ERASMUS UNIVERSITEIT ROTTERDAM

ERASMUS SCHOOL OF ECONOMICS

The Effect of Estimation Error on Portfolio Weights

Author:
J. Oudenaarden
358270
Supervisor:
B.F. Diris

Abstract

Using sample moments of asset returns as inputs to construct optimal portfolio weights results in volatile and extreme portfolio weights. This research therefore compares a number of methods that reduce the estimation error in the mean or in the covariance matrix. The methods are evaluated over different estimation periods with the help of simulated and empirical stock data, where the $1/N$ portfolio is used as a benchmark, because DeMiguel et al. (2009) found no portfolio that could significantly outperform the $1/N$ portfolio. In contrast to (DeMiguel et al., 2009), who used diversified equity portfolios to construct their portfolios, this paper uses individual stocks as available assets, which have higher idiosyncratic risk. It is found that the global minimum variance portfolios with short-sale constraints perform as good or better than the $1/N$ rule consistently across all datasets and estimation periods.

January 11, 2017

1 Introduction

The optimal mean-variance portfolio weights introduced by Markowitz (1952) are the most popular portfolio weights used in practice. The mean-variance portfolio weights are computed such that the resulting portfolio has the lowest risk for a given level of expected return. However, when sample moments (mean and covariance matrix) are used as inputs, they are notorious for resulting in extreme and non intuitive weights (Best and Grauer, 1991). The main cause for this is the estimation error in the estimators. Therefore mean-variance portfolios perform bad out-of-sample.

Jobson and Korkie (1980) demonstrated, as one of the first, the problems of the plug-in estimates for different sample sizes. They found in a simulation experiment that the mean–variance frontiers obtained by the sample mean and covariance matrix are excessively volatile and on average perform worse than the true mean–variance frontier. Small adjustments in the input (the historical mean and covariance matrix) result in large changes in the portfolio weights. These issues cause problems in practice. Extreme portfolio weights are difficult to implement and large changes result in high transaction costs. It is clear that assets with large positive historical risk premia tend to get a large positive weight assigned to them, and vice-versa. However, this large positive/negative risk premia can be a result of a large positive/negative estimation error. This means that assets of which the risk premium contains a large positive estimation error is given a large positive weight, and vice-versa. Michaud (1989) claims therefore, that the mean-variance optimal weights are in a way "error maximizers". Green and Hollifield (1992) state that the resulting extreme portfolio weights of optimal portfolios based on sample moments are due to the dominance of a single factor in equity returns.

Many methods have been introduced that reduce the estimation error in the mean or the covariance matrix. James and Stein (1961) showed that his James-Stein estimator dominates the sample mean in terms of the mean squared error. Black and Litterman (1992) developed a method that combines the investor's views with the market equilibrium, which results in diversified and intuitive portfolio weights. Ledoit and Wolf (2003b) introduced an estimate of the covariance matrix that 'shrinks' the unstructured sample estimate to the structured implied estimate by the single index model. This way the extreme values in the covariance matrix, which most likely cause the estimation error, are pulled towards the more stable estimates of the single index model. Factor models reduce the statistical error of the plug-in estimates by reducing the number of elements that have to be estimated for the covariance matrix, by imposing a factor structure for the covariance among assets. Pástor (2000) introduced the idea to incorporate a prior on α that determines the strength of the belief in the factor model. Portfolio restrictions are used to reduce the effect the estimation error has on the portfolio weights. Jagannathan and Ma (2003) showed that imposing portfolio restrictions is a form of shrinkage and give an explanation why these constraints reduce the estimation error. Tu and Zhou (2011) derived an optimal combination rule between the $1/N$ rule and the Markowitz portfolio, that reduces the estimation error in the portfolio weights.

This paper investigates what methods perform the best in producing estimates for the mean and the covariance with as little estimation error as possible. These estimates can then be used to compute accurate portfolio weights. This is done by computing tangency and minimum variance portfolios, using the plug-in inputs estimated by the various methods, and evaluating them by their out-of-sample Sharpe-ratio and turnover. A similar research has been done by DeMiguel et al. (2009), where they compared 14 estimation methods to construct portfolio weights. They did this by constructing out of sample optimal portfolios using the moments estimated by various methods. They compared the portfolios to each other and to the naive $1/N$ portfolio. They found that the $1/N$ portfolio is not outperformed consistently, in terms of Sharpe-ratios, across the seven datasets that they used by the 14 methods. They say that the main reason that the $1/N$ portfolio performs so well, is that the $1/N$ does not suffer from the estimation error in the mean and covariance matrix.

Since (DeMiguel et al., 2009), a lot of research has been done with the goal to find a strategy that outperforms the $1/N$ rule. Tu and Zhou (2011) derived an optimal combination rule that combines the $1/N$ rule with a more complicated model, such as the Markowitz portfolio weights. They state that the resulting weights theoretically strictly outperform both counterparts of the combination. However, due to estimation errors in the optimal combination coefficient, this is not always the case in practice. Both DeMiguel et al. (2009) and Tu and Zhou (2011) use diversified equity portfolios to construct their portfolios. These equity portfolios contain less idiosyncratic risk than individual stocks. Because tangency and minimum variance portfolios diversify away the idiosyncratic risk, the $1/N$ rule performs better relative to the optimal portfolios when the available assets have low idiosyncratic risk.

That is why, in contrast to DeMiguel et al. (2009), this paper constructs portfolios for individual stocks instead of diversified factors. This paper also investigates how the methods perform in different states of the economy, such as high or low volatile periods. The goal of this paper is thus to investigate whether the $1/N$ rule still belongs to the top portfolios as it did in the paper by DeMiguel et al. (2009), or that the higher idiosyncratic risk in the individual stocks results in a worse performance of the $1/N$ rule compared to the mean-variance optimized portfolios. Also a higher estimation window, of 180 months, is used compared to DeMiguel et al. (2009), who use estimation windows of 60 and 120 months. A higher estimation window should result in more accurate estimates of the moments and thus more accurate portfolio weights. The portfolios are evaluated in different states of the market, such as periods of recession or expansion, which leads to interesting results as some portfolios perform better than others in different periods.

The portfolios are based on various datasets that consist of stocks from both the New York and NASDAQ stock exchange. Estimation error reduction methods will be used for the moments on which the optimal portfolios are based, some of which are also used by DeMiguel et al. (2009). Simulated stock returns are used to compare the different estimation techniques. The estimated moments and resulting portfolios are evaluated in terms of different loss functions. In the simulation experiments, the 'Model Confidence Set' is used to determine which method performs significantly the best out of all methods. This is a set of methods that contains the 'best' method for a given level of confidence in terms of a defined loss function. The best methods are combined as inputs for the tangency portfolios. All the different portfolios are used to compute the portfolio weights will be explained in the Methodology section. To test whether there is a significant difference between the out-of-sample Sharpe ratios, the two-sample statistic for comparing Sharp-ratios from Opdyke (2007) is used.

It is found that the $1/N$ still performs good relative to the other methods when the available assets are individual stocks. For the various datasets, it is almost never outperformed significantly by the other methods. However, in contrast to DeMiguel et al. (2009) there are also portfolios that are not outperformed by the $1/N$ portfolio consistently across all datasets and estimation periods. For instance, the minimum-variance portfolios with short-sale constraints are never outperformed by the benchmark in any of the estimation periods and datasets. Also the unconstrained minimum-variance portfolios, that use the shrinkage estimator from Ledoit and Wolf (2003a) for the covariance matrix, outperform the $1/N$ portfolio for the estimation periods of 120 and 180 months for one of the datasets that uses stocks from the NASDAQ stock exchange. In general, the strategies that ignore the expected returns as an input perform the best in terms of out-of-sample Sharpe-ratio. This suggests that the estimation error in the mean has a much larger effect on the portfolio weights than that of the covariance matrix. This concurs with Chopra and Ziemba (2011) who found that the effect of the estimation error in the mean is about 10 times larger than that of the covariances and variances combined.

The portfolio turnover of the short-sale constrained minimum-variance portfolios is also lower than that of the $1/N$ portfolio across all estimation period and datasets. For higher estimation periods, also the unconstrained minimum-variance portfolios and the mean-variance portfolios that uses CAPM implied moments outperform the benchmark strategy in terms of turnover. Still the $1/N$ portfolio is one of the top portfolios in terms of turnover as it outperforms most portfolios and is only slightly outperformed when it is. In periods of low market volatility, the $1/N$ is mostly outperformed by the short-sale constrained minimum-variance portfolios. Also in periods of high volatility the tangency portfolios that use CAPM or Fama and French implied moments on average have a higher out-of-sample Sharpe-ratio.

The set-up of the rest of this paper is as follows. First the methods that are used to estimate the moments and evaluate the portfolios are explained in the methodology section. Then the data is described and the summary statistics of the datasets are shown. Then the results of the simulation experiments are presented in the results section. The evaluation of the tangency portfolios is also displayed in the results section. Finally the conclusion laid out in the final section.

2 Methodology

A portfolio manager has a menu of available assets that consists of a number (K) of risky assets and a risk-free asset. Then this investor tries to compute the vector of portfolio weights (w) for these assets. This determines what percentage of his capital he should invest in each asset. The most used and well known method is the Markowitz (1952) model. Here, the investor maximizes the following objective function:

$$\max_{w} w'\tilde{\mu} - \frac{\gamma}{2} w'\Sigma w, \tag{1}$$

where $\tilde{\mu} \equiv \mu - \iota R_f$ is the expected excess return for the assets (risk premia), Σ is the covariance matrix of the risky assets and γ is the coefficient of relative risk aversion. The solution to this optimization problem can be found by setting the first derivative of the function that needs to be maximized equal to zero:

$$\tilde{\mu} - \gamma \Sigma w = 0 \tag{2}$$

$$\Rightarrow w = \frac{1}{\gamma} \Sigma^{-1} \tilde{\mu}, \tag{3}$$

The vector w contains the optimal weights for the K risky assets and $1 - w'\iota$ is the weight for the risk-free asset. The Markowitz portfolio is a combination of the risk-free rate and the tangency portfolio, which is fully invested in risky assets:

$$w_{TGP} = \frac{\Sigma^{-1}\tilde{\mu}}{|\iota'\Sigma^{-1}\tilde{\mu}|} \tag{4}$$

Another portfolio used in this paper is the global-minimum-variance (GMV) portfolio. These portfolio weights are found by minimizing the variance of the portfolio without accounting for the expected returns:

$$\min_{w} \frac{1}{2} w'\Sigma w \text{ s.t.} \tag{5}$$

$$w'\iota = 1 \tag{6}$$

This problem can be solved by using the Lagrangian:

$$L = w'\Sigma w - \lambda(w'\iota - 1) \tag{7}$$

$$\Rightarrow \frac{\delta L}{\delta w} = \Sigma w - \lambda \iota = 0 \tag{8}$$

$$\Rightarrow \frac{\delta L}{\delta \lambda} = w'\iota - 1 = 0 \tag{9}$$

$$\text{Rewrite (8)} \Rightarrow w = \lambda \Sigma^{-1} \iota \tag{10}$$

$$\text{Substitute (10) in (9)} \Rightarrow \lambda = \frac{1}{\iota'\Sigma^{-1}\iota} \tag{11}$$

$$\text{Substitute (11) in (10)} \Rightarrow w_{GMV} = \frac{\Sigma^{-1}\iota}{\iota'\Sigma^{-1}\iota} \tag{12}$$

$$\tag{13}$$

Both the tangency and the GMV portfolio weights are evaluated in this paper. To compute the weights of both the tangency and the GMV portfolio, the μ and Σ have to be estimated.

2.1 Sample Moments

The most common and easy way to estimate the expected returns (μ) is to use the sample mean ($\bar{y}$):

$$\bar{y} = \frac{1}{T} \sum_{t=1}^{T} y_t \tag{14}$$

The standard statistical method of estimating the covariance matrix is to use the sample covariance matrix $S = [\hat{\sigma}^2_{i,j}]$ with the following entries:

$$\hat{\sigma}^2_{i,j} = \frac{1}{T-1}\sum_{t=1}^{T}(y_{t,i} - \bar{y}_i)(y_{t,j} - \bar{y}_j) \tag{15}$$

These sample estimates are unbiased, which means that on average the estimator is equal to the true value. One characteristic of the sample estimators is that the variance of the estimators lowers as the estimation period increases. And thus the sample moments are very accurate when the estimation period is very long. However, due to time varying parameters, the estimation period, that is used in practice, is very short. Because of this high estimation variance for the small estimation periods, the sample mean and covariance matrix contain large estimation errors in practice. Also, when the number of stocks increases by 1, from N to $N+1$, the number of different elements in the covariance matrix increases by $N+1$. Put together, the estimation error in the covariance matrix is big when the number of stocks is large, relative to the estimation period.

The Markowitz portfolio inherits the estimation error from the estimated moments. This often results in very large short positions or corner solutions, which are not feasible in practice or diverse, respectively. For this reason, the Markowitz portfolio does not perform well in practice. That is why in this paper different estimators are used as plug-in values for the formula that determines the portfolio weights. The collection of estimators that are used in this paper consist well known estimators that are commonly used in the industry. The methods that are used can be categorized in Shrinkage, Factor models, portfolio constraints and mixed estimation. The decision theoretic approach with a diffuse prior or a normal distributed prior won't be used, because it leads to similar results as the sample moments or James-Stein estimator, respectively.

2.2 James-Stein estimator

The first estimator is the James-Stein estimator. As discussed above, the sample mean is unbiased. However, James and Stein (1961) found that for $N \geq 3$ independent normal random variables the sample mean is dominated by a biased estimator in terms of the joint mean squared error (JMSE). The JMSE is a combination of the bias and the variance of the corresponding estimator $\hat{\mu}$:

$$\begin{aligned}\text{JMSE}(\hat{\mu}) &= E((\hat{\mu} - \mu)'(\hat{\mu} - \mu)) \\ &= \text{Bias}(\hat{\mu})'\text{Bias}(\hat{\mu}) + \text{tr}(\text{Var}(\hat{\mu}))\end{aligned}$$

This biased estimator is called the James-Stein estimator and it is a convex combination of the sample mean ($\bar{y}$) and a fixed constant value (μ_0), which is called the shrinkage target:

$$\mu^* = \delta\mu_0\iota + (1-\delta)\bar{y} \tag{16}$$

for $0 < \delta < 1$. The idea of the James-Stein is that it 'shrinks' the extreme values in the estimated sample mean ($\bar{y}$) towards a constant fixed value (μ_0). This is also why μ_0 is called the shrinkage target. This shrinkage target can be any constant value, where it is chosen to be the grand mean across all assets in this paper. This is also a popular shrinkage target in practice. The James-Stein estimator is a biased estimator. However, the variance of the estimator is reduced, relative to the sample mean, such that this new estimator has a lower estimation error in terms of the JMSE. This is done by choosing the optimal shrinkage factor δ:

$$\delta = \min\left[1, \frac{(N-2)/T}{(\bar{y} - \mu_0)'\Sigma^{-1}(\bar{y} - \mu_0)}\right] \tag{17}$$

Looking at the formula of the optimal shrinkage factor, δ increases when the number of assets (N) increases. This is intuitive, because when there are more means to be estimated, the estimation error in the sample mean will also increase. Also, δ decreases when the sample size (T) increases, which is logical due to the fact that a larger sample size results in a more precise sample mean. Finally, δ decreases when the distance between $\bar{y}$ and μ_0 increases. The James-Stein estimator of the excess returns is used as the plug-in value for $\tilde{\mu}$ for the tangency portfolio.

2.3 Black-Litterman

Another method that reduces the estimation error in the expected returns, is the Black-Litterman model (Black and Litterman, 1992). It uses Bayesian estimation to combine the equilibrium returns with a number (K) of subjective views of the investor about the expected returns of a number of assets. The resulting new vector of expected returns, the posterior distribution, causes logical portfolio weights and avoids extreme portfolio weights. According to Lee (2000), the problem of 'error-maximization', explained earlier, is also reduced when using the Black-Litterman model, because it spreads the estimation errors over the vector of expected returns. Equilibrium returns are the expected returns that correspond with the situation where the market is in equilibrium. The equilibrium returns are derived from the reverse formula of the mean variance portfolio optimization method, where the market capitalization weights are used as the vector of weights:

$$\mu_{equil} = \gamma \Sigma w_{mkt}$$

However, sometimes the investor disagrees with the implied equilibrium excess returns, because he has his own number (K) of views on the expected returns of some stocks. That is why the user is able to specify his views, which can be used as an input in the model:

$$p(\nu|\mu) = N(P\mu, \Omega),$$

where the $K \times N$ matrix P determines which stocks are involved in the views of the investor and how they are related to each other. The $K \times 1$ vector ν determines the absolute value of the views. The $K \times K$ covariance matrix Ω is a diagonal matrix with variances of the views representing the uncertainty in each view. The off-diagonal elements are 0, because it is assumed that the views are independent of each other. This paper uses the identity matrix as the P matrix, which means that the investor only has absolute views on the stocks expected returns. This paper uses the sample mean as the investor's views on the expected returns and the sample variances as the uncertainties in the views. Finally, the resulting vector of expected returns that leads to sensible and understandable portfolio weights is:

$$E[\mu|\nu] = [((\lambda\Sigma)^{-1} + P'\Omega^{-1}P]^{-1}[\lambda\Pi + P'\Omega^{-1}\nu] \tag{18}$$

Σ is the covariance matrix of the excess returns ($N \times N$ matrix).

γ is the risk aversion coefficient.

w_{mkt} is the market capitalization weight ($N \times 1$ column vector) of the assets

λ is a very small constant (usually $\frac{1}{T}$).

The posterior mean is a combination of the equilibrium returns (Π) and the investor's views (ν). When the uncertainty in the views (Ω) increases, the posterior mean will be further away from the investor's views and closer to the equilibrium returns. λ is defined as $\frac{1}{T}$, which means that when the sample size increases, less weight is given to the equilibrium returns. The resulting posterior mean of the excess returns is used as the plug-in value for $\tilde{\mu}$ in the tangency portfolio.

The Investor's Views

Besides the sample means and variances, this paper also uses more advanced investor's views from Beach and Orlov (2007). They use EGARCH-M(1,1) (exponential GARCH (Nelson, 1991) in combination with ARM-M (Engle et al., 1987)) models to generate the investor's view matrix ν and the covariance matrix representing the uncertainty in each view Ω. GARCH (generalized autoregressive conditional heteroscedasticity (Bollerslev, 1986)) models the conditional mean and variance and can be used to forecast them with the following equations:

$$y_t = \mu + \epsilon_t \tag{19}$$

$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2, \tag{20}$$

where y_t is the vector of excess stock returns and x_t a vector of exogenous variables. Because downward movements in the stock price usually result in a higher volatility than upward movements, Beach and Orlov use EGARCH to model the conditional variance. This model is able to capture asymmetric shocks in the variance. For the conditional mean, they use ARCH-M, which adds the conditional variance in the equation of the mean. They also include a set of extra regressors z_1 for the conditional mean (Premium, Term, Dividend Yield, Spread and Oil) and z_2 for the conditional variance (Premium, Term, Inflation, Dividend Yield and Spread):

$$y_t = \mu + \delta\hat{\sigma}_t^2 + \psi z_{1t} + \epsilon_t \tag{21}$$

$$\log(\sigma_t^2) = \omega + \beta\log(\sigma_{t-1}^2) + \alpha\left|\frac{\epsilon_{t-1}}{\sigma_{t-1}}\right| + \gamma\frac{\epsilon_{t-1}}{\sigma_{t-1}} + \phi z_{2t} \tag{22}$$

This model is used to create absolute views on the individual stocks by forecasting the conditional mean for period $t+1$. The variances which represent the uncertainty in the views is obtained by forecasting the conditional variance for period $t+1$.

2.4 Factor Models

Another approach to reduce the statistical error of the plug-in estimates used in this paper, is to impose a factor structure for the stock returns $y_{i,t}$, which are regressed on a number of (K) factors:

$$y_{i,t} = \alpha_i + \beta_i' f_t + \epsilon_{i,t},$$

where f_t contains the K factors. The implied mean and covariance matrix are:

$$\mu = a + B\mu_f$$
$$\Sigma = B\Sigma_F B' + \Sigma_\epsilon$$

Σ_F and Σ_ϵ are the covariance matrices of the factors and the errors respectively, where a and B are the OLS estimates of α and β respectively. In this paper it is assumed that the residuals are uncorrelated. This results in a drop in the number of unique parameters, which is reduced from $N(N+1)/2$ to $N(K+2) + K(K+1)/2$. This is done by setting $\Sigma_\epsilon = D_\epsilon$ (diagonal matrix). It is also assumed that the factors are uncorrelated, which reduces the number of unique parameters further to $K + N(K+2)$. These assumptions are not very realistic in most cases and results in an increase in estimation bias. However, empirical evidence shows that the decrease in estimation variance more than offsets the increase in estimation bias. To further reduce the number of parameters, the restriction that the expected returns do not deviate from the model is imposed (no mispricing). This is done by setting $\alpha = 0$. When this restriction is used, the investor is sure that the asset pricing model explains the returns of the assets. The implied mean and covariance matrix will be used together as estimators for $\tilde{\mu}$ and Σ.

Pástor

It is not very realistic that an asset pricing model perfectly explains the returns of the individual assets. For this reason, Pástor (2000) introduced a prior on α that determines how strong the belief is in the asset pricing model. Similar to Black and Litterman (1992), this is a form of mixed estimation. However, this time the economic views originate from an asset pricing model such as CAPM or Fama and French. Wang (2005) developed a new shrinkage like method, which they use to derive formulas for the implied mean and covariance matrix of the asset pricing model. These formulas are also used in this paper to compute the implied moments. Consider the following prior for α:

$$p(\alpha|\Sigma_\epsilon) = N(0, \theta\Sigma_\epsilon) \tag{23}$$

Then, the expression for the implied mean is as follows:

$$\mu = \omega\bar{\beta}\hat{\mu}_2 + (1-\omega)\hat{\mu}_1, \tag{24}$$

where $\bar{\beta}$ is the maximum likelihood estimate for β, when α is 0, and where $\hat{\mu}_1$ and $\hat{\mu}_2$ are the sample means of the excess return ($y_{i,t}$) and the factors (f_t), respectively. The expression for the implied covariance matrix is as follows:

$$\Sigma \;=\; b\big(\omega\bar{\beta} + (1-\omega)\hat{\beta}\big)\hat{\Omega}_{22}\big(\omega\bar{\beta} + (1-\omega)\hat{\beta}\big)' + h\big(\omega\bar{\delta} + (1-\omega)\hat{\delta}\big)\big(\omega\bar{\Sigma} + (1-\omega)\hat{\Sigma}\big), \tag{25}$$

where $\hat{\beta}$ is the maximum likelihood estimate for β, when α is not 0, and where $\bar{\Sigma}$ and $\hat{\Sigma}$ are the implied covariance matrices for Σ when α is 0 and not 0, respectively. The rest of the formulas needed for the expressions above are:

$$\begin{aligned}
\omega &= \frac{1}{1 + T\theta/(1+\hat{S}^2)} \\
\bar{\delta} &= \frac{T(T-2)+k}{T(T-k-2)} - \frac{k+3}{T(T-k-2)} \cdot \frac{\hat{S}^2}{1+\hat{S}^2} \\
\hat{\delta} &= \frac{(T-2)(T+1)}{T(T-k-2)} \\
b &= \frac{T+1}{T-k-2} \\
h &= \frac{T}{T-m-k-1}
\end{aligned}$$

Where $\hat{S}^2$ is the highest Sharpe-ratio of a portfolio that lies on the efficient frontier, which is computed using the sample mean and covariance matrix of the factor portfolios. And because the variance of the prior should not be higher than the assumed variance σ_α^2, the following equation should hold:

$$\theta \times \max(\mathrm{diag}(\Sigma_\epsilon)) = \sigma_\alpha^2 \tag{26}$$

In this paper different choices of factors are made. The first choice is based on economic theory and is the CAPM model. Here the market portfolio is used as the only factor. The second choice is the Fama and French (FF) model, which is based on empirical work. This is a three factor model which contains also the market portfolio. Besides the market portfolio, the other two factors are "Small Minus Big" (SMB) and "High Minus Low" (HML). SMB is the return on smallest stocks minus returns on largest stocks and HML is the returns on high B/M (Book to Market ratio) stocks minus returns on low B/M stocks. The last choice are the first couple of principal components, which explain enough of the variance of the data and still reduces the number of unique parameters sufficiently.

2.5 Shrinkage of the Covariance Matrix

The CAPM factor model implied covariance matrix described in the previous section has a lot of bias, because the model is strict and misspecified. However, it has a small estimation variance. The sample covariance matrix (S) on the other hand, has a large estimation variance but is unbiased (asymptotically). Ledoit and Wolf (2003b) developed an estimator that tries to take the best of both worlds. Similar to what the James-Stein estimator does for the expected return, they propose to estimate the covariance matrix of stock returns by an optimally weighted average of an unbiased estimator with a high estimation variance, the sample covariance matrix, and a biased estimator with a lower estimation variance, the single-index model from Sharpe (1963). The idea is that the decrease in estimation variance outweighs the increase in bias. The single-index implied covariance matrix is more structured and stable compared to the sample covariance matrix. Therefore, the optimal weight α determines how much structure is set on the estimator.

The following shrinkage estimator ($\hat{\Sigma}_{Shrink}$) is an optimal combination between these two extreme estimates.

$$\hat{\Sigma}_{Shrink} = \hat{\alpha}F + (1 - \hat{\alpha})S \tag{27}$$

S is the sample covariance matrix and F is shrinkage target. Ledoit and Wolf (2003b) used the covariance matrix implied by the single-index model as their shrinkage target:

$$y_{i,t} = \alpha_i + \beta_i r_{m,t} + \epsilon_{i,t},$$

which implies the following covariance matrix:

$$F = \sigma_m^2 \hat{\beta}\hat{\beta}' + \hat{\Sigma}_\epsilon,$$

where σ_m^2 is the variance of the market factor $r_{m,t}$, $\hat{\beta}$ is the vector containing the estimated values of β and $\hat{\Sigma}_\epsilon$ is the diagonal matrix, which contains residual variance estimates.

This paper also uses another shrinkage target, namely the constant correlation model from Ledoit and Wolf (2003a). The constant correlation model replaces all the correlations of the asset returns to the other assets, with the average of all correlations. This way all correlations are the same. On the diagonal of the shrinkage target are the sample variances. Then $\hat{\alpha}$ is defined as:

$$\hat{\alpha} = \frac{1}{T}\frac{\pi - \rho}{\gamma} = \frac{\kappa}{T}$$

$$\pi = \sum_{i=1}^{N}\sum_{j=1}^{N} AsyVar[\sqrt{T}s_{ij}]$$

$$\rho = \sum_{i=1}^{N}\sum_{j=1}^{N} AsyCov[\sqrt{T}f_{ij}, \sqrt{T}s_{ij}]$$

$$\gamma = \sum_{i=1}^{N}\sum_{j=1}^{N} (\phi_{ij} - \sigma_{ij})^2$$

The optimal shrinkage factor α is defined such that the expected quadratic distance between the true covariance matrix and $\hat{\Sigma}_{Shrink}$ is minimized.

If κ were known, $\frac{\kappa}{T}$ can be used as the optimal shrinkage factor. However, κ is unknown. This means that a consistent estimator for κ is needed. This is an estimator that converges to the true value of κ, when the number of observations increases. The consistent estimators of π, ρ, and γ, which lead to the consistent estimator of κ, can be found in the Appendix.

This results in the consistent estimator for κ:

$$\hat{\kappa} = \frac{\hat{\pi} - \hat{\rho}}{\hat{\gamma}}, \tag{28}$$

which leads to the formula used for the shrinkage intensity:

$$\hat{\alpha}^* = \max\left(0, \min\left(\frac{\hat{\kappa}}{T}, 1\right)\right) \tag{29}$$

The shrinkage estimators, with both the single-index model and the constant correlation model as shrinkage targets, are used as plug-in values for Σ. They are used in the GMV portfolio and the tangency portfolio, where it is combined with the Black-Litterman and James-Stein estimators for $\tilde{\mu}$.

2.6 Portfolio Restrictions

This paper uses short-sale constraints as a way to reduce the estimation error in the portfolio weights, which means that the individual weights can't go below zero. Empirical findings suggest that imposing portfolio restrictions reduces the estimation error of the portfolio weights based on sample moments (Frost and Savarino, 1988). In the previous sections shrinkage was used to compute estimators for the expected returns (James and Stein, 1961) and the covariance (Ledoit and Wolf, 2003b,a). Jagannathan and Ma (2003) show that imposing short-sale constraints are actually quite similar to these methods, because it can be explained as a form of shrinkage. When the following mean-variance optimization is used to determining the portfolio weights, where the variance is minimized for a certain target portfolio return μ_p:

$$\min_w \frac{1}{2} w' \Sigma w \text{ s.t.} \tag{30}$$

$$w' \mu = \mu_p \tag{31}$$

$$w \leq \bar{w} \tag{32}$$

$$w \geq 0 \tag{33}$$

This leads to the following Lagrangian:

$$L = \frac{1}{2} w' S w - \theta(w'\mu - k) - \lambda' w + \delta'(w - \bar{w}) \tag{34}$$

Jagannathan and Ma (2003) show that this is the same problem as the unconstrained mean-variance problem when μ is replaced by $\mu + \frac{1}{\theta}\lambda - \frac{1}{\theta}\delta$ with target return $\mu_p - \frac{1}{\theta}\delta'\bar{w}$. This means that if the non negativity constraint is binding, the expected return of asset i is moved upward by $\frac{1}{\theta}\lambda_i$. The same goes for the upper bound, which lowers the expected return by $\frac{1}{\theta}\delta_i$ if the constraint is binding. This is similar to shrinkage, because stocks with low expected returns will likely have negative weights, low expected returns will be adjusted upwards. The opposite goes for stocks with high expected returns.

Jagannathan and Ma (2003) give a similar explanation for why imposing portfolio constraints in a global minimum variance portfolio is a form of shrinkage of the covariance matrix. They show and prove that the following adjusted covariance matrix:

$$\tilde{S} = S + (\delta\iota' + \iota\delta') - (\lambda\iota' + \iota\lambda') \tag{35}$$

is symmetric and positive semi-definite. And that the solution to the unconstrained GMV problem using $\tilde{S}$ is the same solution as for the constrained GMV problem using S.

So imposing non negativity constraints is the same as reducing the covariance of stock i with the other stocks, σ_{ij}, by $\lambda_i + \lambda_j$, which is positive if $\omega_i < 0$, and reducing the variance of stock i, σ^2_{ii}, by $2\lambda_i$. In the unconstrained portfolio variance minimization, stocks with large negative weights tend to have high covariances with other stocks. And since the highest covariance estimates more likely contain a positive estimation error, imposing non negativity constraints reduces the estimation error. The opposite goes for stocks stocks with low covariance and low variance. This adjustment is quite similar to a shrinkage effect. In this research, both the mean-variance optimized portfolio as the global minimum variance portfolio are used in combination with short-sale restrictions in the hope that the reduction in estimation error outweighs the specification error caused by the constraints.

2.7 1/N Rule as a Shrinkage Target

Besides the methods that reduce the estimation error in the inputs, this paper also uses a method that reduces the estimation error in the portfolio weights. As discussed earlier, the Markowitz portfolio weights perform bad out-of-sample due to the estimation error in the plug-in values. Tu and Zhou (2011) suggest a combination rule between the Markowitz portfolio ($\frac{1}{\gamma}\Sigma^{-1}\mu$, where $\gamma = 5$) using the sample moments as the plug-in values and the $1/N$ portfolio:

$$w_c = (1 - \delta)w_e + \delta\bar{w}, \tag{36}$$

where w_e are the $1/N$ weights, $\bar{w}$ is the Markowitz rule using the sample moments and δ is the combination coefficient that lies between 0 and 1. Compared to the true Markowitz portfolio, $\bar{w}$ is asymptotically unbiased but has a high variance in small samples, whereas the $1/N$ has a high bias but has zero variance. They derive an optimal δ that minimizes the expected distance between the expected utility of the true Markowitz portfolio and of $\hat{w}_c$. This true δ is unknown in practice, which is why the following estimator is used:

$$\hat{\delta} = \frac{\hat{\pi}_1}{\hat{\pi}_1 + \hat{\pi}_2}, \tag{37}$$

where

$$\hat{\pi}_1 = w_e'\hat{\Sigma}w_e - \frac{2}{\gamma}w_e'\hat{\mu} + \frac{1}{\gamma^2}\tilde{\theta}^2 \tag{38}$$

$$\hat{\pi}_1 = \frac{1}{\gamma^2}(c_1 - 1)\tilde{\theta}^2 + \frac{c_1}{\gamma^2}\frac{N}{T} \tag{39}$$

Theoretically $\hat{w}_c$ should improve on both w_e and $\bar{w}$. However, due to estimation error in $\hat{\delta}$, it does not always outperform both counterparts of the combination. Nevertheless, Tu and Zhou (2011) show that the estimation error in δ is generally small, which means that it outperforms $\bar{w}$ and performs close to or better than the $1/N$ rule. To let the weights in w_c sum to 1, it is divided by the sum of the weights:

$$\frac{w_c}{|\iota' w_c|}$$

2.8 Evaluation Techniques

Different evaluation methods are used to evaluate the different methods, described in Table 1, to estimate the expected returns and the covariance. These evaluation methods are described in the following sections. The methods that perform the best in reducing the estimation error, are combined to be used as inputs for the tangency portfolios based on the real stock data.

2.8.1 Simulation Experiment

To evaluate the estimation error of the mean and covariance matrix, a simulation experiment, similar to what Jobson and Korkie (1980) have done, is executed. For every model in Table 1,

Table 1: Methods to reduce the estimation error in the mean and covariance matrix

#	**Methods**
1	The sample mean and covariance matrix
2	The James Stein estimator with the following shrinkage targets: Grand mean across all assets, mean of the GMV portoflio, implied mean of the single index model
3	Black and Litterman combined with the sample covariance matrix
4	Ledoit and Wolf estimator with the following shrinkage targets: Single index model and the constant correlation model
5	Factor Models (with and without the prior of Pastor (2000)) with the following factor choices: the CAPM model, Fama and French and the principal components
6	Portfolio restrictions: short-sale and upper bound restrictions

This table lists the methods used to estimate the moments in the simulation experiments

the following simulation experiment is done: Stock data is simulated, $N = 10$ stocks and length T, for a certain mean and covariance matrix. Estimators for mean and covariance matrix are obtained by using the historical sample moments for N stock returns. Then these estimated moments are assumed to be the true moments for the simulations. 250 independent datasets are simulated, which results in a sufficiently low Monte-Carlo error, using a multivariate normal distribution for different periods T (25, 75, 150). From these datasets, tangency portfolios are constructed for the 250 simulated datasets using the plug-in estimates of the different methods. To evaluate the estimated tangency portfolios, the difference between the true Sharpe-ratio and the Sharpe-ratio that belongs to the estimated tangency portfolio is used as a loss function:

$$L = SR_{true} - \hat{SR}, \tag{40}$$

where the true Sharpe-ratio is the Sharpe-ratio that belongs to the tangency portfolio that uses the moments that are used to simulate the 250 datasets.

To evaluate the estimation methods of the mean separately, methods 1, 2, 3 and 5 from Table 1 are used to estimate the mean of the simulated samples, where the return of the market weighted portfolio is used as the market factor for the CAPM model. The joint mean squared error is used as a loss function to evaluate and compare the methods:

$$MSE_J(\hat{\mu}) = E\big((\hat{\mu} - \mu)'\Sigma^{-1}(\hat{\mu} - \mu)\big) \tag{41}$$

To evaluate the estimation methods of the covariance matrix separately, methods 1, 4 and 5 from Table 6 are used to estimate the covariance matrix of the simulated samples, where the return of the market weighted portfolio is used as the market factor for the CAPM model. The joint mean squared error, from Patton and Sheppard (2007), is used as a loss function for the covariance matrix to evaluate and compare the methods:

$$L(\hat{\Sigma}, H; b) = \begin{cases} \frac{2}{b+2}\mathrm{tr}\big(\hat{\Sigma}^{b+2} - H^{b+2}\big) - \mathrm{tr}\big(H^{b+1}(\hat{\Sigma} - H)\big), & \text{if } b \neq -1, -2 \\ \mathrm{tr}\big(H^{-1}\hat{\Sigma}\big) - \log|H^{-1}\hat{\Sigma}| - K, & \text{if } b = -2 \end{cases} \tag{42}$$

$\hat{\Sigma}$ is the estimated $N \times N$ covariance matrix, H is the true covariance matrix and b is a shape parameter. b is chosen to be 0 because, according to Patton and Sheppard (2007) it results in the robust loss-function mean squared error(MSE).

2.8.2 Tangency and Minimum-Variance Portfolios

To study the out of sample performance of the estimated moments, tangency and minimum-variance portfolios for a set of assets described in Table 6 are used. On the basis of the results of the simulation experiment, combinations of the best performing methods are used to estimate the mean and the covariance matrix.

A "rolling-window" approach is used: To estimate μ and Σ a sample of M months will be used. Starting in period $t = M + 1$, the portfolio weights are obtained using the estimated parameters that are estimated using the sample up to period M. These weights are used to compute the return of month $M + 1$. Then the first observation is deleted from the estimation sample and the next period is included (First In First Out). This is continued until the end of the dataset is reached. At the end, there will be $T - M$ out-of-sample portfolio returns, which can be evaluated. These out of sample returns of every portfolio will be evaluated by their Sharpe-ratio:

$$SR = \frac{E(y_{t+1} - R_f)}{\sqrt{Var(y_{t+1} - R_f)}}, \tag{43}$$

To test if the Sharpe-ratios are significantly different from each other, the 'two-sample statistic for comparing Sharpe-Ratios' from Opdyke (2007) is used.

In addition to the Sharpe-Ratios, the average expected excess return and volatility of the portfolios are shown. One thing to note is that the out-of-sample Sharpe-ratio is not the same as dividing the average expected excess return by the average volatility, because:

$$\frac{\frac{1}{N}\sum_{i=1}^{50}\mu_i}{\frac{1}{N}\sum_{i=1}^{50}\sigma_i} \neq \frac{1}{N}\sum_{i=1}^{50}\frac{\mu_i}{\sigma_i}, \tag{44}$$

where μ_i and σ_i are the expected excess return and the volatility of the portfolio of dataset i.

To evaluate if the portfolios perform different relative to each other in different states of the economy, they will be evaluated conditional on market volatility. This is done by estimating the realized volatility of the NASDAQ and NYSE market indices. And splitting the period of the dataset in to 5 quantiles based on volatility.

Another measure that is used to evaluate the portfolios is the portfolio turnover. The portfolio turnover measures how frequently and how much the assets in the portfolio are bought and sold. High portfolio turnover is undesirable in the presence of transaction costs, because it reduces the profit that is made from the portfolio return. This paper uses the same definition for the portfolio turnover as DeMiguel et al. (2009):

$$\text{Turnover} = \frac{1}{T-M}\sum_{t=1}^{T-M}\sum_{j=1}^{N}\Big(|\hat{w}_{k,j,t+1} - \hat{w}_{k,j,t^+}|\Big), \tag{45}$$

where $\hat{w}_{k,j,t}$ is the portfolio weight in asset j at time t under strategy k, $\hat{w}_{k,j,t^+}$ is the portfolio weight before rebalancing at $t + 1$; and $\hat{w}_{k,j,t+1}$ is the desired portfolio weight at time $t + 1$, after rebalancing.

2.8.3 Model Confidence Set

A Model Confidence Set is a method that finds a subset out of all the models, which contain the best models in terms of a given loss function for a certain level of confidence. Consider a set M^0 which contains a finite number of objects indexed by $i = 1, ..., m_0$. The objects in the set are evaluated in terms of their loss function, L_t. This loss function is variable that is defined by the use and represents how 'bad' the model performs on time t. Let us define the relative performance variables $d_{ij,t}$ at time t for all $i, j \in M^0$ in terms of their loss function $d_{ij,t} = L_{i,t} - L_{j,t}$, which are defined in section 2.8.1, and assume that $u_{ij,t} = E(d_{ij,t})$ is finite and does not depend on t. The set containing the superior models is then defined by $M^* = i \in M^0$: $u_{ij} \leq 0$ for all $j \in M^0$: u_{ij} . In words this means that alternative $i \in M^0$ is preferred to alternative $j \in M^0$ if $u_{ij} \leq 0$. The objective of the MCS process is to determine this set M^*. This is done by using an equivalence test, δ_m and an elimination rule e_m. The equivalence test δ_M tests for a confidence level of α if $H_0, M : u_{ij} = 0$ for all $i, j \in M$, M sub sample of M^0. This is done by the t-statistic:

$$t_{i,j} = \frac{d_{i,j}}{\sqrt{v\hat{a}r(d_{i,j})}}, \tag{46}$$

where $v\hat{a}r(d_{i,j})$ is estimated using stationary bootstrap with a block length of $l = 2$ (the results are robust for different block lengths) and the number of bootstrap samples of $B = 10000$. The elimination rule e_M gives the object of M that has to be removed from M in case $H_{0,M}$ is rejected. The algorithm behind the MCS procedure is as follows:

Step 1 Initially set $M = M^0$.

Step 2 Test $H_{0,M}$ using δ_M at level α.

Step 3 Define $M^*_{1-\alpha} = M$ if $H_{0,M}$ rejected, use e_M to remove an object from M and repeat the algorithm from step 1.

The MCS algorithm stops if $H_{0,M}$ is accepted and gives us a set estimate $M^*_{1-\alpha}$ consisting of the set of surviving objects as a result. This $M^*_{1-\alpha}$ is referred as the model confidence set.

3 Data

The data consists of monthly excess (over the risk free rate) returns of American stock data, which is obtained from the Datastream database. The 3-month US T-bill is used as the risk free rate. The Fama-French factors are obtained from Ken French's website. The individual stocks are obtained from two different stock exchanges:

NYSE 3652 stocks, 6/1/1976-5/1/2016, 481 monthly observations

NASDAQ 4722 stocks, 5/1/1986-5/1/2016, 361 monthly observations

Financial firms (SIC code between 6000 and 6799) are excluded, because "high leverage is normal for financial firms, where it more likely indicates distress in non-financial companies" (Fama and French, 2008, 1992).

3.1 Dynamic Dataset

The datasets described above consist of (N) thousands of individual stocks, therefore the estimation window has to consist of ($T > N$) observations in order to compute the covariance matrix and to make sure it is not singular. This amount of monthly observations is not available and because of this, a random subset of companies is created from the entire dataset. The "rolling-window" approach, explained in section 2.8.2 above, is applied on this subset of stocks. The subset of companies is constant over the whole period T and for all the methods. The companies in this subset don't have a lifespan that is equal to the length of the dataset. Some companies go bankrupt or are delisted from the stock exchange for another reason. New companies enter the stock exchange in the midst of the dataset. For every period only the companies that are available during the whole estimation window of 180 months will be included in the portfolio. This means that new companies can enter the portfolio when they arise and companies that go bankrupt or are not traded anymore leave the portfolio. This 'dynamic' dataset is used to avoid selection (survivorship) bias. This means that the dataset is not limited to stocks that are available during the whole time span of the dataset. To not be limited to one random subset of companies, 50 random subsets are used to compute portfolios. This is done for both the NYSE and NASDAQ dataset, which results in 100 datasets that contain a random selection of stocks. The number of stocks that are in the portfolio (N) can change every period. The number of stocks in the portfolio ranges approximately between 5 - 20 stocks and 4 - 25 stocks for the NYSE and NASDAQ dataset, respectively. The size of the random selection of stocks is chosen so that the portfolios contain not too little and not too much stocks.

3.2 Descriptive Statistics

In Table 2 the average of the mean, volatility, skewness and kurtosis of the stocks are reported. From Table 2 and Figure 1 it is clear that the stocks on the NASDAQ stock exchange have a higher volatility and mean than the stocks on the New York stock exchange.

Table 2: averages of μ, σ, skewness and kurtosis of excess monthly stock returns for both datasets.

	NYSE	NASDAQ
μ	0.92%	1.55%
σ	11.54%	19.20%
skewness	0.67	1.48
kurtosis	9.04	13.23

Figure 1: Plots of μ, σ of the excess stock returns.

(a) NYSE **(b)** NASDAQ

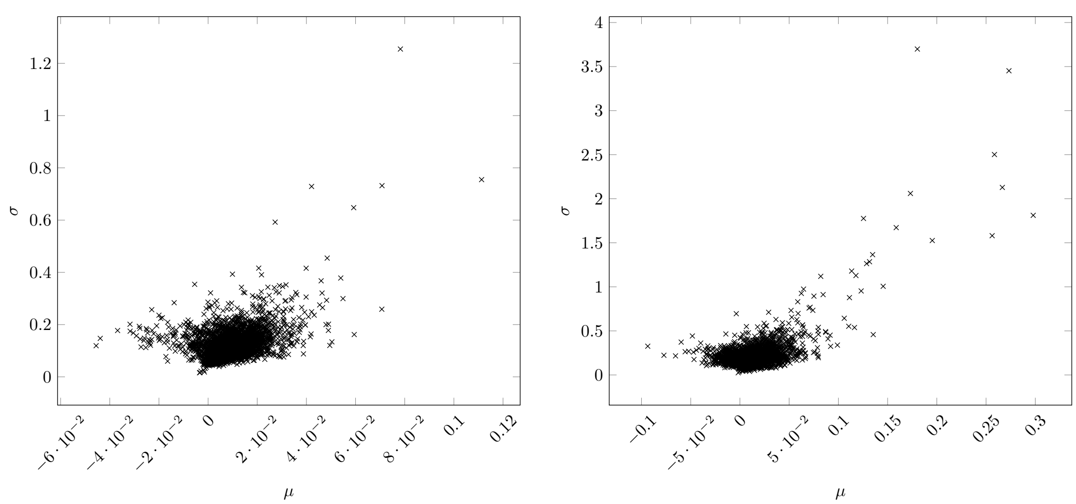

These figures (a and b) plot the mean against the volatility of the excess monthly stock returns.

Figure 2 shows the realized monthly volatility of the market indices from both the New York and the NASDAQ stock exchanges. The realized volatility is estimated using daily returns and then multiplied by $\sqrt{21}$ to compute the monthly volatility. These plots show the state of the economy as the indices are capitalization weighted indices of all the stocks traded on the corresponding stock exchanges. The high volatility spikes usually correspond with a bear market (recession), whereas periods of low volatility mostly appear in bull markets (expansion). Both the stock exchange have high volatility spike in 1987 and 2009, which both correspond with known recessions. However the NASDAQ stock exchange has another high volatility spike around the period of the year 2001, which corresponds with the burst of 'dot-com' bubble. The New York stock exchange index also has a increase of volatility in this period. However, because the NASDAQ trades mostly stocks from technology companies, the effect of this crisis is larger for this stock exchange. Because of these differences in characteristics between the two stock exchanges, they are evaluated separately in the results section. This is done because it is interesting to see how the methods perform in more volatile stock exchange such as the NYSE compared to the NASDAQ.

Figure 2: Plots of the Market Volatility

(a) NASDAQ **(b)** NYSE

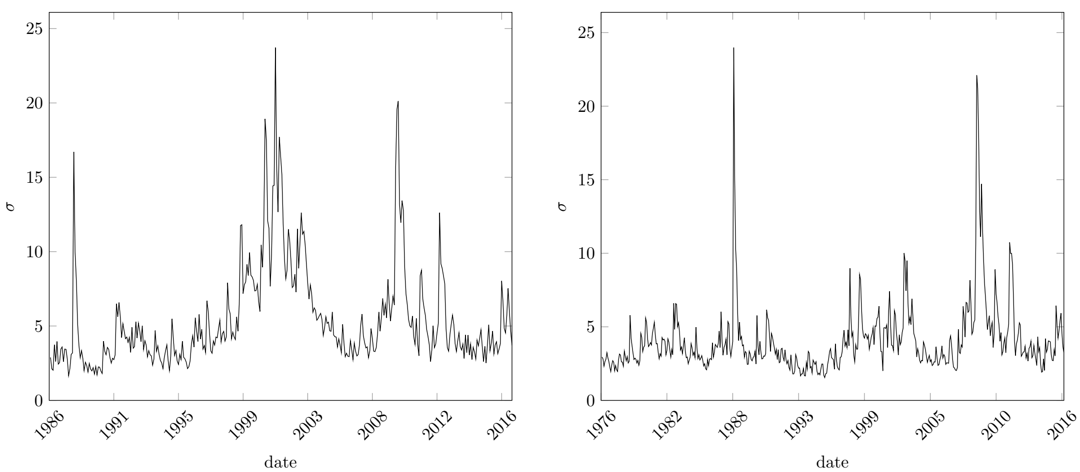

These plots show the estimated realized volatility of the capitalization weighted indices of both the New York and the NASDAQ stock exchange. The volatility is estimated using daily returns and is multiplied by $\sqrt{21}$ to compute the monthly volatility

4 Results

4.1 Simulation Experiment

This section contains the results of the simulations experiment described in section 2.8.1. Table 3 reports the average Sharpe-ratios of the tangency portfolios using the moments estimated by the different models. In the figures 3 to 7, in the appendix, histograms of these Sharpe-ratios can be found for each method and each estimation period.

Table 3: Sharpe-ratios of simulated data

Panel A: $T = 25$

	B-L	*CAPM*	$w > 0$	*J-S*	*Ledoit SI*	*Ledoit CC*	*Pastor*	*PCA*	*sample*
μ	17.8%	18.1%	19.2%*	10.6%	8.3%	7.1%	18.1%	8.0%	8.7%
σ	1.5%	1.7%	3.1%	7.5%	8.8%	8.0%	1.8%	8.1%	8.0%
min	13.7%	12.6%	8.4%	-11.0%	-14.8%	-16.0%	12.6%	-13.5%	-10.8%
max	21.5%	22.7%	24.9%	26.3%	25.9%	24.3%	22.7%	23.9%	26.2%
MCS	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00

Panel B: $T = 75$

	B-L	*CAPM*	$w > 0$	*J-S*	*Ledoit SI*	*Ledoit CC*	*Pastor*	*PCA*	*sample*
μ	17.7%	18.4%	21.0%*	17.2%	15.5%	14.3%	18.6%	14.5%	15.7%
σ	1.9%	0.9%	2.8%	6.1%	7.2%	7.3%	1.0%	7.2%	7.1%
min	11.3%	16.2%	12.0%	-10.3%	-14.7%	-14.0%	15.1%	-13.1%	-13.8%
max	21.8%	20.8%	25.6%	25.3%	25.5%	25.0%	21.0%	25.1%	25.2%
MCS	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00

Panel C: $T = 150$

	B-L	*CAPM*	$w > 0$	*J-S*	*Ledoit SI*	*Ledoit CC*	*Pastor*	*PCA*	*sample*
μ	17.7%	18.5%	22.4%*	20.6%	19.7%	19.0%	18.8%	18.4%	19.8%
σ	2.6%	0.7%	2.1%	3.5%	4.4%	4.8%	0.8%	5.8%	4.4%
min	9.8%	15.9%	14.5%	-1.5%	-8.7%	-13.9%	16.0%	-11.1%	-9.2%
max	22.7%	20.4%	26.5%	26.7%	26.6%	26.4%	20.7%	27.3%	26.7%
MCS	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00

This table shows the mean, standard deviation, minimum and maximum of the 250 Sharpe ratios of the tangency portfolios for the methods used in the simulation experiment. The Sharpe ratio that belongs to the true mean and covariance is 28.1%. The p-value of the model confidence set determines if the method sits in the set of the best methods (5%) confidence level. If the p-value is lower than 5%, the method is significantly different from the best method in terms of the defined loss function.

From Table 3 it is clear that the method that uses short-sale constraints performs the best for all the estimation periods. Because it is the method with the highest average Sharpe-ratio and is the only method that $M^*_{1-\alpha}$ set consists of. The tangency portfolio that uses the sample moments without any constraints performs bad for the small estimation periods, but performs better as the estimation period increases. This is true for more methods like Ledoit and Wolf 2003 and 2004 and principal component implied moments. The moments that these methods produce, result in volatile portfolio weights, which results in volatile Sharpe-ratios. They perform better as the sample size increases, but they still produce volatile portfolio weights compared to methods like Black-Litterman, short-sale constraints, CAPM (with Pástor prior) implied moments. The James-Stein expected returns performs poorly for a low estimation period of $T = 25$, but performs much better for higher estimation periods. It even is the next best method for the estimation period of 150 in terms of average Sharpe-ratio.

The histograms 3 to 7, in the appendix, give more insight to the way in which certain portfolios perform better than others. For example Figure 3 clearly shows how short-sale constraints can reduce the volatility and the sensitivity to changes in the inputs of the weights. The Sharpe-ratios for the constrained portfolios are notably closer to the true Sharpe-ratio and have a considerably lower volatility. Figure 4 shows that the Black-Litterman posterior mean and CAPM implied moments result in very stable portfolio weights even for small estimation periods. Figure 6 shows that using the sample mean in combination with a the shrinkage estimators for the covariance matrix as inputs does not improve the performance of the portfolios much. The results look similar to that of the sample moments. The estimation error in the mean has a much bigger influence on the portfolio weights than that of the covariance matrix.

Table 4: Simulation experiment for μ

μ	T=25		T=75		T=150	
	JMSPE	mcs	JMSPE	mcs	JMSPE	mcs
sample	0.395	0.00	0.134	0.00	0.066	0.00
B-L	0.074*	1.00	0.068	0.00	0.056	0.00
Pastor (5%)	0.129	0.00	0.070	0.00	0.045	0.00
CAPM	0.129	0.00	0.070	0.00	0.046	0.00
PCA	0.093	0.00	0.086	0.00	0.076	0.00
J-S	0.186	0.00	0.045*	1.00	0.034*	1.00

This table reports the joint mean squared error of the estimated μ for the different methods. An * denotes the method that performs the best in terms of the JMSPE. The p-values in the 'mcs' columns show if the corresponding significantly differ from the best model. A critical value of 5% is used.

Table 4 shows that the best performing method for estimating the μ in terms of the JMSPE is the Black-Litterman model for estimation period of $T = 25$ and the James-Stein estimator for estimation periods of $T = 75$ and $T = 150$.

Table 5: Simulation experiment for Σ

Σ	T=25		T=75		T=150	
	JMSPE	mcs	JMSPE	mcs	JMSPE	mcs
Pastor (5%)	2.07E-03	0.00	3.00E-03	0.00	8.07E-04	0.00
CAPM	3.26E-04	0.00	9.02E-05	0.00	2.38E-06*	1.00
Ledoit SI	4.30E-04	0.00	2.76E-04	0.00	9.68E-05	0.00
Ledoit CC	1.31E-04*	1.00	4.21E-05*	1.00	6.84E-06	0.00
PCA	7.17E-04	0.00	5.21E-04	0.00	1.69E-04	0.00
sample	5.58E-04	0.00	2.74E-04	0.00	9.30E-05	0.00

This table reports the joint mean squared error of the estimated covariance matrix for the different methods. An * denotes the method that performs the best in terms of the JMSPE. The p-values in the 'mcs' columns show if the corresponding significantly differ from the best model. A critical value of 5% is used.

Table 5 shows that the best performing method for estimating the Σ in terms of the JMSPE is the method from Ledoit and Wolf (2003a) for the estimation periods of 25 and 75 months and the CAPM implied Σ for the estimation period of 150 months.

The best performing methods from the simulation experiments are the Black-Litterman model and the James-Stein estimator for estimating the μ and the CAPM implied covariance matrix and the estimator from Ledoit and Wolf (2003a) for estimating the covariance matrix. Also the short-sale constraints perform the best in terms of providing low volatile and small biased portfolio weights. In the next section these methods will be combined to compute tangency portfolio weights for empirical stock data next to all the other methods from Table 6. This results in the following collection of portfolios that will be tested on the stock data:

Table 6: Techniques that are used to compute the portfolio weights

#	Methods
0	$1/N$ rule
	Tangency Portfolios
1	Sample Moments
2	Shrinkage Weights from Tu and Zhou (2011)
3	Black-Littermean μ, Sample Σ
4	Black-Littermean μ with Investor's Views from Beach and Orlov (2007), Sample Σ
5	James-Stein μ, Sample Σ
6	CAPM Implied Moments
7	CAPM Implied Moments Using the Prior from Pástor
8	Fama and French Implied Moments
9	Fama and French Implied Moments Using the Prior from Pástor
10	PCA Implied Moments
11	Short-Sale Constraints, Using Sample Moments
	Minimum Variance Portfolios
12	Sample Covariance
13	Σ from Ledoit and Wolf (2003b)
14	Σ from Ledoit and Wolf (2003a)
15	Short-Sale Constraints, Sample Σ
16	Short-Sale Constraints, Σ from Ledoit and Wolf (2003b)
17	Short-Sale Constraints, Σ from Ledoit and Wolf (2003a)
	Short-Sale Constrained Tangency Portfolios
18	Black-Litterman μ, CAPM Implied Σ
19	Black-Litterman μ, Σ from Ledoit and Wolf (2003a)
20	James-Stein μ, CAPM Implied Σ
21	James-Stein μ, Σ from Ledoit and Wolf (2003a)

This Table shows a list of all methods that are used to compute the portfolio weights. It also shows what techniques are used to estimate the moments that are required as inputs to compute the portfolio weights

4.2 Stock Data Results

Table 7 shows the average over the 50 datasets of the out-of-sample Sharpe-ratios. The columns '$> 1/N$' and '$< 1/N$' show how often, out of the 50 datasets, the corresponding method outperformed the $1/N$ rule and was outperformed by $1/N$ rule, respectively. To test if the Sharpe-ratios are significantly different from each other, the 'two-sample statistic for comparing Sharpe-Ratios' from Opdyke (2007) is used. Table 8 shows the average over the 50 datasets of the expected excess returns and average volatility of the portfolio returns.

The results differ slightly between the two stock exchanges NASDAQ and NYSE. For the NASDAQ stock exchange, the $1/N$ is outperformed by the tangency portfolio that uses Fama-French 3-factor model implied mean and covariance matrix in 1 out of the 50 datasets for the estimation window of 120 months. It also has a higher average Sharpe-ratio for the estimation period of 180 months. Another portfolio that outperforms the $1/N$ portfolio in 1 out of the 50 datasets is the GMV portfolio that uses the shrinkage estimator from Ledoit and Wolf (2003a) with the constant correlation model as a shrinkage target. Also for the estimation periods of 120 and 180 months it is never outperformed by the $1/N$ portfolio in any of the 50 datasets. For the portfolios that are based on the stocks from the New York stock exchange, none of the considered portfolios outperform the $1/N$ portfolio across all estimation periods. Hence, the $1/N$ portfolio performs worse for the NASDAQ dataset relative to the other portfolios. The stocks traded on the NASDAQ stock exchange on average have a considerably higher volatility

Table 7: Sharpe-ratios for stock data

Panel A: NASDAQ

	T=60			T=120			T=180		
	S-R	$> 1/N$	$< 1/N$	*S-R*	$> 1/N$	$< 1/N$	*S-R*	$> 1/N$	$< 1/N$
$1/N$	13.22%			13.22%			13.22%		
MV/in sample	35.22%	42%	0%	137.94%	78%	0%	248.12%	90%	0%
MV/sample	-3.97%	0%	36%	0.46%	0%	30%	4.03%	0%	22%
MV/combination	11.23%	0%	0%	12.62%	0%	0%	12.77%	0%	0%
MV/B-L	7.36%	0%	14%	7.45%	0%	16%	7.47%	0%	16%
MV/B-L (inv. views)	5.19%	0%	16%	7.79%	0%	20%	8.14%	0%	16%
MV/J-S	-0.30%	0%	28%	9.92%	0%	8%	11.04%	0%	4%
MV/CAPM	3.85%	0%	2%	3.63%	0%	4%	12.46%	0%	0%
MV/CAPM/Pastor (5%)	0.31%	0%	18%	0.28%	0%	30%	11.99%	0%	0%
MV/FF	1.63%	0%	32%	11.36%	2%	0%	13.57%	0%	0%
MV/FF/Pastor (5%)	0.50%	0%	28%	10.45%	2%	2%	13.06%	0%	0%
MV/PCA	-1.67%	0%	26%	6.04%	0%	10%	2.53%	0%	22%
MV/sample/w>0	7.11%	0%	20%	9.91%	0%	8%	9.23%	0%	12%
GMV	11.65%	0%	6%	13.08%	0%	2%	13.09%	0%	0%
GMV/Ledoit SI	12.02%	0%	6%	13.27%	0%	2%	13.12%	0%	0%
GMV/Ledoit CC	12.53%	2%	8%	13.66%	2%	0%	13.44%	2%	0%
GMV/w>0	13.49%	0%	0%	13.75%	0%	0%	13.72%	0%	0%
GMV/Ledoit SI/w>0	13.54%	0%	0%	13.79%	0%	0%	13.74%	0%	0%
GMV/Ledoit CC/w>0	13.57%	0%	0%	13.78%	0%	0%	13.72%	0%	0%
MV/B-L/CAPM/w>0	8.50%	0%	8%	8.68%	0%	6%	8.74%	0%	10%
MV/B-L/Ledoit CC/w>0	8.42%	0%	10%	8.03%	0%	10%	7.98%	0%	14%
MV/J-S/CAPM/w>0	10.72%	0%	10%	12.53%	0%	6%	12.64%	0%	2%
MV/J-S/Ledoit CC/w>0	10.42%	0%	12%	12.40%	0%	4%	12.16%	0%	0%

Panel B: NYSE

	T=60			T=120			T=180		
	S-R	$> 1/N$	$< 1/N$	*S-R*	$> 1/N$	$< 1/N$	*S-R*	$> 1/N$	$< 1/N$
$1/N$	13.87%			13.87%			13.87%		
MV/in sample	30.38%	44%	0%	60.91%	54%	0%	161.88%	78%	0%
MV/sample	-2.53%	0%	62%	-0.93%	0%	48%	-0.64%	0%	66%
MV/combination	11.23%	0%	2%	12.62%	0%	2%	12.77%	0%	0%
MV/B-L	7.98%	0%	22%	8.06%	0%	20%	8.05%	0%	20%
MV/B-L (inv. views)	5.19%	0%	34%	7.79%	0%	22%	8.14%	0%	16%
MV/J-S	-1.19%	0%	54%	5.89%	0%	28%	7.98%	0%	22%
MV/CAPM	-1.26%	0%	48%	8.37%	0%	0%	12.83%	0%	0%
MV/CAPM/Pastor (5%)	-1.34%	0%	48%	6.84%	0%	10%	12.52%	0%	0%
MV/FF	1.28%	0%	38%	6.01%	0%	20%	12.81%	0%	0%
MV/FF/Pastor (5%)	0.50%	0%	42%	4.55%	0%	24%	11.52%	0%	2%
MV/PCA	-1.51%	0%	50%	4.14%	0%	36%	-0.19%	0%	48%
MV/sample/w>0	8.59%	0%	14%	8.02%	0%	26%	8.19%	0%	28%
GMV	9.80%	0%	4%	10.64%	0%	6%	10.94%	0%	6%
GMV/Ledoit SI	9.74%	0%	6%	10.46%	0%	12%	10.79%	0%	8%
GMV/Ledoit CC	9.93%	0%	10%	10.63%	0%	10%	10.87%	0%	8%
GMV/w>0	13.75%	0%	0%	13.85%	0%	0%	13.87%	0%	0%
GMV/Ledoit SI/w>0	13.75%	0%	0%	13.85%	0%	0%	13.86%	0%	0%
GMV/Ledoit CC/w>0	13.81%	0%	0%	13.87%	0%	0%	13.86%	0%	0%
MV/B-L/CAPM/w>0	8.79%	0%	24%	8.75%	0%	24%	8.72%	0%	18%
MV/B-L/Ledoit CC/w>0	8.81%	0%	16%	8.52%	0%	18%	8.34%	0%	18%
MV/J-S/CAPM/w>0	9.81%	0%	12%	10.53%	0%	14%	10.83%	0%	12%
MV/J-S/Ledoit CC/w>0	9.47%	0%	12%	10.15%	0%	14%	10.40%	0%	18%

This table reports the mean, over the 50 datasets, of the out-of-sample Sharpe-ratios of the tangency/GMV portfolios created using the estimated moments of the different methods. The columns '$> 1/N$' and '$< 1/N$' show how often, out of the 50 datasets, the corresponding method outperformed the $1/N$ rule and was outperformed by $1/N$ rule, respectively. To test if the Sharpe-ratios are significantly different from each other, the 'two-sample statistic for comparing Sharpe-Ratios' from Opdyke (2007) is used. Panel A reports the results for the stocks from the NASDAQ, and Panel B reports the results for the stocks from the NYSE.

Table 8: The average of the mean and volatility of the portfolios

Panel A: NASDAQ

	T=60		T=120		T=180	
	μ	σ	μ	σ	μ	σ
1/N	1,00%	7,67%	1,00%	7,67%	1,00%	7,67%
MV/in sample	42,18%	423,62%	8,01%	43,96%	4,55%	20,54%
MV/sample	-63,29%	1156,73%	-1,28%	199,41%	-0,86%	44,69%
MV/combination	1,13%	12,92%	0,87%	7,65%	0,91%	7,44%
MV/B-L	0,61%	8,31%	0,61%	8,30%	0,61%	8,29%
MV/B-L (inv. views)	-0,38%	25,11%	0,60%	8,40%	0,62%	8,30%
MV/J-S	0,42%	64,64%	1,51%	22,00%	0,80%	8,76%
MV/CAPM	0,09%	12,01%	-0,12%	14,54%	0,94%	7,68%
MV/CAPM/Pastor	0,56%	63,91%	1,02%	32,47%	0,90%	7,62%
MV/FF	0,85%	40,50%	0,87%	8,31%	0,98%	7,33%
MV/FF/Pastor (5%)	1,24%	56,10%	0,76%	10,51%	0,97%	7,54%
MV/PCA	-1,61%	100,81%	1,73%	45,19%	-70,36%	971,18%
MV/sample/w>0	0,60%	8,37%	0,74%	7,73%	0,67%	7,48%
GMV	0,72%	6,44%	0,80%	6,41%	0,81%	6,43%
GMV/Ledoit SI	0,72%	6,27%	0,81%	6,38%	0,81%	6,43%
GMV/Ledoit CC	0,74%	6,16%	0,83%	6,33%	0,83%	6,39%
GMV/w>0	0,94%	7,04%	0,97%	7,06%	0,97%	7,08%
GMV/Ledoit SI/w>0	0,94%	7,02%	0,97%	7,05%	0,97%	7,08%
GMV/Ledoit CC/w>0	0,95%	7,06%	0,97%	7,09%	0,97%	7,12%
MV/B-L/CAPM/w>0	0,65%	7,74%	0,67%	7,77%	0,68%	7,82%
MV/B-L/Ledoit CC/w>0	0,67%	8,09%	0,65%	8,20%	0,65%	8,23%
MV/J-S/CAPM/w>0	0,72%	6,75%	0,79%	6,46%	0,81%	6,54%
MV/J-S/Ledoit CC/w>0	0,71%	6,82%	0,78%	6,45%	0,78%	6,53%

Panel B: NYSE

	T=60		T=120		T=180	
	μ	σ	μ	σ	μ	σ
1/N	0,82%	5,96%	0,82%	5,96%	0,82%	5,96%
MV/in sample	39,53%	509,83%	7,81%	46,37%	5,67%	47,29%
MV/sample	-53,38%	1083,75%	-2,79%	137,76%	2,93%	148,29%
MV/combination	0,72%	7,10%	0,72%	5,73%	0,73%	5,71%
MV/B-L	0,46%	5,91%	0,47%	5,91%	0,47%	5,92%
MV/B-L (inv. views)	0,50%	12,66%	0,58%	8,16%	0,48%	5,91%
MV/J-S	0,04%	70,33%	0,87%	24,30%	0,35%	10,20%
MV/CAPM	-0,08%	5,63%	0,47%	5,70%	0,72%	5,70%
MV/CAPM/Pastor	-0,32%	13,52%	0,32%	13,77%	0,69%	5,57%
MV/FF	1,44%	72,74%	2,05%	42,77%	0,84%	6,55%
MV/FF/Pastor (5%)	-1,41%	166,03%	2,20%	72,92%	0,90%	8,66%
MV/sample/w>0	0,58%	6,71%	0,49%	6,11%	0,48%	5,87%
GMV	0,46%	4,70%	0,49%	4,67%	0,51%	4,70%
GMV/Ledoit SI	0,44%	4,55%	0,48%	4,63%	0,50%	4,69%
GMV/Ledoit CC	0,44%	4,48%	0,48%	4,61%	0,50%	4,70%
GMV/w>0	0,76%	5,56%	0,78%	5,63%	0,78%	5,67%
GMV/Ledoit SI/w>0	0,76%	5,55%	0,77%	5,62%	0,78%	5,66%
GMV/Ledoit CC/w>0	0,77%	5,61%	0,78%	5,67%	0,78%	5,70%
MV/B-L/CAPM/w>0	0,50%	5,80%	0,51%	5,85%	0,51%	5,88%
MV/B-L/Ledoit CC/w>0	0,52%	6,02%	0,51%	6,00%	0,49%	5,99%
MV/J-S/CAPM/w>0	0,51%	5,26%	0,50%	4,80%	0,51%	4,77%
MV/J-S/Ledoit CC/w>0	0,50%	5,34%	0,48%	4,81%	0,49%	4,78%

This table reports the average, over the 50 datasets, of the out-of-sample expected return and volatility of the tangency/GMV portfolios created using the estimated moments of the different methods. Panel A reports the results for the stocks from the NASDAQ, and Panel B reports the results for the stocks from the NYSE.

than the stocks traded on the New York stock exchange, which could be a reason for the this difference in performance. This will be researched later on.

Another thing that stands out from Table 7 is that the unconstrained mean-variance portfolios that uses the sample moments or PCA implied moments performs very bad across all estimation periods and both datasets. Out of the 50 datasets, they never outperform the $1/N$ portfolio and on average have the lowest Sharpe-ratios of all the methods. One thing to note is that the results of these portfolios are dominated by extreme portfolio returns, which arise from extreme portfolio weights. This problem exists for most unconstrained tangency portfolios for the lowest estimation period of 60 months, where the results are dominated by extreme returns that arise in the more volatile periods. The Sharpe-ratios of the sample based tangency portfolios increase when the estimation period increases, which is due to both an increase in the mean and a decrease in the volatility of the portfolio returns, however it still does not come near the Sharpe-ratios of the $1/N$ portfolio. The sample based tangency portfolio that uses short-sale constraints was the best performing method in the simulation experiment in terms of in-sample Sharpe-ratios. This method performs much better than the unconstrained sample based portfolio weights, but is still outperformed by the $1/N$ portfolio in terms of out-of-sample Sharpe-ratio. Also the combination rule that shrinks the Markowitz portfolio to the $1/N$ rule performs very good compared to other tangency portfolios even for the small estimation period of 60 months. It is the best performing tangency portfolio and even outperforms the short-sale constrained tangency portfolios. However it does not outperform the $1/N$ portfolio in any dataset, which can be blamed on the estimation error in the estimated $\hat{\delta}$.

The unconstrained mean-variance portfolios that use factor implied mean and covariance matrix perform bad for the estimation window of 60 months in both datasets. The portfolios that use CAPM implied moments performs good in terms of portfolio volatility, but bad in terms of expected returns. The portfolios that use FF implied moments, perform bad in terms of both the mean and the volatility of the portfolio return. This can be due to the fact that the FF factors has more parameters that need to be estimated and thus needs a larger estimation sample than the CAPM model. The mean-variance portfolios that use factors based on economic theory, CAPM and FF, perform much better when the estimation period increases. They perform similar to the $1/N$ portfolio for an estimation period of 180 months, because they never have a Sharpe-ratio that is significantly different from the Sharpe-ratio of the $1/N$ portfolio out of the 50 datasets for the NASDAQ dataset. For the NYSE dataset, the portfolio that uses FF implied factors in combination with the prior from Pástor (2000) is outperformed in 1 of the 50 datasets.

The unconstrained mean-variance portfolios that use the shrinkage estimator of James and Stein (1961) for the mean in combination with the sample covariance matrix performs quite bad in all estimation periods for the NYSE dataset, in terms of out-of-sample Sharpe-ratio. This is largely due to the high portfolio volatility, which can be partly blamed on the sample covariance matrix that is used. It performs better for the NASDAQ dataset but it still performs average at best. It performs much better in the constrained tangency portfolio in combination with the CAPM implied covariance matrix and the shrinkage estimator from Ledoit and Wolf (2003a). These portfolios belong to the best performing mean-variance portfolio for the estimation periods of 60 and 120 months. Hence, they do a very good job at reducing the estimation error in the moments.

It stands out that the tangency portfolios that use the Black-Litterman model to estimate the mean performs quite constant across the estimation periods and it is the second best performing unconstrained mean-variance portfolio for the lowest estimation period of 60 months. The out-of-sample Sharpe-ratio does increase slightly as the estimation period increases. This could suggest that for lower estimation periods, where the estimation error is high, the posterior mean of the Black-Litterman model is closer to the implied equilibrium weights. The B-L portfolio performs worse when the EGARCH(1,1) investor views are used for low estimation periods of 60 months. For the estimation period of 120 months it performs much better in terms of reducing the portfolio variance for both datasets, which results in a higher Sharpe-ratio. And for $T = 180$ it performs slightly better for both datasets. This suggests that only for large estimation periods, where the parameters from the EGARCH-M(1,1) model can be estimated accurately, these investor views add extra value.

The unconstrained global minimum variance portfolios perform much better compared to the tangency portfolios in terms of out-of-sample Sharpe-ratios. The higher Sharpe-ratio can be mostly credited to lower portfolio variance, which is the aim of the GMV portfolio. Even for the relative small estimation period of 60 months, the unconstrained GMV portfolios are outperformed rarely by the $1/N$ portfolio. For higher estimation periods, the unconstrained GMV portfolios perform better and perform similar to the $1/N$ portfolio. For the NASDAQ dataset it is never outperformed by the $1/N$ portfolio for the estimation period of 180 months. The unconstrained GMV that uses the shrinkage estimator from Ledoit and Wolf (2003a) with the constant correlation model as a shrinkage target even outperforms the $1/N$ portfolio across all estimation periods in 1 data(sub)set.

The GMV portfolios with short-sale constraints perform even better than the unconstrained GMV portfolios. They are never outperformed by the $1/N$ portfolio across all estimation periods and datasets. This is mostly due to an increase in the expected excess return of the portfolios compared to the unconstrained portfolios. This is due to the fact that short-sale constraints alleviate the the extreme portfolio weights, which lead to extreme portfolio returns. For the NASDAQ dataset, they have an average Sharpe-ratio that is higher than that of the $1/N$ portfolio across all estimation periods. For the NYSE dataset, it still has a Sharpe-ratio that is lower or as big as that of the $1/N$ portfolio.

Overall the portfolios that ignore the estimate of the expected excess return perform the best in terms of out-of-sample Sharpe-ratio. The portfolios that do use the expected returns as an input perform clearly worse than the $1/N$ portfolio and the GMV portfolios. This suggests that portfolio weights are much more sensitive to estimation error in the expected returns than that of the covariance matrix. This corresponds with the findings from Chopra and Ziemba (2011), who state that the effect of the estimation error in the mean on the portfolio weights is about 10 times larger than that of the covariances and variances combined. Therefore, it is better to not use the expected returns as an input, but focus on reducing the estimation error in the covariance matrix and combining it with short-sale constraints.

Table 9: Relative turnover

	NYSE			NASDAQ		
	T=60	T=120	T=180	T=60	T=120	T=180
$1/N$ absolute turnover	0.08	0.08	0.08	0.11	0.11	0.11
Relative turnover for each strategy						
MV/sample	745.16	145.41	146.61	313.93	76.18	34.58
MV/combination	3.27	1.38	1.20	2.81	1.28	1.21
MV/B-L	1.07	1.04	1.03	1.04	1.02	1.02
MV/B-L (inv. views)	10.68	2.23	1.44	10.55	1.45	1.20
MV/J-S	43.17	8.94	4.64	29.48	9.90	1.76
MV/CAPM	2.31	1.45	0.95	3.97	4.85	0.99
MV/CAPM/Pastor (5%)	26.95	13.84	1.01	30.13	10.58	1.25
MV/FF	77.89	12.32	1.84	18.30	2.09	1.15
MV/FF/Pastor (5%)	118.73	40.56	3.22	21.52	3.27	1.33
MV/PCA	95.61	32.01	136.88	36.51	12.98	161.41
MV/sample/w>0	3.20	2.24	1.92	2.38	1.69	1.44
GMV	1.85	1.17	1.00	1.56	1.03	0.89
GMV/Ledoit SI	1.50	1.09	0.96	1.29	0.97	0.87
GMV/Ledoit CC	1.22	0.96	0.89	1.10	0.88	0.82
GMV/w>0	0.98	0.97	0.96	0.96	0.94	0.94
GMV/Ledoit SI/w>0	0.98	0.96	0.96	0.95	0.94	0.94
GMV/Ledoit CC/w>0	0.97	0.97	0.97	0.95	0.94	0.94
MV/B-L/CAPM/w>0	1.38	1.17	1.12	1.21	1.08	1.05
MV/B-L/Ledoit CC/w>0	1.34	1.13	1.09	1.19	1.08	1.07
MV/J-S/CAPM/w>0	2.28	1.44	1.18	1.68	1.09	0.95
MV/J-S/Ledoit CC/w>0	2.38	1.50	1.22	1.76	1.15	1.00

This table shows the turnover for the strategies from each optimizing method relative to the turnover of the $1/N$ portfolio. The first row shows the absolute turnover of the $1/N$ portfolio.

The first row of Table 9 reports the absolute turnover (equation 45) of the $1/N$ portfolio. For the other portfolios, the turnover relative to the $1/N$ portfolio is shown. For the estimation periods of 60 and 120 months, most of the mean-variance portfolios (constrained and unconstrained) have turnover that is considerably higher than that of the $1/N$ portfolio. The mean-variance portfolios that use the Black-Litterman posterior mean as the input for the mean, is the only mean-variance portfolio that has a turnover that is almost as good as the benchmark for these estimation periods. This means that the Black-Litterman posterior mean performs well in terms of resulting in non-extreme and steady portfolio weights. Short-sale constraints naturally result in more stable weights and thus a lower turnover.

For the estimation period of 180 months, most of the mean-variance portfolios are a lot closer to the benchmark in terms of turnover. The portfolio that uses CAPM implied moments even has a lower turnover than the benchmark for both stock exchanges. The sample moments and the principal component implied moments still perform poorly. This was also found in the simulation experiment, where they were both very sensitive to small changes in the inputs. A reason that the principal components perform bad compared to the other factor models, such as CAPM and FF, is that the principal components do not add any economic information in contrast to the other models.

The GMV portfolios perform much better in computing portfolio weights that are stable, as they result in a much lower turnover than most mean-variance portfolios for all estimation periods.

The GMV portfolios with short-sale constraints have a lower turnover compared to the $1/N$ portfolio. The unconstrained GMV portfolios also perform better than the $1/N$ portfolio for higher estimation periods. The unconstrained GMV portfolio that uses the shrinkage estimator from Ledoit and Wolf (2003a) performs the best in terms of turnover for the estimation periods of 120 and 180 months. This is noteworthy, because it even has a lower turnover than the GMV portfolios with short-sale constraints. This means that the unconstrained GMV portfolio that uses the shrinkage estimator from Ledoit and Wolf (2003a) reduces the estimation error in the covariance matrices. However, for the estimation period of 60 months, the GMV portfolios with short-sale constraints have the lowest turnover.

The tables 11, 12 and 13 show how the portfolios from Table 6 perform, in terms of Sharpe-ratios, in different states of the market. These states are defined by 5 quantiles of volatility that range from low to high. The first thing that stands out from these tables is that there is a negative relationship between the portfolio return and the volatility. This is due to the fact that high (low) volatile periods usually go together with dropping (rising) prices, which is called a bearish (bullish) state of the economy.

When the results in Table 7 and tables 11, 12 and 13 are compared, it is clear that the methods that outperform the $1/N$ rule, like the GMV portfolios with short-sale constraints, mostly outperform the $1/N$ rule in periods with lower volatility for both datasets. Even for the estimation period of 60 months, they outperform the $1/N$ portfolio in terms of out-of-sample Sharpe-ratio in these low volatile periods, where they are overall outperformed by the $1/N$ rule. A reason for this could be that in the volatile periods, the estimated covariance matrices contain more extreme estimates. This results in corner solutions for the portfolio weights in contrast to the $1/N$ rule, which avoids concentrated positions.

The GMV portfolios without short-sale constraints do generally have a higher Sharpe-ratio than the $1/N$ rule for higher estimation periods of 120 and 180 months in periods with the highest market volatility. This can be due to the property of the $1/N$ rule that it always takes a long position in any of the available assets. This is not desirable when the market is in a bearish state, because the prices are dropping. The unconstrained GMV portfolios without short-sale constraints can take short positions and therefore outperform the $1/N$ rule in a bear market state. It is also clear that the portfolios that use short-sale constraints perform worse than the unconstrained portfolios in the periods with the highest volatility. This is because of the bear state of the market, where it is more desirable to go short, which is impossible with short-sale constraints. Another notable result is that some of the tangency portfolios without short-sale constraints, like the one that uses CAPM or FF implied moments, usually have a higher average Sharpe-ratio than the $1/N$ and GMV portfolios for the highest market volatility quantile. This can be credited to the fact that the expected returns implied by the CAPM and FF factors contain important information, especially in times of recession, that is ignored in the $1/N$ and GMV portfolios. And thus in high volatile periods it is not a matter of course to ignore the expected returns as an input. Because the NASDAQ stock exchange suffers much more from the burst of the 'dot-com' bubble, the tangency portfolios that use CAPM and FF implied portfolios perform on average better for the NASDAQ dataset relative to the NYSE dataset.

5 Conclusion

The goal of this paper was to find the methods that performed the best in reducing the estimation error in the mean and covariance matrix of stock returns. These estimators were used to compute portfolio weights and were tested in terms of the out-of-sample Sharpe-ratio. The difference with previous research, such as DeMiguel et al. (2009); Tu and Zhou (2011), was that the portfolios were constructed using individual stocks instead of diversified portfolios.

It was found that mean-variance optimized portfolios that use the estimator of the expected returns as an input, perform badly compared to minimum-variance portfolios, which ignores the data of the expected returns. This confirms the existing evidence in the literature, such as Jobson and Korkie (1980), that the sample moments result in extreme and unstable portfolio weights. Minimum-variance portfolios result in more stable weights, because the estimation error in the mean results in extreme and volatile weights even when the estimation error is reduced by the various methods. This concurs with Chopra and Ziemba (2011) who state that the estimation error in the mean has an effect on the portfolio weights which is 10 times larger than that of the covariances and variances combined. In high volatile periods however, the tangency portfolios that use CAPM and FF implied moments do result in an average higher out-of-sample Sharpe-ratio than both the benchmark and the GMV portfolios. Therefore, in periods of high volatility, the information in the mean should not be neglected. Portfolios with short-sale constraints in general perform much better than unconstrained portfolios, because they alleviate the extreme portfolio weights, which lead to extreme portfolio returns. The short-sale constrained portfolio especially outperform the unconstrained portfolios in periods with low volatility, which correspond with periods of expansion. In periods of recession, short-sale constraints hurt the performance of the portfolios more, because it is more desirable to go short.

The $1/N$ portfolio still sits among the top strategies. However, in contrast to DeMiguel et al. (2009), it does not outperform, or is outperformed by, the minimum-variance portfolio with short-sale constraints across all estimation periods and datasets. This can be credited to the higher idiosyncratic risk in the individual stocks relative to the diversified equity portfolios used by DeMiguel et al. (2009). Also, for the higher estimation period of 180 months most of the GMV portfolios and some of the tangency portfolios, like the ones that use CAPM or FF implied moments, are never outperformed significantly by the $1/N$ rule. The higher idiosyncratic risk benefits the optimized portfolio weights, which diversifies the idiosyncratic risk, compared to the equally weighted weights. Also the minimum-variance portfolios perform better in terms of turnover, relative to the benchmark strategy.

Future research: This paper used a lot of methods to reduce the estimation error in the estimated moments. However, there are more methods to reduce the estimation error that were not used in this paper. Future research can therefore exist out of evaluating newer methods like using a higher frequency of data, like daily or weekly, to estimate the covariance matrix more precise. Also evaluate the performance conditional on the amount of stocks in the portfolio to see if the increase in estimation error, due to an increase in the number of parameters, outweighs the increasing ability to diversify.

6 Appendix

First, a consistent estimator for π is:

$$\hat{\pi} = \sum_{i=1}^{N}\sum_{j=1}^{N}\hat{\pi}_{ij} \text{ with}$$

$$\hat{\pi}_{ij} = \frac{1}{T}\sum_{t=1}^{T}\left((y_{it}-\bar{y}_i)(y_{jt}-\bar{y}_j)-s_{ij}\right)^2$$

This result is proven by Ledoit and Wolf (2003b). And the consistent estimator for γ is:

$$\hat{\gamma} = \sum_{i=1}^{N}\sum_{j=1}^{N}\hat{\gamma}_{ij}$$

$$\hat{\gamma}_{ij} = \sum_{i=1}^{N}\sum_{j=1}^{N}(f_{ij}-s_{ij})^2,$$

where f_{ij} and s_{ij} are consistent estimators ϕ_{ij} and σ_{ij}, respectively.

The consistent estimator for ρ depends on the shrinkage target. When the single index model, used by Ledoit and Wolf (2003b), is the shrinkage target, the consistent estimator for ρ is as follows:

$$\hat{\rho} = \sum_{i=1}^{N}\sum_{j=1}^{N}\hat{\rho}_{ij}$$

$$\hat{\rho}_{ij} = \sum_{i=1}^{N}\hat{\pi}_{ii} + \sum_{i=1}^{N}\sum_{j=1,j\neq i}^{N} r_{ijt}$$

$$\hat{\pi}_{ii} = \frac{1}{T}\sum_{t=1}^{T}\left((y_{it}-\bar{y}_i)^2-s_{ij}\right)^2$$

$$r_{ijt} = \frac{s_{j0}s_{00}(y_{it}-\bar{y}_i)+s_{i0}s_{00}(y_{jt}-\bar{y}_j)-s_{i0}s_{j0}(y_{0t}-\bar{y}_0)}{s_{00}^2}(y_{0t}-\bar{y}_0)(y_{it}-\bar{y}_i)(y_{jt}-\bar{y}_j)-f_{ij}s_{ij}$$

,where y_{0t} is the market return at time t and s_{i0} is the sample covariance between stock i's return and the market return.

When the constant correlation model, used by Ledoit and Wolf (2003a), is the shrinkage target, the consistent estimator for ρ is as follows:

$$\hat{\rho} = \sum_{i=1}^{N}\sum_{j=1}^{N}\hat{\rho}_{ij}$$

$$\hat{\rho}_{ij} = \sum_{i=1}^{N}\hat{\pi}_{ii} + \sum_{i=1}^{N}\sum_{j=1,j\neq i}^{N}\frac{\bar{r}}{2}\left(\sqrt{\frac{s_{jj}}{s_{ii}}}\hat{\upsilon}_{ii,ij} + \sqrt{\frac{s_{ii}}{s_{jj}}}\hat{\upsilon}_{jj,ij}\right)$$

$$\hat{\pi}_{ii} = \frac{1}{T}\sum_{t=1}^{T}\big((y_{it}-\bar{y}_i)^2 - s_{ij}\big)^2$$

$$\bar{r} = \frac{2}{(N-1)N}\sum_{i=1}^{N-1}\sum_{j=i+1}^{N} r_{ij}$$

$$r_{ij} = \frac{s_{ij}}{\sqrt{s_{ii}sjj}}$$

$$\hat{\upsilon}_{ii,ij} = \frac{1}{T}\sum_{t=1}^{T}\big((y_{it}-\bar{y}_i)^2 - s_{ii}\big)\big((y_{it}-\bar{y}_i)(y_{jt}-\bar{y}_j) - s_{ij}\big)$$

$$\hat{\upsilon}_{jj,ij} = \frac{1}{T}\sum_{t=1}^{T}\big((y_{jt}-\bar{y}_j)^2 - s_{jj}\big)\big((y_{it}-\bar{y}_i)(y_{jt}-\bar{y}_j) - s_{ij}\big),$$

where $\hat{\upsilon}_{ii,ij}$ and $\hat{\upsilon}_{jj,ij}$ are consistent estimators for $AsyCov\big(\sqrt{T}s_{ii}, \sqrt{T}s_{ij}\big)$ and $AsyCov\big(\sqrt{T}s_{jj}, \sqrt{T}s_{ij}\big)$, respectively.

Table 10: Averages of δ

	NYSE			NASDAQ		
	T=60	T=120	T=180	T=60	T=120	T=180
$\hat{\delta}$	0.18 (0.06)	0.17(0.06)	0.16(0.07)	0.30(0.11)	0.25(0.10)	0.21(0.10)

This table contains the averages of the estimated parameters δ from section 2.7 for each dataset (NYSE and NASDAQ) and estimation period (T). Behind the averages, the standard deviation of the estimated parameters over time is shown between brackets.

Table 11: Sharpe-ratios in different market volatility periods

T=60:

	NYSE					NASDAQ				
volatility quantile	1	2	3	4	5	1	2	3	4	5
1/N	38.51%	31.15%	37.64%	11.79%	-9.03%	52.88%	21.32%	9.92%	19.13%	-8.78%
MV/sample	-0.68%	-2.01%	0.63%	-4.71%	-4.93%	2.81%	-3.86%	-0.62%	-6.18%	-8.97%
MV/combination	31.06%	26.22%	32.38%	11.34%	-11.07%	50.90%	17.06%	8.25%	15.88%	-9.81%
MV/B-L	27.65%	25.77%	22.55%	4.14%	-13.20%	37.71%	14.66%	3.47%	12.51%	-9.58%
MV/B-L (inv. views)	20.98%	21.39%	15.18%	1.99%	-11.20%	26.53%	8.92%	4.16%	9.21%	-5.67%
MV/J-S	6.24%	7.62%	8.17%	1.61%	-9.23%	28.97%	9.55%	5.59%	0.36%	-11.66%
MV/CAPM	12.56%	18.73%	17.26%	-8.33%	-20.29%	42.56%	-0.63%	6.97%	-1.27%	-5.12%
MV/CAPM/Pastor (5%)	9.92%	14.16%	11.96%	-4.46%	-18.13%	38.64%	-4.24%	3.21%	-6.69%	-5.38%
MV/FF	18.30%	10.10%	11.18%	4.04%	-8.27%	39.75%	16.83%	4.11%	5.54%	-13.84%
MV/FF/Pastor (5%)	15.59%	7.83%	7.23%	4.96%	-7.50%	34.33%	13.83%	1.32%	4.54%	-11.88%
MV/PCA	3.08%	3.85%	5.89%	-3.06%	-5.32%	23.11%	5.76%	-0.66%	-4.16%	-5.88%
MV/sample/w>0	22.26%	22.39%	24.05%	4.82%	-11.30%	40.08%	10.99%	8.21%	11.80%	-13.53%
GMV	27.76%	24.31%	22.46%	7.11%	-11.45%	38.27%	18.01%	11.53%	18.37%	-10.21%
GMV/Ledoit SI	27.83%	24.77%	22.15%	6.52%	-11.32%	40.12%	18.57%	12.01%	18.98%	-9.94%
GMV/Ledoit CC	29.64%	24.83%	22.34%	6.39%	-10.94%	41.54%	19.79%	11.77%	20.37%	-9.03%
GMV/w>0	38.60%	31.57%	37.47%	11.57%	-9.28%	53.67%	22.11%	10.68%	19.81%	-9.21%
GMV/Ledoit SI/w>0	38.63%	31.62%	37.50%	11.55%	-9.29%	53.85%	22.13%	10.57%	19.99%	-9.04%
GMV/Ledoit CC/w>0	38.74%	31.74%	37.79%	11.62%	-9.26%	41.86%	17.25%	4.60%	13.38%	-9.79%
MV/B-L/CAPM/w>0	29.28%	26.78%	24.24%	5.25%	-13.03%	39.84%	17.10%	3.44%	14.24%	-8.88%
MV/B-L/Ledoit CC/w>0	28.74%	26.83%	24.46%	5.41%	-12.27%	53.50%	21.95%	10.62%	19.71%	-9.23%
MV/J-S/CAPM/w>0	26.86%	24.75%	26.30%	6.82%	-10.29%	46.29%	18.07%	12.15%	16.93%	-11.93%
MV/J-S/Ledoit CC/w>0	26.74%	24.01%	25.61%	5.81%	-9.68%	45.33%	16.56%	11.49%	17.17%	-11.19%

This table reports the mean, over the 50 datasets, of the out-of-sample Sharpe-ratios of the tangency/GMV portfolios created using the estimated moments of the different methods in different periods of levels of volatility. The monthly volatility quantile borders for the NASDAQ capitalization weighted index are: (0.036 0.044 0.055 0.080). The monthly volatility quantile borders for the NASDAQ capitalization weighted index are: (0.026 0.032 0.041 0.053)

Table 12: Sharpe-ratios in different market volatility periods

T=120:

	NYSE					NASDAQ				
volatility quantile	1	2	3	4	5	1	2	3	4	5
1/N	38.51%	31.15%	37.64%	11.79%	-9.03%	52.88%	21.32%	9.92%	19.13%	-8.78%
MV/sample	1.07%	4.38%	4.74%	0.57%	-6.96%	18.13%	5.33%	4.42%	1.54%	-10.62%
MV/combination	36.14%	29.52%	36.04%	10.32%	-10.77%	52.21%	19.80%	9.58%	17.80%	-9.92%
MV/B-L	27.64%	25.62%	22.88%	4.09%	-13.03%	37.97%	14.73%	3.35%	12.60%	-9.58%
MV/B-L (inv. views)	26.68%	25.00%	21.79%	3.68%	-11.60%	36.35%	14.78%	3.12%	12.41%	-9.46%
MV/J-S	22.41%	18.02%	17.32%	3.97%	-11.71%	40.54%	17.32%	10.93%	15.32%	-9.28%
MV/CAPM	38.35%	31.11%	31.23%	-2.91%	-12.68%	41.70%	15.12%	-6.15%	0.27%	-8.34%
MV/CAPM/Pastor (5%)	37.87%	31.04%	28.22%	-2.83%	-8.90%	37.79%	13.91%	-5.72%	-1.27%	-10.77%
MV/FF	28.40%	26.01%	30.29%	10.35%	-5.85%	45.19%	19.47%	8.21%	19.04%	-10.23%
MV/FF/Pastor (5%)	24.06%	20.86%	24.44%	9.28%	-5.62%	41.88%	17.26%	8.80%	17.65%	-10.18%
MV/PCA	18.86%	16.54%	21.03%	3.28%	-10.66%	42.06%	16.42%	8.00%	10.89%	-11.61%
MV/sample/w>0	22.13%	21.34%	25.91%	3.89%	-11.50%	42.78%	14.25%	8.48%	16.77%	-11.09%
GMV	29.69%	25.74%	23.22%	7.10%	-8.91%	41.94%	18.90%	13.16%	21.37%	-8.38%
GMV/Ledoit SI	29.48%	25.52%	22.42%	6.77%	-8.65%	42.37%	19.48%	13.21%	21.74%	-8.25%
GMV/Ledoit CC	30.57%	25.60%	22.60%	7.12%	-8.08%	44.14%	20.52%	12.68%	22.81%	-7.85%
GMV/w>0	38.74%	31.69%	37.44%	11.51%	-9.08%	54.01%	22.13%	10.88%	20.12%	-8.99%
GMV/Ledoit SI/w>0	38.74%	31.71%	37.44%	11.49%	-9.08%	54.15%	22.23%	10.93%	20.18%	-8.97%
GMV/Ledoit CC/w>0	38.81%	31.77%	37.66%	11.54%	-9.07%	54.25%	22.26%	10.88%	20.27%	-8.93%
MV/B-L/CAPM/w>0	29.00%	26.76%	24.14%	4.95%	-12.58%	41.50%	16.76%	5.53%	13.92%	-9.64%
MV/B-L/Ledoit CC/w>0	28.59%	26.46%	24.13%	4.92%	-12.71%	39.53%	15.87%	4.04%	13.41%	-9.32%
MV/J-S/CAPM/w>0	29.46%	25.14%	25.86%	7.33%	-9.74%	47.23%	20.81%	12.10%	20.55%	-10.07%
MV/J-S/Ledoit CC/w>0	29.72%	24.30%	25.32%	6.54%	-9.57%	46.73%	19.94%	11.89%	20.98%	-9.88%

This table reports the mean, over the 50 datasets, of the out-of-sample Sharpe-ratios of the tangency/GMV portfolios created using the estimated moments of the different methods in different periods of levels of volatility. The monthly volatility quantile borders for the NASDAQ capitalization weighted index are: (0.036 0.044 0.055 0.080). The monthly volatility quantile borders for the NASDAQ capitalization weighted index are: (0.026 0.032 0.041 0.053)

Table 13: Sharpe-ratios in different market volatility periods

T=180:

	NYSE					NASDAQ				
volatility quantile	1	2	3	4	5	1	2	3	4	5
1/N	38.51%	31.15%	37.64%	11.79%	-9.03%	52.88%	21.32%	9.92%	19.13%	-8.78%
MV/sample	5.30%	3.82%	7.89%	-0.45%	-11.74%	29.98%	12.84%	2.79%	6.74%	-8.75%
MV/combination	36.74%	29.43%	36.15%	10.19%	-10.72%	53.95%	20.94%	9.09%	18.24%	-9.44%
MV/B-L	27.71%	25.68%	22.74%	3.99%	-12.96%	38.06%	14.79%	3.38%	12.63%	-9.59%
MV/B-L (inv. views)	27.70%	25.68%	22.78%	4.24%	-12.91%	37.74%	14.89%	3.20%	12.79%	-9.31%
MV/J-S	27.32%	21.66%	19.57%	4.72%	-10.55%	44.49%	20.29%	10.73%	17.55%	-9.65%
MV/CAPM	38.56%	30.42%	37.00%	9.53%	-8.66%	50.51%	20.08%	8.42%	20.15%	-7.37%
MV/CAPM/Pastor (5%)	38.04%	30.29%	36.25%	9.17%	-9.69%	49.89%	20.69%	8.11%	19.88%	-8.34%
MV/FF	29.20%	24.56%	32.26%	11.89%	-6.61%	48.09%	22.23%	8.88%	22.25%	-7.28%
MV/FF/Pastor (5%)	23.98%	20.19%	28.12%	11.68%	-6.65%	45.20%	21.19%	8.49%	21.21%	-7.74%
MV/PCA	2.97%	0.95%	5.43%	-1.42%	-8.11%	27.10%	10.92%	1.55%	4.00%	-8.76%
MV/sample/w>0	23.38%	22.20%	26.93%	3.19%	-12.80%	44.88%	15.68%	5.94%	16.53%	-11.33%
GMV	31.93%	26.70%	22.89%	7.05%	-8.41%	43.39%	20.59%	12.97%	21.07%	-8.88%
GMV/Ledoit SI	31.62%	26.41%	22.35%	6.86%	-8.22%	43.36%	20.76%	12.90%	21.28%	-8.81%
GMV/Ledoit CC	32.32%	26.46%	22.86%	7.10%	-7.83%	44.59%	21.33%	12.79%	22.32%	-8.41%
GMV/w>0	38.88%	31.71%	37.34%	11.48%	-9.02%	54.12%	22.30%	11.07%	20.07%	-9.05%
GMV/Ledoit SI/w>0	38.89%	31.71%	37.34%	11.46%	-9.02%	54.19%	22.37%	11.11%	20.11%	-9.05%
GMV/Ledoit CC/w>0	38.91%	31.68%	37.54%	11.48%	-9.01%	54.26%	22.35%	10.97%	20.18%	-9.02%
MV/B-L/CAPM/w>0	28.97%	26.57%	24.23%	4.90%	-12.40%	42.41%	16.43%	5.33%	13.88%	-9.41%
MV/B-L/Ledoit CC/w>0	28.35%	25.82%	23.87%	4.59%	-12.66%	39.94%	15.29%	3.68%	13.16%	-9.25%
MV/J-S/CAPM/w>0	32.19%	26.13%	25.81%	6.83%	-9.89%	48.64%	21.69%	11.74%	20.67%	-9.69%
MV/J-S/Ledoit CC/w>0	32.51%	25.22%	25.22%	6.07%	-9.90%	47.93%	20.77%	10.94%	20.39%	-9.64%

This table reports the mean, over the 50 datasets, of the out-of-sample Sharpe-ratios of the tangency/GMV portfolios created using the estimated moments of the different methods in different periods of levels of volatility. The monthly volatility quantile borders for the NASDAQ capitalization weighted index are: (0.036 0.044 0.055 0.080). The monthly volatility quantile borders for the NASDAQ capitalization weighted index are: (0.026 0.032 0.041 0.053)

(a) Sample moments, T=25

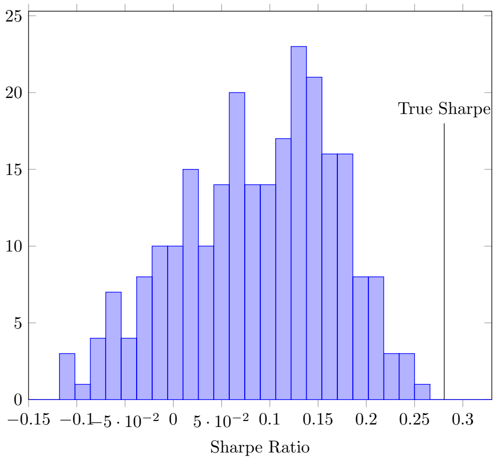

(b) Sample moments constrained w>0, T=25

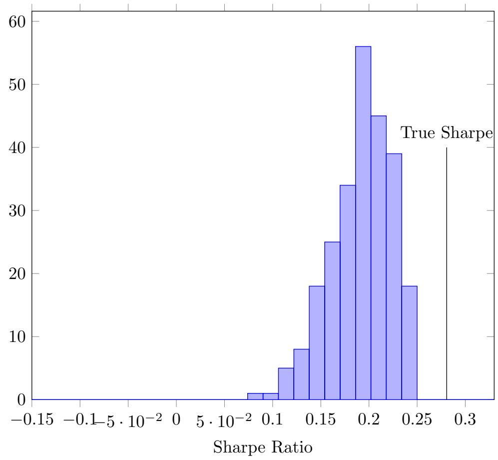

(c) Sample moments, T=75

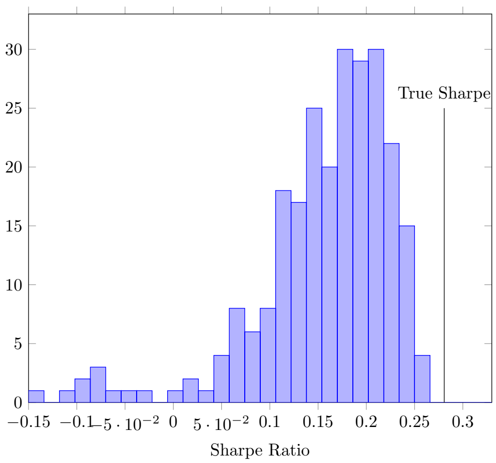

(d) Sample moments constrained w>0, T=75

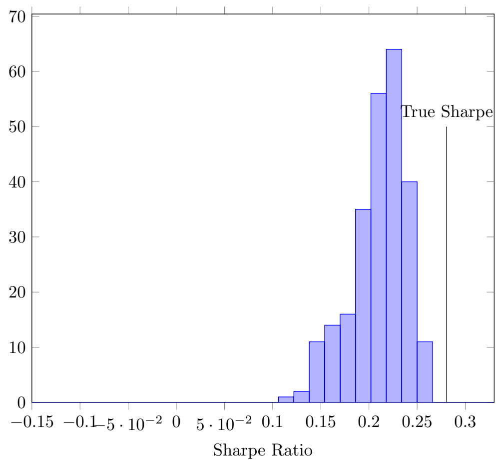

(e) Sample moments, T=150

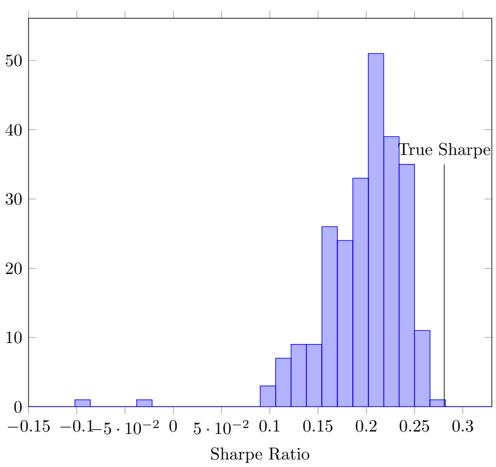

(f) Sample moments constrained w>0, T=150

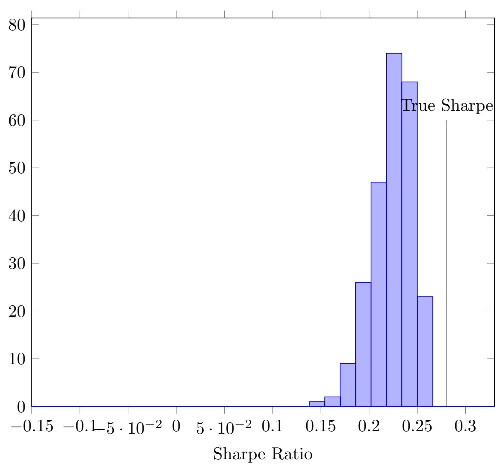

Figure 3: These figures are histograms showing the Sharpe-ratios of 250 simulations for the tangency portfolios using the moments estimated by the different methods. The "True Sharpe" is the Sharpe-ratio that belongs to the tangency portfolio that uses the true moments by which the data sets are simulated.

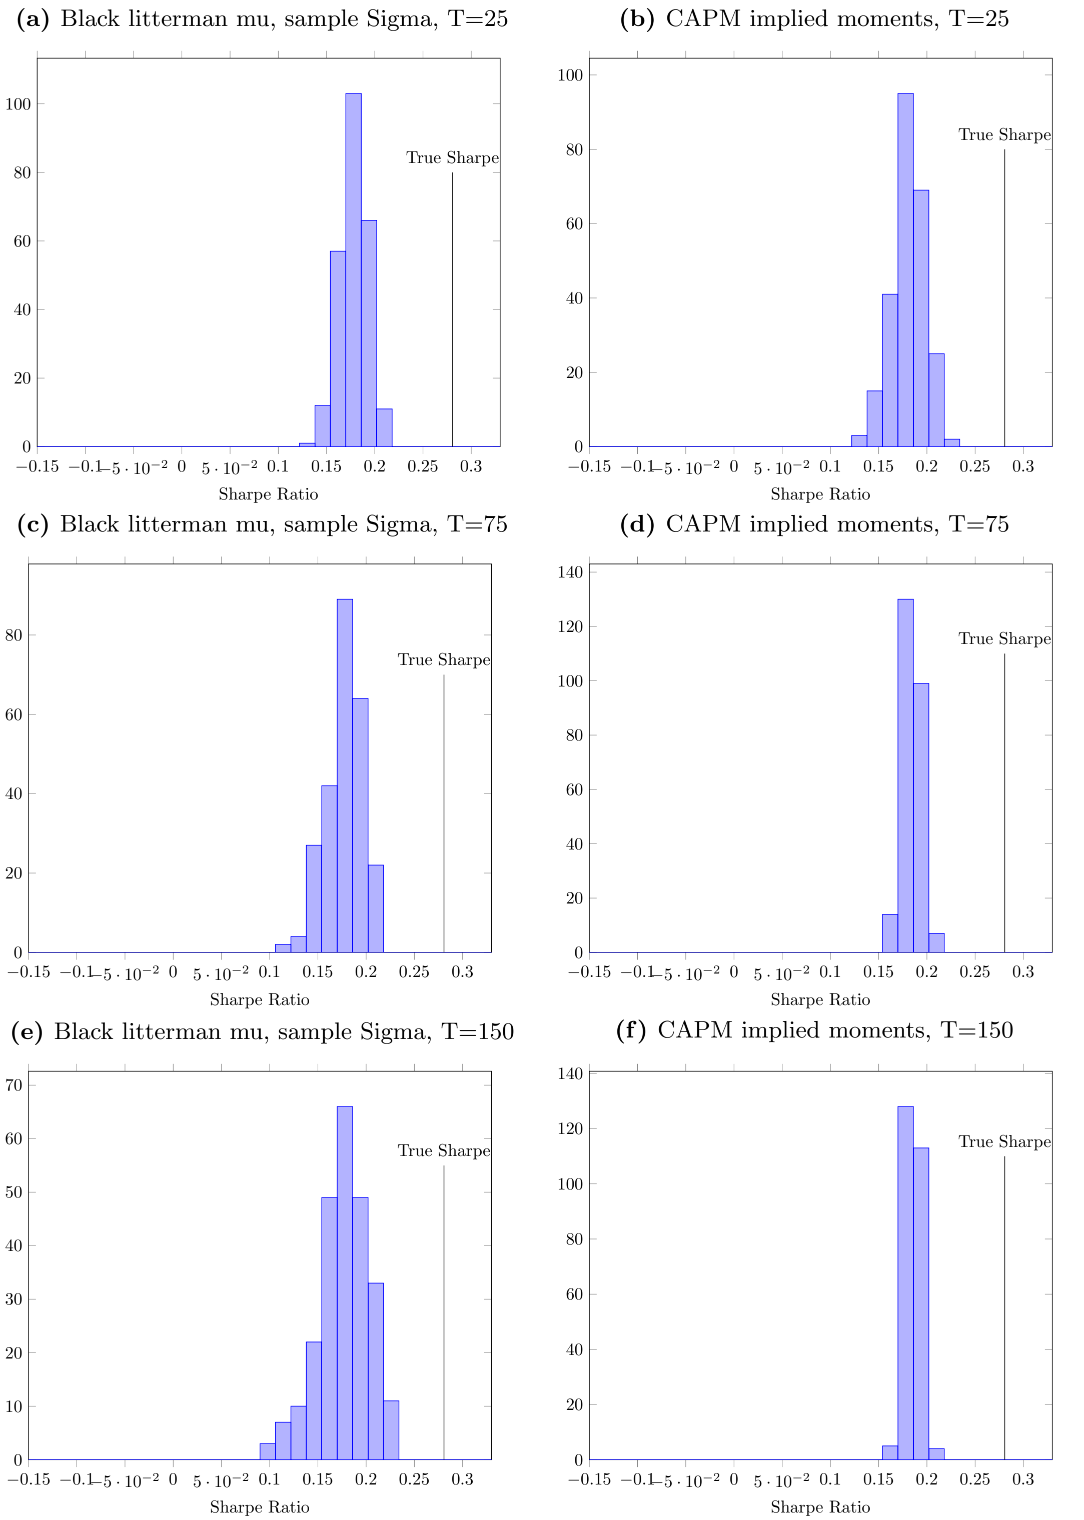

Figure 4: These figures are histograms showing the Sharpe-ratios of 250 simulations for the tangency portfolios using the moments estimated by the different methods. The "True Sharpe" is the Sharpe-ratio that belongs to the tangency portfolio that uses the true moments by which the data sets are simulated.

(a) CAPM pastor 5% implied moments, T=25

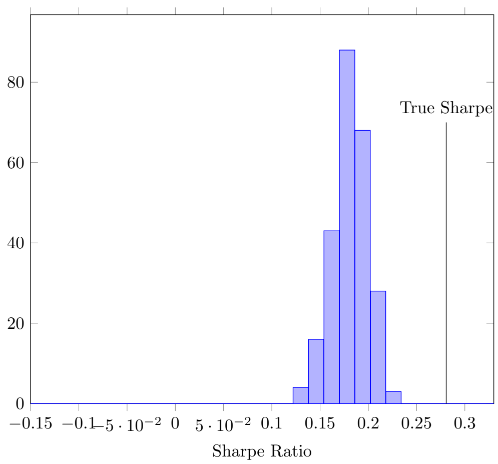

(b) James Stein mean, sample covariance T=25

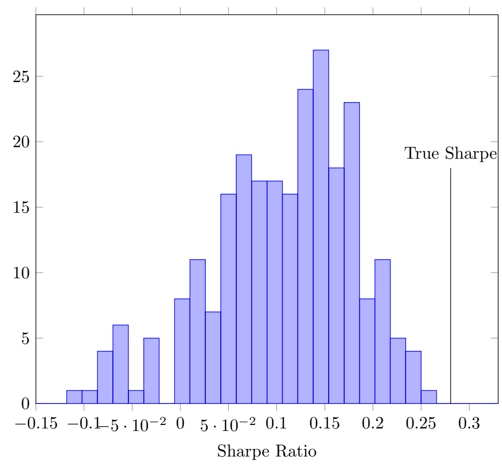

(c) CAPM pastor 5% implied moments, T=75

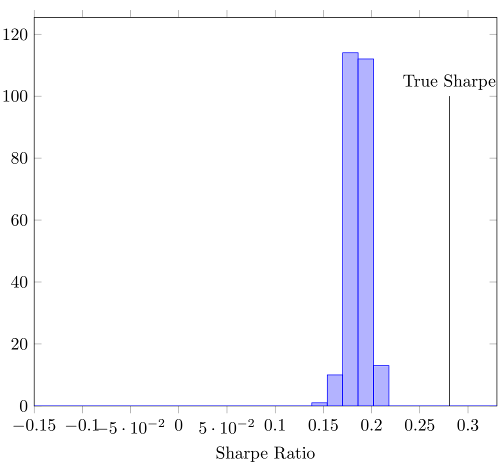

(d) James Stein mean, sample covariance T=75

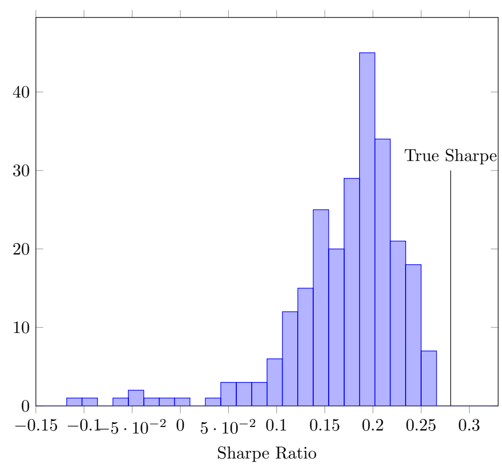

(e) CAPM pastor 5% implied moments, T=150

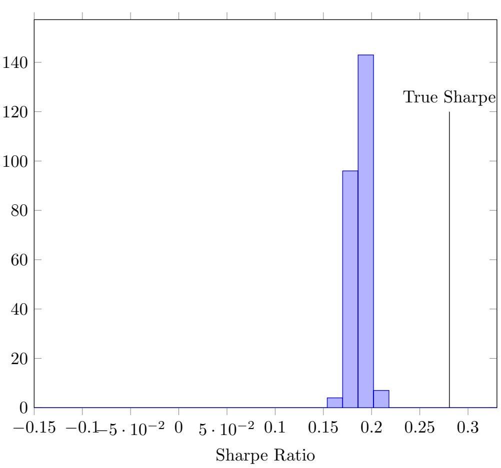

(f) James Stein mean, sample covariance T=150

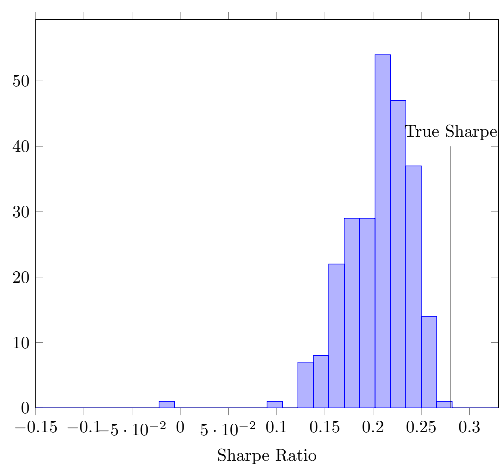

Figure 5: These figures are histograms showing the Sharpe-ratios of 250 simulations for the tangency portfolios using the moments estimated by the different methods. The "True Sharpe" is the Sharpe-ratio that belongs to the tangency portfolio that uses the true moments by which the data sets are simulated.

(a) sample mean, Ledoit 2003 covariance T=25

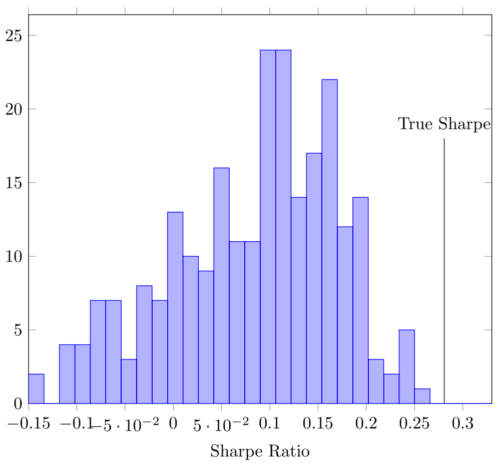

(b) sample mean, Ledoit 2004 covariance T=25

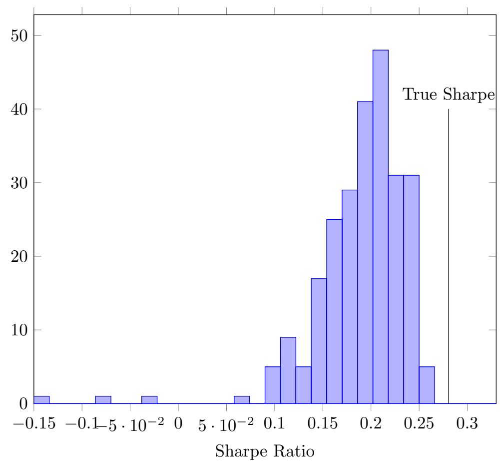

(c) sample mean, Ledoit 2003 covariance T=75

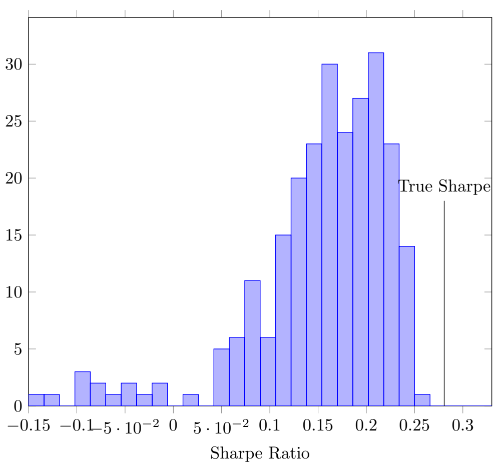

(d) sample mean, Ledoit 2004 covariance T=75

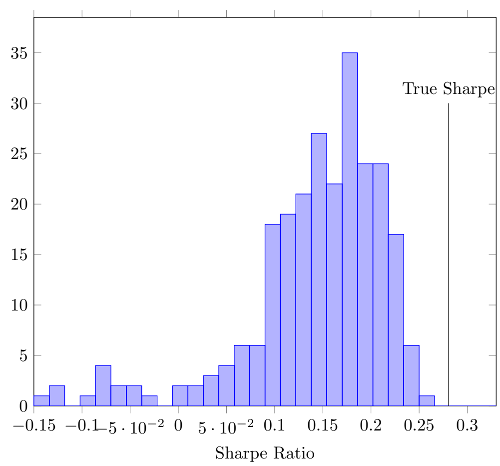

(e) sample mean, Ledoit 2003 covariance T=150

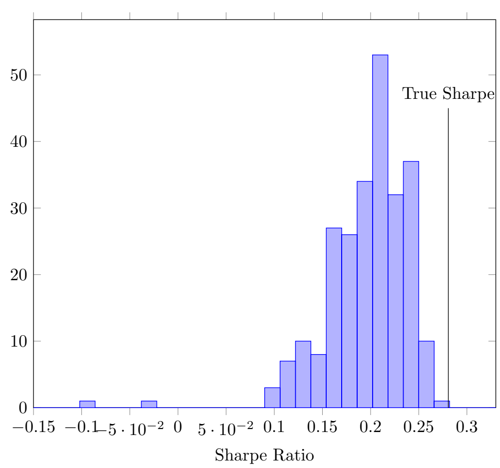

(f) sample mean, Ledoit 2004 covariance T=150

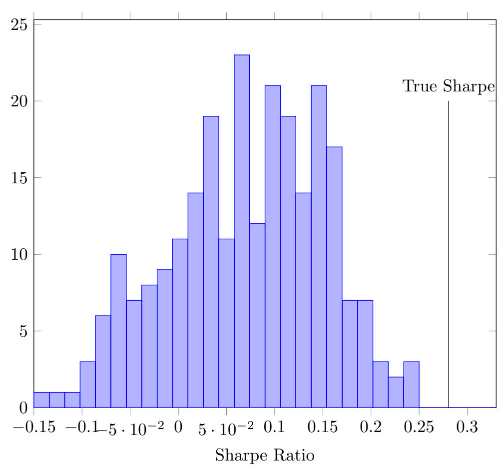

Figure 6: These figures are histograms showing the Sharpe-ratios of 250 simulations for the tangency portfolios using the moments estimated by the different methods. The "True Sharpe" is the Sharpe-ratio that belongs to the tangency portfolio that uses the true moments by which the data sets are simulated.

(a) pca implied moments T=25

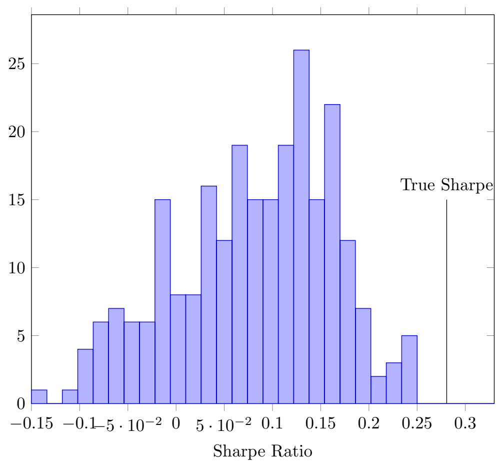

(b) pca implied moments T=75

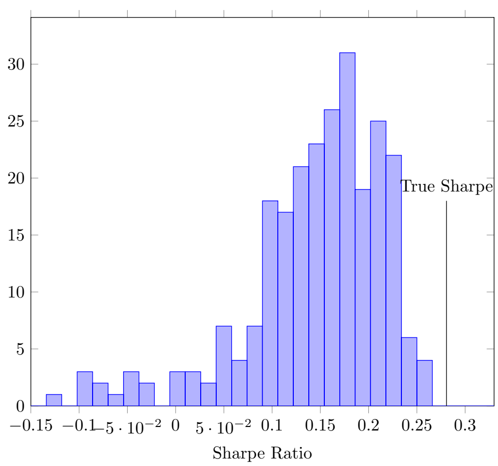

(c) pca implied moments T=150

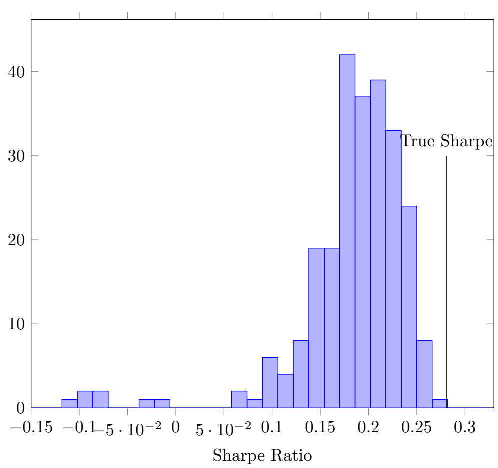

Figure 7: These figures are histograms showing the Sharpe-ratios of 250 simulations for the tangency portfolios using the moments estimated by the different methods. The "True Sharpe" is the Sharpe-ratio that belongs to the tangency portfolio that uses the true moments by which the data sets are simulated.

References

Steven L Beach and Alexei G Orlov. An application of the black–litterman model with egarch-m-derived views for international portfolio management. *Financial Markets and Portfolio Management*, 21(2):147–166, 2007.

Michael J Best and Robert R Grauer. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results. *Review of Financial Studies*, 4(2):315–342, 1991.

Fischer Black and Robert Litterman. Global portfolio optimization. *Financial Analysts Journal*, 48(5):28–43, 1992.

Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics*, 31(3):307–327, 1986.

Vijay K Chopra and William T Ziemba. The effect of errors in means, variances, and covariances on optimal portfolio choice. *The Kelly Capital Growth Investment Criterion: Theory and Practice*, 3:249, 2011.

Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *Review of Financial Studies*, 22(5):1915–1953, 2009.

Robert F Engle, David M Lilien, and Russell P Robins. Estimating time varying risk premia in the term structure: The arch-m model. *Econometrica: Journal of the Econometric Society*, pages 391–407, 1987.

Eugene F Fama and Kenneth R French. The cross-section of expected stock returns. *the Journal of Finance*, 47(2):427–465, 1992.

Eugene F Fama and Kenneth R French. Dissecting anomalies. *The Journal of Finance*, 63(4): 1653–1678, 2008.

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

Richaed C Green and Burton Hollifield. When will mean-variance efficient portfolios be well diversified? *The Journal of Finance*, 47(5):1785–1809, 1992.

Ravi Jagannathan and Tongshu Ma. Risk reduction in large portfolios: Why imposing the wrong constraints helps. *The Journal of Finance*, 58(4):1651–1684, 2003.

William James and Charles Stein. Estimation with quadratic loss. In *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 361–379, 1961.

J David Jobson and Bob Korkie. Estimation for markowitz efficient portfolios. *Journal of the American Statistical Association*, 75(371):544–554, 1980.

Olivier Ledoit and Michael Wolf. Honey, i shrunk the sample covariance matrix. *UPF Economics and Business Working Paper*, 2003a.

Olivier Ledoit and Michael Wolf. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of empirical finance*, 10(5):603–621, 2003b.

Wai Lee. *Theory and Methodology of Tactical Asset Allocation*, volume 65. John Wiley & Sons, 2000.

Harry Markowitz. Portfolio selection. *The journal of finance*, 7(1):77–91, 1952.

Richard O Michaud. The markowitz optimization enigma: is' optimized'optimal? *Financial Analysts Journal*, 45(1):31–42, 1989.

Daniel B Nelson. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica: Journal of the Econometric Society*, pages 347–370, 1991.

John Douglas JD Opdyke. Comparing sharpe ratios: so where are the p-values? *Journal of Asset Management*, 8(5):308–336, 2007.

L'uboš Pástor. Portfolio selection and asset pricing models. *The Journal of Finance*, 55(1): 179–223, 2000.

AJ Patton and K Sheppard. Evaluating volatility and correlation forecasts. handbook of financial time series. edited y tg andersen, ra, davis, j.'p. kreiss and t. mikosch, 2007.

William F Sharpe. A simplified model for portfolio analysis. *Management science*, 9(2):277–293, 1963.

Jun Tu and Guofu Zhou. Markowitz meets talmud: A combination of sophisticated and naive diversification strategies. *Journal of Financial Economics*, 99(1):204–215, 2011.

Zhenyu Wang. A shrinkage approach to model uncertainty and asset allocation. *Review of Financial Studies*, 18(2):673–705, 2005.