay and Pastor (2000). Averages are based on 10,000 datasets and 50,000 portfolios. Figures are given in percentages.

	n			
Risk level (σ_k)	10	15	20	25
Low (σ_1)	17.3725	22.1677	27.0596	30.4131
σ_2	38.1947	55.4474	68.5012	75.7535
σ_3	66.1569	78.7477	89.7152	96.2691
σ_4	78.4324	90.2979	99.9294	105.7627
High (σ_5)	84.1478	95.3503	104.3091	109.9552

Table 6.7: Average out-of-sample Sharpe-ratios (p.a.) of the volatility constrained ex-post mean-variance optimal investment strategy

This table presents the average annualised out-of-sample Sharpe ratios of the shortsale constrained ex-post mean-variance optimal investment strategy under varying ex-post volatility levels. The values are obtained through Monte-Carlo methods, with the parameters estimated over T=900 days and the portfolio constituents held in their optimal proportion over the following O=90 days out-of-sample period. The underlying return series is generated from a one-factor model, following MacKinlay and Pastor (2000). Averages are based on 10,000 datasets and 50,000 portfolios.

	n			
Risk level (σ_k)	10	15	20	25
Low (σ_1)	0.9126	1.0567	1.3792	1.5282
σ_2	0.7771	0.7597	0.6607	0.6691
σ_3	0.5155	0.4918	0.4021	0.3895
σ_4	0.3823	0.3426	0.2930	0.3045
High (σ_5)	0.2978	0.2747	0.2419	0.2587

However, this shall not imply that choosing the riskier portfolio naturally leads to worse ex-ante performance in terms of utility. As pointed out previously, decreasing γ can positively affect the performance of short-sale constrained mean-variance optimal strategies in terms of certainty equivalent return. Hence, a relatively larger inherent degree of estimation error does not necessarily constitute a relatively worse ex-ante performance. Which strategy performs best, out-of-sample, largely depends on the chosen method of performance evaluation. Therefore, when investment strategies are long only and maximising utility is the objective, choosing the riskier portfolio can yet be beneficial.

7 Conclusion

Research on parameter estimation error is concerned with the question of "how much worse off an investor is, if the distribution of returns is estimated with an error". Our analysis provides two very different answers to this question. On the one hand, we find that with increasing size of the opportunity set (portfolio sizes) the measurement error in strategies that rely on parameter estimates increases, confirming Hypothesis 1.1. On the other hand, we show that there exists no causal relationship between the total amount of estimation error and the out-of-sample performance of a mean-variance investment strategy. When estimation error is within reasonable bounds, increasing the portfolio size leads to improved ex-ante performance. Furthermore, we observe that short-sale constrained mean-variance strategies, which directly rely on an estimate of expected returns, do on average outperform the naive 1/n diversification rule in terms of Sharpe ratio, just as minimum-variance portfolios do. We, thus, reject Hypothesis 1.2 that was derived based on results of DeMiguel et al. (2009b).

Most notably, our work also highlights the importance of imposing short-sale constraints in mean-variance optimal investment rules in order to control estimation error and to improve the out-of-sample performance. In further analysis we explicitly show that when maximizing Sharpe ratio is the objective, constraining risk levels in ex-post optimal portfolios helps to control estimation error and leads to superior ex-ante performance. This finding also serves as a valid explanation, as to why minimum-variance portfolios are generally more robust against estimation error and frequently outperform mean-variance optimal portfolios.

References

Benartzi, S. and Thaler, R. (2001). Naive Diversification Strategies in Defined Contribution Saving Plans. *The American Economic Review*, 91(1):79–98.

Bodnar, T., Parolya, N., and Schmid, W. (2013). On the Equivalence of Quadratic Optimization Problems Commonly Used in Portfolio Theory. *European Journal of Operational Research*, 229(3):637–644.

Broadie, M. (1993). Computing Efficient Frontiers using Estimated Parameters. *Annals of Operations Research*, 45:21–58.

Brown, S. (1976). *Optimal portfolio choice under uncertainty: A Bayesian approach*. PhD thesis, University of Chicago.

Chopra, V. K. and Ziemba, W. T. (1993). The Effect of Errors in Means , Variances, and Covariances on Optimal Portfolio Choice. *The Journal of Portfolio Management*, 19(2):6–11.

DeMiguel, V., Garlappi, L., Nogales, F. J., and Uppal, R. (2009a). A Generalized Approach to Portfolio Optimization: Improving Performance by Constraining Portfolio Norms. *Management Science*, 55(5):798–812.

DeMiguel, V., Garlappi, L., and Uppal, R. (2009b). Optimal Versus Naive Diversification : How Inefficient is the 1 / N Portfolio Strategy ? *The Review of Financial Studies*, 22(5):1916–1953.

Duchin, R. and Levy, H. (2009). Markowitz Versus the Talmudic Diversification Strategies. *The Journal of Portfolio Management*, 35(2):71–75.

Engle, R. F. (1993). Statistical Models for Financial Volatility. *Financial Analysts Journal*, 49(1):72–78.

Frost, P. and Savarino, J. (1988). For better performance: Constrain portfolio weights. *The Journal of Portfolio Management*, 15(1):29–34.

Green, R. and Hollifield, B. (1992). When Will Mean-Variance Efficient Portfolios be Well Diversified? *The Journal of Finance*, 47(5):1785–1809.

Haugen, R. A. and Baker, N. L. (1991). The efficient market inefficiency of capitalization-weighted stock portfolios. *The Journal of Portfolio Management*, 17(3):45–50.

Jagannathan, R. and Ma, T. (2003). Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps. *The Journal of Finance*, 58(4):1651–1683.

Jobson, J. D. and Korkie, B. (1980). Estimation for Markowitz Efficient Portfolios. *Journal of the American Statistical Association*, 75(371):544–554.

Jobson, J. D. and Korkie, B. M. (1981). Performance Hypothesis Testing with the Sharpe and Treynor Measures. *The Journal of Finance*, 36(4):889–908.

Jorion, P. (1986). Bayes-Stein Estimation for Portfolio Analysis. *The Journal of Financial and Quantitative Analysis*, 21(3):279–292.

Kallberg, J. G. and Ziemba, W. T. (1984). Mis-specifications in portfolio selection problems. In *Risk and capital*, pages 74–87. Springer Berlin Heidelberg.

Kan, R. and Zhou, G. (2007). Optimal Portfolio Choice with Parameter Uncertainty. *Journal of Financial and Quantitative Analysis*, 42(3):621–656.

Kempf, A. and Memmel, C. (2006). Estimating the Global Minimum Variance Portfolio. Technical Report October.

Kroll, Y., Levy, H., and Markowitz, H. (1984). Mean-Variance Versus Direct Utility Maximization. *The Journal of Finance*, 39(1):47–61.

Lai, T. L., Xing, H., and Chen, Z. (2011). Mean-variance portfolio optimization when means and covariances are unknown. *Annals of Applied Statistics*, 5(2):798–823.

Ledoit, O. and Wolf, M. (2004). Honey , I Shrunk the Sample Covariance Matrix. *The Journal of Portfolio Management*, 30(4):110–119.

Levy, H. and Duchin, R. (2003). Markowitz's Mean-Variance Rule and the Talmudic Investment Recommendation. In *Handbook of Portfolio Construction*, pages 97–123.

Levy, H. and Sarnat, M. (1972). Safety First - An Expected Utility. *The Journal of Financial and Quantitative Analysis*, 7(3):1829–1834.

Lo, A. W. (2002). The Statistics of Sharpe Ratios. *Financial Analysts Journal*, 58(4):36–52.

Mackinlay, C. and Pastor, L. (2000). Asset Pricing Models : Implications for Expected Returns and Portfolio Selection. *The Review of Financial Studies*, 13(4):883–916.

Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, 7(1):77–91.

Michaud, R. (1989). The Markowitz Optimization Enigma : Is ' Optimized ' Optimal? *Financial Analysts Journal*, 45(1):31–42.

Pástor, L. (2000). Portfolio Selection and Asset Pricing Models. *The Journal of Finance*, 55(1):179–223.

Pratt, J. (1964). Risk Aversion in the Small and in the Large. *Econometrica*, 32(1):122–136.

Scherer, B. (2007). *Portfolio Construction and Risk Budgeting*. Risk Books.

Sortino, F. and Price, L. (1994). Performance Measurement in a Downside Risk Framework. *Journal of Investing*, 3(3):59–64.

Tobin, J. (1958). Liquidity Preference as Behavior Towards Risk. *The Review of Economic Studies*, 25(2):65–86.

Tu, J. and Zhou, G. (2004). Data-generating process uncertainty : What difference does it make in portfolio decisions? *The Journal of Financial Economics*, 72(2):385–421.

Tu, J. and Zhou, G. (2011). Markowitz meets Talmud : A combination of sophisticated and naive diversification strategies. *Journal of Financial Economics*, 99:204–215.

Zellner, A. and Chetty, K. (1965). Prediction and Decision Problems in Regression Models from the Bayesian Point of View. *Journal of the American Statistical Association*, 60(310):608–616.