Risk–Reward Optimisation for Long-Run Investors: an Empirical Analysis*

Manfred Gilli[a] & Enrico Schumann[b]

[a] Department of Econometrics, University of Geneva and Swiss Finance Institute

[b] Department of Econometrics, University of Geneva

21 August 2009

Abstract

This is a draft version and must not be quoted. Comments are very welcome.

A common approach in portfolio selection is to characterise a portfolio of assets by a desired property, the 'reward', and something undesirable, the risk. These properties are often identified with mean and variance of returns, respectively, even though, given the non-Gaussian nature of financial time series, alternative specifications like partial and conditional moments, quantiles, and drawdowns seem theoretically more appropriate. We analyse the empirical performance of portfolios selected by optimising risk–reward ratios constructed from such alternative functions. We find that in many cases these portfolios outperform our benchmark (minimum-variance), in particular when long-run returns are concerned. We also find, however, that all the strategies tested (including minimum-variance) are sensitive to relatively small changes in the data. The main theme throughout our analysis is that minimising risk, as opposed to maximising reward, leads to good out-of-sample performance. Adding a reward-function to the selection criterion usually improves a given strategy only marginally.

Keywords: Portfolio optimisation, Optimisation heuristics, Downside risk

JEL codes: G11

1 Introduction

In modern portfolio theory (Markowitz, 1952), a portfolio of financial assets is characterised by a desired property, the 'reward', and something undesirable, the 'risk'. Balancing financial aspects against statistical and, particularly, computational limitations, Markowitz identified reward and risk with the expectation and the variance of returns, respectively. An alleged weakness of the resulting mean–variance optimisation is that selecting portfolios

*The authors gratefully acknowledge financial support from the EU Commission through MRTN-CT-2006-034270 COMISEF.

only on the basis of the first two moments of portfolio returns should not be appropriate, given the considerable body of evidence of the non-Gaussian nature of financial time series (Cont, 2001).

This paper investigates this criticism. We empirically evaluate portfolio selection criteria that have been proposed as alternatives to the mean–variance rule, thus we replace the mean and variance by alternative measures of 'reward' and 'risk'. These alternative functions explicitly take into account certain empirical regularities ('stylised facts') of financial prices like fat tails or asymmetric return distributions.

Our paper contributes to the literature by giving evidence of the empirical effectiveness of selection criteria for financial assets portfolios. We also provide robustness checks for our empirical results by trying to capture the uncertainty around the point estimates that are usually presented in empirical studies on portfolio optimisation. To solve the portfolio problems we use a heuristic optimisation technique, Threshold Accepting, which is capable of optimising portfolios under all the different selection criteria discussed. Threshold Accepting works directly on the empirical distribution function of portfolio returns, it does not require approximating the data by a parametric distribution.

The paper links to two strands of the financial literature. Firstly, there is the large number of studies investigating the empirical performance of mean–variance optimisation. The main finding here is that a straightforward estimation of the required parameters, that is the assets' means and their variance–covariance matrix, and their 'plugging-into' the objective function, often lead to badly diversified portfolios that perform poorly out-of-sample. Secondly, there are several theoretical studies on desirable properties of risk and performance measures, in particular such measures that capture non-Gaussian properties of the data. Related is the literature on performance measurement and performance attribution, which often comes from a more practical or accounting background. We briefly review some results of these studies here.

With highly-correlated data series, portfolios obtained from mean–variance optimisation are very sensitive to the required input parameters (means, variances and correlations), with small changes in the estimated parameters leading to large changes in portfolio holdings. Thus, even under a stable data-generating process, the in-sample efficient frontier is very difficult to estimate, which makes it often a poor predictor of the out-of-sample frontier (in which we are naturally more interested). Hence, portfolios that appear mean–variance efficient in one period are often far from efficient and perform poorly in the next period.

These estimation difficulties are by now well documented; early studies are Cohen and Pogue (1967); Frankfurter et al. (1971), see also Jobson and Korkie (1980); Jorion (1985, 1986); Best and Grauer (1991); Chopra et al. (1993); Board and Sutcliffe (1994). Brandt (forthcoming) gives a very good overview. The general result of these studies is that in finite samples parameter estimates are usually very imprecise; the means in particular are most difficult to estimate (though also most rewarding if measured correctly). One extreme reaction to these estimation problems is to completely disregard any historical information, a strategy that has recently received a lot of attention with DeMiguel et al. (2009). Empirically, however, there is evidence of persistent characteristics in financial time series. For instance, when compared with constant predictors, recent past variances seem to be good predictors of future variances, so it seems strange not to exploit this knowledge. An optimal data strategy will thus very likely include at least some historical information, though (just as likely) not all information. This is a well-known principle for forecasts based on noisy data: to reduce overfitting and to obtain better out-of-sample results, sample information needs to be ignored; otherwise, we will fit ('overfit') the model to the noise. Under this principle, seemingly contradictory strategies (like removing noise by applying principal components analysis, or adding additional noise by resampling) achieve the same result.

Different less extreme methods have been proposed in the literature to improve estimation, for instance shrinkage estimation or the usage of factor models (Briner and Connor, 2008; Brandt, forthcoming). Though often motivated differently, the prescription of these techniques is similar: they either constrain the input parameters to be 'more equal', or set up maximum holding sizes. Either approach results in less extreme portfolio weights and more diversification. There are several recurrent findings concerning these different approaches. Not allowing short sales improves performance, as does disregarding the means altogether and only minimising a portfolio's variance (see Board and Sutcliffe (1994) for UK stocks, Chan et al. (1999) for US stocks). In fact, the robust empirical performance of this minimum-variance (MV) portfolio led to considerable interest in improving the estimation of the variance–covariance matrix. Various shrinkage and 'portfolio of estimators' methods seem to improve performance, even though empirically it seems hard to distinguish a truly best method (Chan et al., 1999; Disatnik and Benninga, 2007), with negligible advantages of more elaborate models over simpler ones. Still, this suggests a long-only MV portfolio as a natural benchmark for alternative portfolio construction methods.

If a riskless asset is available, mean–variance optimisation reduces to maximising a portfolio's Sharpe ratio (Sharpe, 1966, 1994), ie, the ratio of a portfolio's excess return over the riskfree rate to its standard deviation. Hence the Sharpe ratio, though probably the most widely used mapping of a portfolio's desirability into a single real number, inherits the alleged weakness of mean–variance optimisation. In recent years, a large number of alternative risk and performance measures have been proposed (for an overview, see for example Rachev et al. (2005)). The development and usage of these measures was driven to a considerable extent by theoretical studies on desirable properties of risk measures (Artzner et al., 1999; Pedersen and Satchell, 1998, 2002; De Giorgi, 2005) (often related to the regulatory discussion on the use of Value-at-Risk and Expected Shortfall in banks), and the growth of asset classes like hedge funds and derivatives which exhibit very non-Gaussian return characteristics.

In practice, these new performance measures are mainly used for ex post comparison of different funds or strategies (Bacon, 2008), but rarely for ex ante optimisation. There are so far only few papers that extend the empirical studies on mean–variance optimisation to alternative specifications of the objective functions, in particular by replacing the variance by an alternative specification like Value-at-Risk (VaR) or drawdown. The main reason is the difficulty to optimise portfolios with such objective functions, in particular in conjunction with constraints and real-world data, since the resulting optimisation problems are often not convex and cannot be solved with standard techniques (like linear or quadratic programming). The few existing papers (eg, Biglova et al. (2004), Farinelli et al. (2008)) consequently either use only a very small number of assets and do not test realistic conditions like transaction costs, or restrict themselves to those models to which standard solvers can be applied (eg, Racheva-Iotova and Stoyanov (2008)). Thus, so far there exists little evidence regarding the general 'value-added' of these measures when applied in portfolio optimisation. On the contrary, there are studies that compare the rankings of funds (in particular hedge funds) according to downside risk measures with those obtained from Sharpe ratios. Even though using rankings may not be equivalent to a veritable portfolio optimisation, the result is quite clear: while many funds demonstrably do not have Gaussian distributions, the Sharpe rankings are virtually identical to rankings based on alternative performance measures (Eling and Schuhmacher, 2007; Brooks and Kat, 2002).

Our approach is to decompose alternative performance measures into their building blocks, like partial moments or quantiles, and then to test whole classes of performance measures (risk–reward ratios) based on these building blocks. As an example, we do not just investigate the ratio of lower to upper partial moment with exponent one (the Omega function, see Keating and Shadwick (2002)), but test such ratios of partial moments for many different exponents. This gives us an indication whether, generally, partial moments are an effective element to be included in portfolio selection criteria.

Our overall results indicate that incorporating alternative reward and risk measures into the portfolio optimisation process does result in improvements over MV. This improvement is clearest when long-run returns are concerned, where for instance portfolios selected by minimising lower partial moments perform very well; even after risk-adjustment alternative selection criteria seem preferable to MV. Our aim is not to conduct a true 'horse race' between different objective functions, but rather to investigate statistical properties of different portfolio strategies. We thus provide statistics on out-of-sample returns, volatilities, and other moments.

We do not discuss estimation issues in this paper. This is not because this issue is not relevant – on the contrary, it is the main issue in portfolio optimisation. Changing the objective function, in particular using selection criteria that only consider functions of the lower tail of the return distribution as risk, considerably exacerbates the estimation problem. It is important to stress that these estimation difficulties are not just a nuisance: without taking care of the data, there is evidence that optimised portfolios give at best results similar to those of simple strategies (eg, portfolios of equal weights). Instead on relying solely on historical data, we thus use a resampling method to approximate the distribution function of returns. This technique greatly mitigates the overfitting, and leads to enormous improvements in out-of-sample performance (see Gilli and Schumann (2009a)).

Finally we should mention optimisation. Many of the portfolio optimisation problems discussed in this paper lead to non-convex optimisation problems that are in most cases infeasible for classical optimisation methods. We use Threshold Accepting, which incidently was, to our knowledge, the first optimisation heuristic used for (non-mean–variance) portfolio optimisation (Dueck and Winker, 1992). This paper focuses on the empirical results, thus we do not discuss the optimisation technique. A detailed description of the application of Threshold Accepting in portfolio management can be found in Gilli and Schumann (in press). Winker (2001) gives a thorough description of the algorithm in general.

The remainder of this paper is structured as follows: Section 2 details the methodology: the investment problem, alternative selection criteria, and the data. Section 3 discusses results, Section 4 concludes.

2 Data and Methodology

2.1 The investment problem

An investor wants to allocate an intial wealth v_0 among $n_\mathcal{A}$ risky assets (there is no riskfree asset). Given a vector of initial prices p_0 of these assets, the budget constraint is

$$v_0 = x' p_0 \, , \tag{1}$$

where x is the vector of holdings in terms of shares or contracts; corresponding weights w can be obtained by dividing Equation (1) by v_0.

The chosen portfolio x is held over a specified horizon until time T. There is no portfolio rebalancing between time 0 and time T, hence end-of-period wealth is

$$v_T = x' p_T \, ,$$

with corresponding portfolio return

$$r_T = \frac{v_T}{v_0} - 1 \, .$$

Since p_T is not known at time 0, final wealth and thus portfolio returns will be a random variable. A selection criterion for the investor is a function that maps this random variable (or a function of it) into a real number.

2.2 Selection criteria

Information about returns can be summarised by computing moments, hence for a sample of $n_\mathcal{S}$ oberservations, let

$$\mathcal{M}_1 = \frac{1}{n_\mathcal{S}} \sum_{s=1}^{n_\mathcal{S}} r_s$$

and

$$\mathcal{M}_2 = \frac{1}{n_\mathcal{S} - 1} \sum_{s=1}^{n_\mathcal{S}} \Big(r_s - \mathcal{M}_1(r) \Big)^2$$

be the mean and variance of returns, respectively. Markowitz's selection rule states to choose portfolios that are mean–variance efficient. A corresponding objective function, to be minimised, can be written as

$$\mathcal{M}_2 - \lambda\,\mathcal{M}_1\,,$$

with λ a measure of risk-aversion. This function includes only mean and variance, without regard to the overall shape of the return distribution. Furthermore, it only cares for final wealth, not for the path that wealth takes between time 0 and T.

More generally, we may argue that a portfolio of risky assets has a desirable property (here mean return) and an undesirable one (here variance of returns), where we refer to these properties as 'reward' and 'risk', respectively. Below we discuss several 'building blocks' for alternative reward and risk functions. We will only look at such selection criteria (we use the term 'objective function' interchangeably with 'selection criterion') that are based on end-of-period returns; we will not investigate criteria which need the time path of wealth (like drawdown, see Gilli and Schumann (2009a)).

Our objective functions Φ will be ratios of risk and reward to be minimised. There are many special cases that have been proposed in the literature, the best-known certainly being the Sharpe ratio (Sharpe, 1966). For this ratio, reward is the mean portfolio return and risk is the standard deviation of portfolio returns. (Throughout this paper, we set the riskfree rate to zero.) Rachev et al. (2005) gives an overview of various alternative measures that have been proposed in the academic literature; further possible specifications come from financial advisors, in particular from the hedge fund and CTA field (Bacon, 2008, ch. 4). Ratios have the advantage of being easy to communicate and interpret (Stoyanov et al., 2007). Even though numerically, linear combinations are often more stable and thus preferable, working with ratios practically never caused problems in our experiments. We generally 'safeguarded' our objective function, though, for cases where numerator or denominator could switch signs while moving through the search space. Ratios that use the mean return for reward, for instance, are not directly interpretable anymore if mean returns are negative.

The optimisation problem can be summarised as follows: let $x = [x_1\ x_2\ \dots\ x_{n_A}]'$ be the holdings of the individual assets, and $\mathcal{J}$ be the set of assets in the portfolio. Then the problem, including constraints, can be written as

$$\begin{aligned} &\min_x\ \Phi \\ &x_j^{\inf} \leq x_j \leq x_j^{\sup} \qquad j \in \mathcal{J}\,, \\ &K_{\inf} \leq \#\{\mathcal{J}\} \leq K_{\sup}\,. \end{aligned} \tag{2}$$

$x_j^{\inf}$ and $x_j^{\sup}$ are minimum and maximum holding sizes, respectively, for those assets included in the portfolio (ie, those in $\mathcal{J}$). $K_{\inf}$ and $K_{\sup}$ are cardinality constraints which set a minimum and maximum number of assets in $\mathcal{J}$.

Partial moments

Partial moments are a convenient way to distinguish between returns above and below a desired return threshold r_d, and to capture potential asymmetry around this threshold. For a sample of $n_{\mathcal{S}}$ return observations, partial moments $\mathcal{P}_\gamma^{(\cdot)}(r_d)$ can be estimated as

$$\mathcal{P}_\gamma^+(r_d) = \frac{1}{n_{\mathcal{S}}} \sum_{r>r_d} (r - r_d)^\gamma \,, \tag{3a}$$

$$\mathcal{P}_\gamma^-(r_d) = \frac{1}{n_{\mathcal{S}}} \sum_{r<r_d} (r_d - r)^\gamma \,. \tag{3b}$$

The superscripts $+$ and $-$ indicate the tail (ie, upside and downside). Partial moments take two more parameters: an exponent γ, and the threshold r_d. The expression '$r > r_d$' indicates to sum only over those returns that are greater than r_d.

A well-known partial moment is the semivariance, given by $\mathcal{P}_2^-(\mathcal{M}_1)$. The square root of this expression, sometimes called 'downside deviation', is used as the risk function in several performance measures like the Sortino and the Upside Potential ratio (Sortino et al., 1999).

Conditional moments

Conditional moments can be estimated by

$$\mathcal{C}_\gamma^+(r_d) = \frac{1}{\#\{r > r_d\}} \sum_{r>r_d} (r - r_d)^\gamma \,, \tag{4a}$$

$$\mathcal{C}_\gamma^-(r_d) = \frac{1}{\#\{r < r_d\}} \sum_{r<r_d} (r_d - r)^\gamma \,, \tag{4b}$$

where again $+$ and $-$ indicate the tail, and '$\#\{r > r_d\}$' is a counter for the number of return observations higher than r_d.

Conditional and partial moments are closely related. For a threshold r_d, the lower partial moment of order γ equals the lower tail's conditional moment of the same order, times the lower partial moment of order 0. That is,

$$\mathcal{P}_\gamma^+(r_d) = \mathcal{C}_\gamma^+(r_d)\mathcal{P}_0^+(r_d)\,,$$
$$\mathcal{P}_\gamma^-(r_d) = \mathcal{C}_\gamma^-(r_d)\mathcal{P}_0^-(r_d)\,.$$

The partial moment of order 0 is simply the probability of obtaining a return beyond r_d. So in words, conditional moments measures the magnitude of returns around r_d, while partial moments also take into account the probability of such returns. For a fixed r_d, both conditional and partial moments convey different information, since both the probability and the conditional moment need to be estimated from the data to obtain a partial moment.

Quantiles

A quantile of a sample $r = [r_1\ r_2\ \dots\ r_{n_\mathcal{S}}]'$ is defined as

$$\mathcal{Q}_q = \mathrm{CDF}^{-1}(q) = \min\{r \mid \mathrm{CDF}(r) \geq q\}\,,$$

where CDF is the cumulative distribution function and q may range from 0% to 100% (we drop the %-sign in subscripts). Thus, the qth quantile is a number $\mathcal{Q}_q$ such that q of the observations are smaller, and $(100\% - q)$ larger than $\mathcal{Q}_q$. Generally, for a given sample several numbers satisfy this definition (Hyndman and Fan, 1996). We always work with the order statistics of the portfolio returns $[r_{[1]}\ r_{[2]}\ \dots\ r_{[n_\mathcal{S}]}]'$, that is, a step function: if k is the smallest integer not less than $q \times n_\mathcal{S}$, then the qth quantile is the $\max(k, 1)$th order statistic. This is consistent with the convention in many statistical packages that $\mathcal{Q}_0$ is the minimum of the sample (an estimator for the worst-case return), and $\mathcal{Q}_{100}$ is the maximum.

Value-at-Risk (VaR) is the loss only to be exceeded with a given, usually small, probability at the end of a defined horizon. Thus, VaR is a quantile of the return distribution; in our notation, VaR for a probability of 1% can be written as $\mathcal{Q}_1$. Quantiles may also be used as reward measures; we could for example maximise a higher quantile (eg, the 90th).

2.3 Scenario generation and optimisation

Our data set consists of more than 500 price series of European companies from the Dow Jones Euro STOXX universe, all denominated in EUR. The data runs from January 1998 to March 2008, thus including phases of rising and declining share prices. For each company, we also have a market capitalisation series; for a given period, we keep only companies with a reasonable minimum market value (more than € 4 billion) as a rough proxy for market liquidity, which leaves between 150 to 200 assets from which to select in a given period.

We use scenario optimisation to obtain portfolio weights. Thus, we firstly construct scenarios, and then find a portfolio that optimises the selection criterion for these scenarios (Dembo, 1991). This approach is not restrictive: if we preferred to work with a parametric model instead, we could use every historical return as one scenario and then estimate the necessary parameters from the scenarios.

In general, the method by which scenarios are generated has great influence on the out-of-sample results of selected portfolios (Gilli and Schumann, 2009a; Gilli et al., 2008). Since our selection criteria only capture cross-sectional dependence (ie, resorting returns would not change the results), we only model the cross-section of returns by a simple regression model, that is

$$r_i = \alpha_i + \beta_i\, r_M + \gamma_i\, r_M^2 + \delta_i\, r_M^3 + \epsilon_i \qquad i = 1, \dots, n_\mathcal{A} \tag{5}$$

where r_i is the return of the ith asset, r_M is the return of an index, and ϵ_i is the part of the return that is uncorrelated with the chosen index. After having estimated α, β, γ and δ for each of the $n_\mathcal{A}$ assets, we resample from the index and from the residuals to obtain new return scenarios. For the results presented here, we used the Euro STOXX 600 as our index.

This approach to scenario generation is not motivated by computational convenience, but the aim is to reduce overfitting by 'throwing away' part of the cross-sectional dependence that is observed in-sample.

In Equations (2), we set $x^{\inf} = 1\%$ and $x^{\sup} = 5\%$. The minimum number of assets in the portfolio $K_{\inf}$ is 10, the upper cardinality is 50. We do not include a riskless asset. Since our algorithm works with actual positions sizes, not weights, a small fraction (order of magnitude of less than 1% of the portfolio) is usually left uninvested.

The portfolio problem described in Equations (2) cannot be solved by standard optimisation techniques like gradient-based methods. There exist approaches to simplify the model until it can be rewritten, for instance, as a linear programme (eg, Mausser et al. (2006)). Such attempts are, however, usually inflexible and cannot accommodate constraints like cardinality restrictions. We will thus use a heuristic method, Threshold Accepting, to solve the problem. Threshold Accepting was introduced by Dueck and Scheuer (1990) and was probably the first heuristic used for non-mean–variance portfolio optimisation (Dueck and Winker, 1992). Winker (2001) gives a thorough overview, Gilli and Schumann (in press) describe the practical implementation for portfolio optimisation.

The optimisation algorithms are written for Matlab R2008a and can be downloaded from `http://comisef.eu`. A single portfolio optimisation over several thousand scenarios takes less than 10 seconds on a PC with an Intel T9300 2.5 GHz with 2 GB of RAM (this includes the generation of the scenarios).

2.4 Moving-window backtest, rebalancing and transaction costs

We implement a rolling-window backtest, thus we optimise the model at point in time t_1 on data from $t_1 - H$ to $t_1 - 1$; H is set to around 260 days (one year). The resulting portfolio is held until $t_2 = t_1 + F$, with F set to around 90 days (three months). At this point, the portfolio is reoptimised, using data from $t_2 - H$ until $t_2 - 1$, and held until $t_3 = t_2 + F$, and so on. In other words, we construct a portfolio using data from the last year, hold the portfolio for three months, and then rebalance. In this manner, we 'walk forward' through the data to compute a wealth trajectory. All results presented later are computed from the out-of-sample paths of wealth, spanning the period from 7 January 1999 to 19 March 2008.

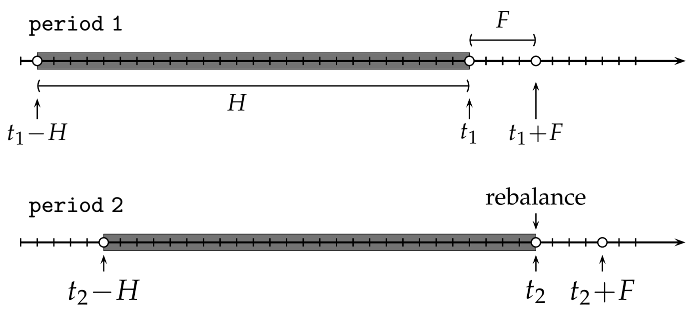

Transaction costs were set to 10 basis points. We also ran backtests with higher transaction costs, and, alternatively, with turnover constraints (not reported here). We found that transaction costs did not significantly influence the results, whereas turnover constraints led to a markedly worse performance.

2.5 Uncertainty

Computing a wealth trajectory from a data set masks a high degree of uncertainty. Firstly, the historical data (both in-sample and out-of-sample) on which the whole procedure is based represents just one realisation of some unknown return-generating process. Secondly, we rely on a resampling procedure to model the data, hence repeatedly generating new scenario sets will generate different optimal portfolios with different ouf-of-sample results. A final source of uncertainty results from the optimisation procedure itself: Threshold Accepting is a stochastic algorithm, hence the resulting optimal portfolios of repeated runs will likely differ from one another (Gilli and Winker, 2009), again producing different results in the out-of-sample periods. These sources of uncertainty are, to our knowledge, rarely addressed in the literature (Gilli and Schumann, 2009b).

To judge the sensitivity of our problem with respect to small data changes, we implement a robustness check based on the following idea: assume a small number of in-sample observations is randomly selected and deleted. The historical return series change, the scenarios will be different, and the composition of the optimal portfolio will change. If the portfolio selection method is robust, we would expect the resulting portfolio to be similar to the original one, as the change in the historical data is only small, and we would also expect the new portfolio to exhibit a similar performance. Repeating this procedure many times, we obtain a sampling distribution of portfolio weights, and hence also a sampling distribution of out-of-sample paths, which gives us an indication of the sensitivity of our selection strategy to a particular data set. This procedure is analogous to repeatedly reestimating a regression equation from jackknifed data to approximate the sampling distribution of the coefficients. We do not compare the differences in the obtained portfolio weights – this would be the equivalent to the regression coefficients – since it is difficult to judge what a given norm of the difference between two weight-vectors practically means. Rather, we look at the changes in out-of-sample results. (In the regression model analogy, the analyst would look at a distribution of forecast errors.) The whole procedure is summarised in Algorithm 1.

Algorithm 1 Robustness check.

set $\mathcal{D} = \{t_1, t_2, \ldots\}$ (rebalancing dates)
set H = length of historical window, F = length of out-of-sample window
set N = number of replications
for $n = 1 : N$ **do**
  for $i = 1 : \#\{\mathcal{D}\}$ **do**
    perturb data from $t_i - H$ to $t_i - 1$
    create scenarios from perturbed data
    optimise portfolio
    store portfolio performance from t_i to $t_i + F$
  end for
  construct nth out-of-sample path
end for

N is set to 100. Algorithm 1 suggests a variety of possible experiments, as different methods for scenario generation may be combined with different criteria to select portfolios.

For this study, scenarios are always generated with the linear regression model (5), and the data perturbation is a jackknifing of about 10% of the observations. Note that we never perturb the out-of-sample data.

3 Results

Our database is not directly comparable with a specific stock market index, so we use as benchmarks minimum-variance (MV) portfolios without short-sales; we also compute 1/N-portfolios (DeMiguel et al., 2009). There is evidence that both of these strategies lead to good out-of-sample results (DeMiguel et al., 2009; Board and Sutcliffe, 1994; Chan et al., 1999; diBartolomeo, 2007). Both selection rules restrict the use of historical data: in the MV case, only variances and covariances are needed, not the means; the 1/N-strategy requires no historical data at all. Figure 1 shows in its upper panel the out-of-sample growth of € 1 invested in the respective strategy from 7. January 1999 to 23. March 2008. (All portfolios, including MV, are computed through resampled scenarios. Rebalancing dates are the same for the benchmarks and the alternative portfolios. The appendix briefly compares the results for MV for historical data; see also Gilli and Schumann (2009a, Section 4.5).) As we stressed above, such a trajectory suggests a precision that is simply not there (for the sake of completeness: the 1/N-strategy returned 7.4% per year, compared with 13.5% for MV). We apply our robustness check, and end up with 100 paths, which give us a 'band' of outcomes, shown in the middle panel. Rescaling final wealth, we obtain a distribution of annual returns for MV.

More compactly, we can put this distribution into a table. The dark grey bar shows the interquartile range, the lighter grey 'whiskers' show the range between the 10th to the 90th quantile of the data. This likely underestimates the true variability of the returns, since the 10th and 90th quantiles from only 100 observations should not be estimated from order statistics, but still, the method gives an idea of the variability of results. The symbol '*c*' stands for 'constant', so reward equal to *c* and risk equal to $\mathcal{M}_2$ gives the MV-portfolio.

reward	risk	median
c	$\mathcal{M}_2$	13.23

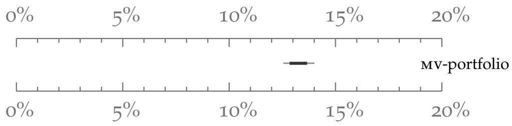

For MV, we have a median return of 13.23%, with a range of about 1.5 percentage points when leaving out the 10 best and the 10 worst paths. We will use two kinds of such tables: the first type (to which the small table just given belongs) contains only information on returns. These tables will appear in the text and are not numbered. Other statistics are presented in full-page tables at the end of the paper; those tables are numbered.

3.1 Partial moments

The following table gives results for selected ratios built from partial moments (results for alternative exponent values can be found in Gilli and Schumann (2009a)).

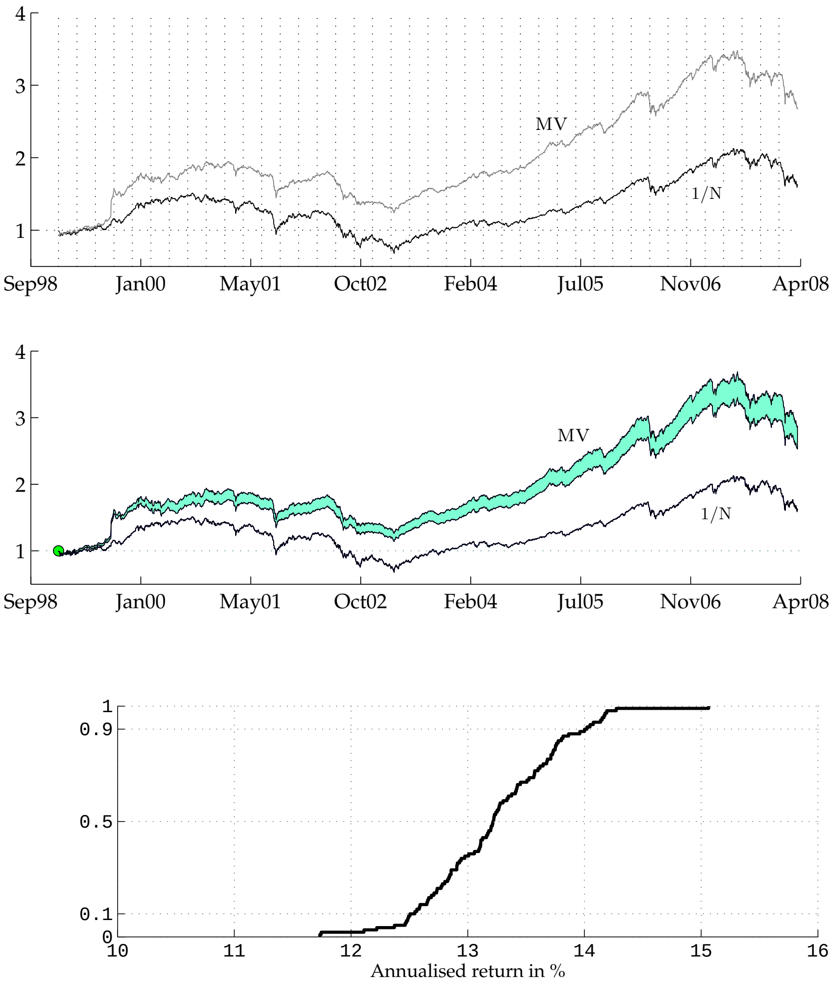

Figure 1: Minimum-variance and 1/N-portfolio.

reward	risk	median	
c	$\mathcal{M}_2$	13.23	MV-portfolio
$\mathcal{P}_0^+$	$\mathcal{P}_0^-$	15.56	
$\mathcal{P}_1^+$	$\mathcal{P}_1^-$	15.83	Omega
$\mathcal{P}_1^+$	$\mathcal{P}_2^-$	15.56	Upside Potential ratio
$\mathcal{P}_{1.5}^+$	$\mathcal{P}_{1.5}^-$	16.23	
$\mathcal{P}_{1.5}^+$	$\mathcal{P}_2^-$	15.95	
$\mathcal{P}_2^+$	$\mathcal{P}_2^-$	15.32	Volatility skewness
$\mathcal{P}_3^+$	$\mathcal{P}_2^-$	6.35	
$\mathcal{P}_4^+$	$\mathcal{P}_2^-$	2.38	
$\mathcal{P}_4^+$	$\mathcal{P}_3^-$	7.76	

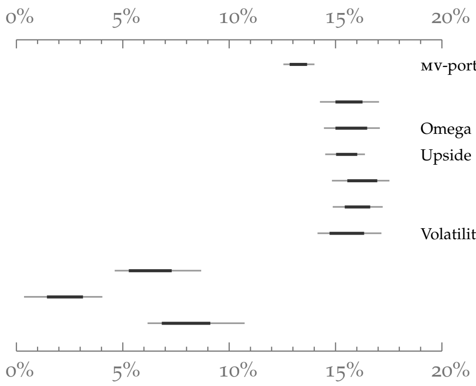

The average annual returns from selection criteria based on partial moments are in many cases 2 to 2.5% higher than for MV. 'In many cases' is to be read as follows: a robust suggestion on how to construct the objective function is not to rely on upper partial moments of a higher order (higher than 1.5, say). Portfolios based on objective functions with $\mathcal{P}^+$ of order 3 or higher perform always poorly, sometimes even resulting in losses over the whole test period. In any case, the curvature of the objective function should be more pronounced for losses than for gains. That is, the exponent γ of the moments should be chosen higher for losses than for gains. The examples at the bottom of the table demonstrate the effect of relying on upper partial moments of higher order.

A further result is that using partial moments of order 0 (ie, frequencies of losses or gains) works well in many instances. The ratio of the frequencies of losses to frequencies of gains, for instance, is a successful objective function even though it ignores any information on the magnitude of returns; combining the frequencies of gains as the reward with a lower partial moment of higher order gives consistently good results.

The higher returns come with a higher variability of returns, at least when measured in terms of volatility. Table 1 shows various statistics for the portfolios. (All statistics are computed from monthly data and are, if applicable, annualised. For the Sharpe and Sortino ratio, the riskfree rate and minimum acceptable return are set to zero. Kurtosis is defined such that a Gaussian variate has a kurtosis of three.) Volatility is mostly similar to the MV case if the lower partial moment included is of order higher then unity and the upper partial moment is of lower order. Interestingly, lower volatility seems to correlate with higher kurtosis and (slightly) higher skewness.

A main result, which is a recurring theme also for other objective functions, is that solely minimising a risk function, and disregarding reward altogether, also leads to the selection of well-performing portfolios. In fact, this is suggested already by the performance of the MV-strategy which is a risk-only strategy. Solely minimising a lower partial moment outperforms MV, even though using upper partial moments of order 0 to 1.5 helps to improve the performance further. The following table and Table 2 present results for lower partial moments. We also added the mean return as the reward function since this conforms with specific objective functions discussed in the literature, for instance the 'Kappa' (Kaplan and Knowles, 2004). As can be seen, including the mean return often improves the average result, but at the price of increased data sensitivity, that is wider distributions, and also higher volatility.

reward	risk	median
c	$\mathcal{M}_2$	13.23
c	$\mathcal{P}_0^-$	15.00
$\mathcal{M}_1$	$\mathcal{P}_0^-$	16.23
c	$\mathcal{P}_1^-$	14.90
$\mathcal{M}_1$	$\mathcal{P}_1^-$	16.18
c	$\mathcal{P}_{1.5}^-$	14.80
$\mathcal{M}_1$	$\mathcal{P}_{1.5}^-$	16.11
c	$\mathcal{P}_2^-$	14.61
$\mathcal{M}_1$	$\mathcal{P}_2^-$	16.18
c	$\mathcal{P}_3^-$	14.83
$\mathcal{M}_1$	$\mathcal{P}_3^-$	15.67

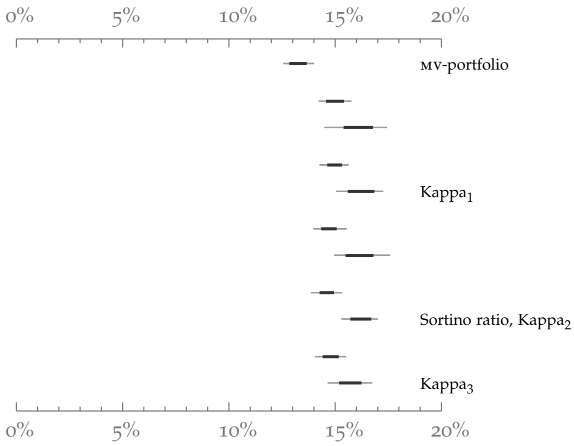

Minimising a lower partial moment of higher order even leads to portfolios with a lower out-of-sample volatility (albeit only marginally so) than MV, but higher returns.

Maximising reward

The results indicate that minimising the historic variability of returns seems an advisable strategy; lower partial moments may be just an alternative way to measure this variability. Upper partial moments are meant to capture reward, but they confound return and risk, since maximising an upper partial moment inevitably also increases variability. In fact, a strategy of solely maximising an upper partial moment (of any order) leads to portfolios that perform very poorly, with a final wealth often barely breaking even over the whole period. In contrast, the rather counterintuitive strategy of minimising an upper partial moment (ie, minimising reward) gives at least positive returns, albeit low ones when compared with other strategies (see table below).

reward	risk	median
c	$\mathcal{M}_2$	13.23
$\mathcal{P}_2^+$	c	−0.70
$\mathcal{P}_3^+$	c	−0.96
c	$\mathcal{P}_2^+$	12.56
c	$\mathcal{P}_3^+$	13.16

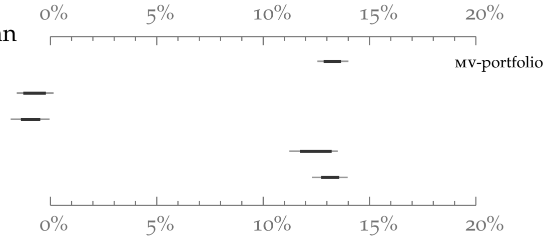

3.2 Conditional moments

The following table gives the results (in terms of annualised returns) for minimising lower conditional moments of order one, where the threshold r_d is defined by a certain quantile. This specification corresponds to minimising Expected Shortfall.

reward	risk	median
c	$\mathcal{M}_2$	13.23
c	$\mathcal{C}_1^-(\mathcal{Q}_1)$	14.88
c	$\mathcal{C}_1^-(\mathcal{Q}_5)$	14.98
c	$\mathcal{C}_1^-(\mathcal{Q}_{10})$	14.58
c	$\mathcal{C}_1^-(\mathcal{Q}_{20})$	14.80

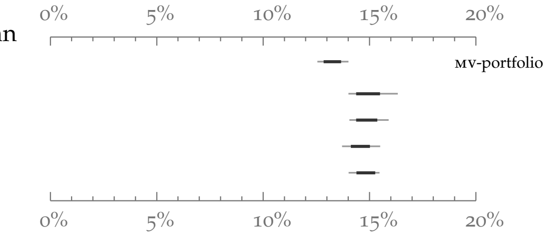

We see that again we improve on MV in terms of returns, with comparable volatility, see Table 3.

Conditional moments offer more possibilities than Expected Shortfall. The following table gives results for ratios of upper to lower conditional moments of different orders. In the literature, this ratio has also been called the Generalised Rachev ratio (Biglova et al., 2004).

reward	risk	median
c	$\mathcal{M}_2$	13.23
$\mathcal{C}_1^+(\mathcal{Q}_{80})$	$\mathcal{C}_1^-(\mathcal{Q}_{20})$	15.98
$\mathcal{C}_1^+(\mathcal{Q}_{80})$	$\mathcal{C}_2^-(\mathcal{Q}_{20})$	15.54
$\mathcal{C}_{1.5}^+(\mathcal{Q}_{80})$	$\mathcal{C}_{1.5}^-(\mathcal{Q}_{20})$	15.84
$\mathcal{C}_2^+(\mathcal{Q}_{80})$	$\mathcal{C}_2^-(\mathcal{Q}_{20})$	15.21
$\mathcal{C}_1^+(\mathcal{Q}_{50})$	$\mathcal{C}_1^-(\mathcal{Q}_{50})$	16.36
$\mathcal{C}_1^+(\mathcal{Q}_{50})$	$\mathcal{C}_2^-(\mathcal{Q}_{50})$	15.50
$\mathcal{C}_{1.5}^+(\mathcal{Q}_{50})$	$\mathcal{C}_{1.5}^-(\mathcal{Q}_{50})$	16.95
$\mathcal{C}_2^+(\mathcal{Q}_{50})$	$\mathcal{C}_2^-(\mathcal{Q}_{50})$	16.57

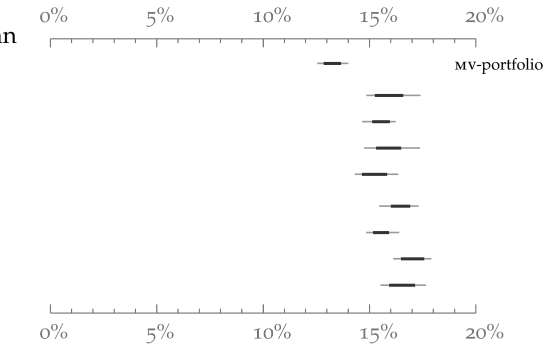

In terms of annualised returns, these ratios improve by two to three percentage points over MV. Table 4 gives additional results. Again, we see that volatility is low if the order (the exponent γ) of the lower moment is higher than the order of the upper moment.

We may also use the mean return as a reward function; results are given below.

reward	risk	median
c	$\mathcal{M}_2$	13.23
$\mathcal{M}_1$	$\mathcal{C}_1^-(\mathcal{Q}_{20})$	16.75
$\mathcal{M}_1$	$\mathcal{C}_{1.5}^-(\mathcal{Q}_{20})$	16.27
$\mathcal{M}_1$	$\mathcal{C}_2^-(\mathcal{Q}_{20})$	15.99
$\mathcal{M}_1$	$\mathcal{C}_3^-(\mathcal{Q}_{20})$	15.83

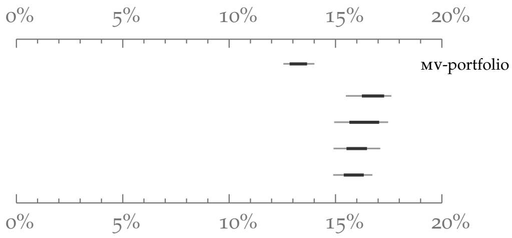

Again, the results suggest that lower conditional moments alone (ie, the risk measures) are already effective selection criteria for portfolios. Including the mean increases the return, but also the volatility. Increasing γ for the lower moment reduces volatility.

3.3 Quantiles

The following table gives results for VaR. The quantiles have been chosen to correspond to the tests for Expected Shortfall, given in the table on page 14. Average returns improve over MV by about less than one percentage points but the results are less stable (ie, wider distributions). Table 6 also shows that skewness is lower, often negative, indicating unfavourably asymmetric distributions.

reward	risk	median
c	$\mathcal{M}_2$	13.23
c	$\mathcal{Q}_1$	14.44
c	$\mathcal{Q}_5$	14.17
c	$\mathcal{Q}_{10}$	14.19
c	$\mathcal{Q}_{20}$	13.70

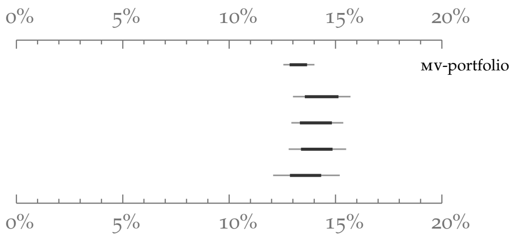

We also form ratios of the form $-\mathcal{Q}_{\text{lo}}/\mathcal{Q}_{\text{hi}}$, where we assume and check that $\mathcal{Q}_{\text{lo}}$ is negative, and $\mathcal{Q}_{\text{hi}}$ is positive. The results, given below, show no improvements over just minimising VaR; average returns deteriorate slightly, and data sensitivity increases.

reward	risk	median
c	$\mathcal{M}_2$	13.23
$\mathcal{Q}_{80}$	$\mathcal{Q}_{10}$	13.30
$\mathcal{Q}_{90}$	$\mathcal{Q}_{10}$	13.26
$\mathcal{Q}_{80}$	$\mathcal{Q}_{20}$	13.46
$\mathcal{Q}_{90}$	$\mathcal{Q}_{20}$	13.72

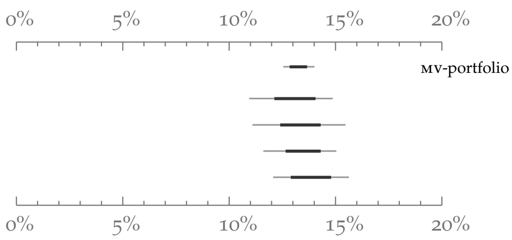

3.4 Asymmetric return distributions

One objection against variance as a risk measure is that is penalises not just the downside, but also positive (desirable) returns. Since the selection criteria presented here are capable of capturing asymmetry in the returns, we would like to see that this asymmetry is reflected in the out-of-sample distributions of portfolios returns.

To make the comparison easier, we rescale the out-of-sample returns of all strategies such that they have an annualised volatility of 14.89% (the median volatility of MV). Next we compute upside and downside deviation, ie, the square roots of $\mathcal{P}_2^+$ and $\mathcal{P}_2^-$. The following table shows the distributions of upside and downside deviation for a selection of objective functions.

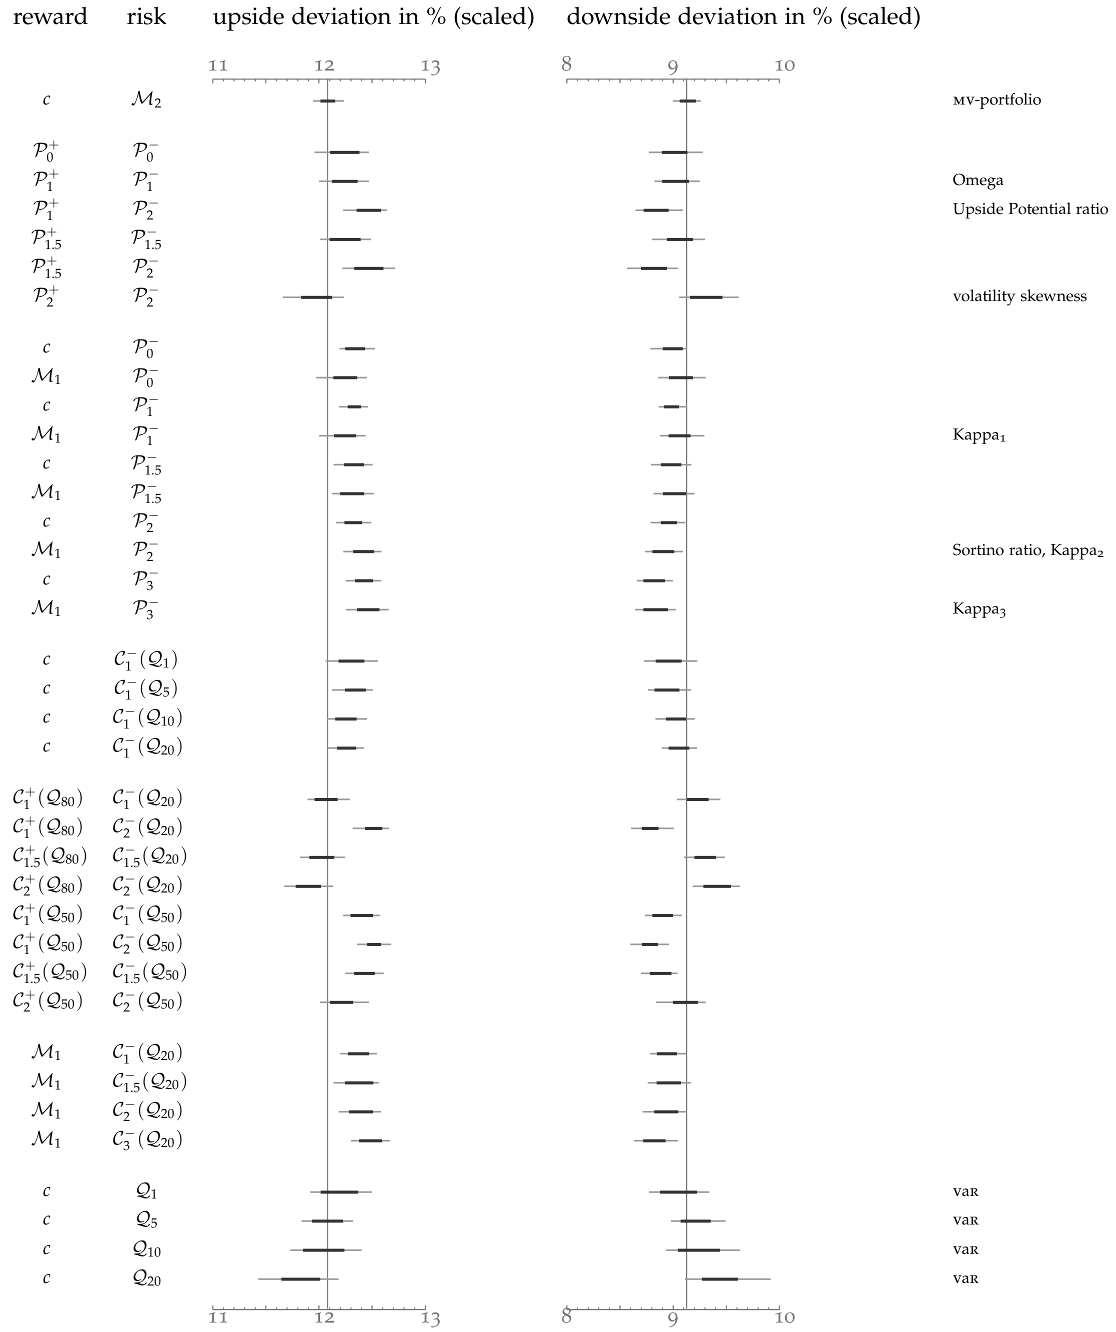

The vertical lines show the median upside and downside deviation of MV. We see that even for MV we have an asymmetry (upside deviation is higher than downside deviation), stemming from the generally positive out-of-sample performance. This asymmetry is more pronounced for many alternative selection criteria. For partial moments, there is again the 'rule' to weigh losses higher than gains, so with γ higher for losses than for gains, we observe this favourable asymmetry. Criteria that minimise lower conditional moments also exhibit this asymmetry.

The results are less clear for ratios of conditional moments; for portfolios built by minimising quantiles we observe unfavourably skewed distributions. The effects are small, anyway: the advantage of an alternative criterion over MV rarely reaches half a percentage point in improvement in either downside or upside. Still, given the general improvement of alternative criteria over MV in terms of returns and variability, this asymmetry is certainly a positive 'side effect' which may be improved further by more dynamic investment strategies.

3.5 Risk, reward and sensitivity

The results given above indicate that solely minimising a risk function leads to portfolios that perform well in terms of returns, while at the same time resulting in portfolio with comparably low return variation (quantiles do not improve on MV, but give reasonable portfolios as well). Adding a reward function often improves the performance of a given portfolio in terms of returns. For instance, for partial moments (see the table on page 13) the average annual return per year was often more than one percentage point higher when a reward function was added. This improvement comes, however, at the cost of higher volatility, and also higher sensitivity to changes in the data. Only looking at the average returns over time cannot demonstrate this, we need to look at single investment periods. In fact, better performance is often caused by only a few periods. Since discussing all selection criteria in detail is well beyond the scope of this paper, we look into concrete example: the Sortino ratio. For convenience we give the returns here again.

reward	risk	median
c	$\mathcal{M}_2$	13.23
c	$\mathcal{P}_2^-$	14.61
$\mathcal{M}_1$	$\mathcal{P}_2^-$	16.18

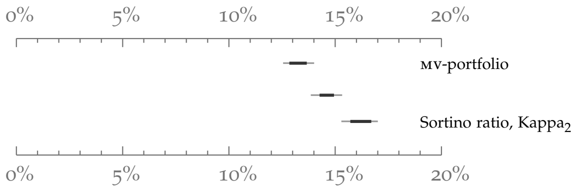

These results suggest that Sortino-optimal portfolios clearly outperform those that solely minimise the downside deviation (the ratio's risk function). We conduct a simple test: we bootstrap from our 40 investment periods, and on the in-sample part (after perturbation) compute a portfolio that only minimises risk, and a portfolio that minimises the risk–reward ratio. We resample 10 000 times, and in each run concatenate the returns of 40 periods (ie, about 10 years, roughly our actual out-of-sample horizon). The distributions of average yearly returns are given in the upper panel in Figure 2.

The picture demonstrates the great variability in returns; the return advantage of the Sortino portfolio is visible with the Sortino ratio's median return more than one percentage points higher. From our 10 000 resampled paths, we get the following probabilities for outperformance of the Sortino ratio.

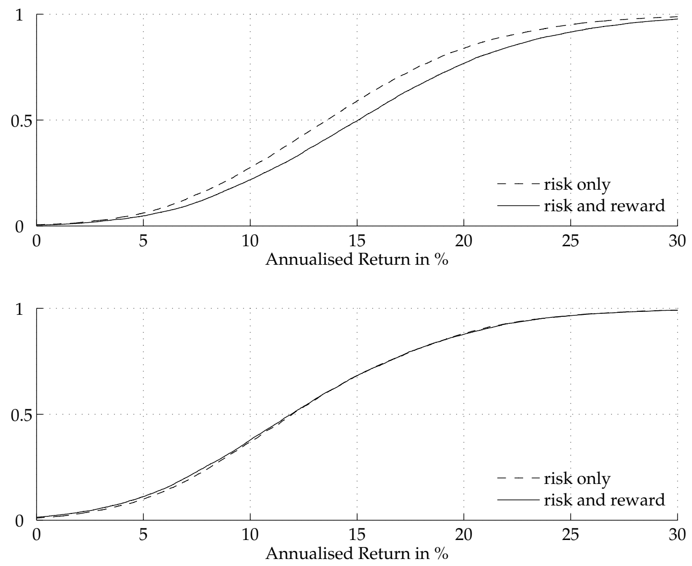

Figure 2: Bootstrapped distributions of yearly returns: All periods (upper panel), best 10% of periods removed (lower panel).

After 1 period (three months):	55% ,
after 4 periods (one year):	63% ,
after 80 periods (twenty years:)	91% .

So even after twenty years, we are not certain that the Sortino ratio outperforms its simpler counterpart without a reward function. In fact, if we remove the best 10% of the periods (4 out of 40), we obtain the following probabilities:

After 1 period (three months):	50% ,
after 4 periods (one year):	51% ,
after 80 periods (twenty years):	47% .

There no more advantage for the Sortino ratio; the lower panel of Figure 2 shows the resulting return distributions. This underlines that there is no stable, continuous outperformance of the risk–reward ratio over the risk-only criterion. Adding the reward ratio is no free lunch in the sense that it does no harm: including reward increases volatility, and also data sensitivity. Recall that for our robustness check we deleted 10% of the original in-sample data, computed a portfolio, and looked at its out-of-sample performance. The average out-of-sample range of returns for one period was 3.4% for risk-only (the same as for MV). For the Sortino ratio, this range grows to 5.5%. In other words, deleting randomly 10% of the in-sample data creates a variation of 5.5% in out-of-sample returns for one period (three months).

4 Conclusion

In this paper we investigated the empirical performance of alternative (ie, non-mean–variance) selection criteria in portfolio optimisation problems. Our main findings are that alternative risk and performance measures in many cases improve on the MV-benchmark. All the strategies tested (including MV) seem very sensitive to relatively small changes in the data.

The results of any study on financial investment depend heavily on the selectable assets and, in particular, the time period under consideration. Hence all quantitative results must be considered conditional on our setting. Still, we can also make some qualitative statements. The recurring theme throughout our study was that minimising risk, as opposed to maximising reward, often lead to good out-of-sample performance; stated differently, low historical variability of portfolio returns was a predictor of good future performance. Our suggestion for constructing objective functions is thus to spend most effort here, as there seem better ways to measure this variability than variance. Selection criteria based on partial and conditional moments performed well, with functions based on quantiles being less satisfying. A careful design of a reward function may improve the strategy in terms of returns, but in many of our tests it also lead to higher return variability and a higher sensitivity to changes in the data sample.

References

Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent Measures of Risk. *Mathematical Finance*, 9(3):203–228, 1999.

Carl A. Bacon. *Practical Portfolio Performance Measurement and Attribution*. Wiley, 2nd edition, 2008.

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.

Almira Biglova, Sergio Ortobelli, Svetlozar Rachev, and Stoyan Stoyanov. Different Approaches to Risk Estimation in Portfolio Theory. *Journal of Portfolio Management*, 31(1): 103–112, 2004.

John L. G. Board and Charles M. S. Sutcliffe. Estimation Methods in Portfolio Selection and the Effectiveness of Short Sales Restrictions: UK Evidence. *Management Science*, 40 (4):516–534, 1994.

Michael W. Brandt. Portfolio Choice Problems. In Y. Aït-Sahalia and L.P. Hansen, editors, *Handbook of Financial Econometrics*. Elsevier, forthcoming.

Beat G. Briner and Gregory Connor. How Much Structure is Best? A Comparison of Market model, Factor Model and Unstructured Equity Covariance Matrices. *Journal of Risk*, 10(4): 3–30, 2008.

Chris Brooks and Harry M. Kat. The Statistical Properties of Hedge Fund Index Returns and Their Implications for Investors . *Journal of Alternative Investments*, 5(2):26–44, 2002.

Louis K. C. Chan, Jason Karceski, and Josef Lakonishok. On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model. *Review of Financial Studies*, 12(5): 937–974, 1999.

Vijay K. Chopra, Chris R. Hensel, and Andrew L. Turner. Massaging Mean–Variance Inputs: Returns from Alternative Global Investment Strategies in the 1980s. *Management Science*, 39(7):845–855, 1993.

Kalman J. Cohen and Jerry A. Pogue. An Empirical Evaluation of Alternative Portfolio-Selection Models. *Journal of Business*, 40(2):166–193, 1967.

Rama Cont. Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. *Quantitative Finance*, 1:223–236, 2001.

Enrico De Giorgi. Reward–Risk Portfolio Selection and Stochastic Sominance. *Journal of Banking & Finance*, 29(4):895–926, 2005.

Ron S. Dembo. Scenario Optimization. *Annals of Operations Research*, 30:63–80, June 1991.

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.

Dan diBartolomeo. The Equity Risk Premium, CAPM and Minimum Variance Portfolios. *Northfield News*, August 2007.

David J. Disatnik and Simon Benninga. Shrinking the Covariance Matrix. *Journal of Portfolio Management*, 33(4):55–63, summer 2007.

Gunter Dueck and Tobias Scheuer. Threshold Accepting. A General Purpose Optimization Algorithm Superior to Simulated Annealing. *Journal of Computational Physics*, 90(1):161–175, 1990.

Gunter Dueck and Peter Winker. New Concepts and Algorithms for Portfolio Choice. *Applied Stochastic Models and Data Analysis*, 8(3):159–178, 1992.

Martin Eling and Frank Schuhmacher. Does the Choice of Performance Measure Influence the Evaluation of Hedge Funds? *Journal of Banking & Finance*, 31:2632–2647, 2007.

Simone Farinelli, Manuel Ferreira, Damiano Rossello, Markus Thoeny, and Luisa Tibiletti. Beyond Sharpe ratio: Optimal Asset Allocation using Different Performance Ratios. *Journal of Banking & Finance*, 32(10):2057–2063, 2008.

George M. Frankfurter, Herbert E. Phillips, and John P. Seagle. Portfolio Selection: The Effects of Uncertain Means, Variances, and Covariances. *Journal of Financial and Quantitative Analysis*, 6(5):1251–1262, 1971.

Manfred Gilli and Enrico Schumann. An Empirical Analysis of Alternative Portfolio Selection Criteria. *Swiss Finance Institute Research Paper No. 09-06*, 2009a.

Manfred Gilli and Enrico Schumann. Optimal enough? *COMISEF Working Paper Series No. 10*, 2009b.

Manfred Gilli and Enrico Schumann. Portfolio optimization with "Threshold Accepting": a practical guide. In Stephen E. Satchell, editor, *Optimizing Optimization: The Next Generation of Optimization Applications and Theory*. Elsevier, in press.

Manfred Gilli and Peter Winker. Heuristic optimization methods in econometrics. In David A. Belsley and Erricos Kontoghiorghes, editors, *Handbook of Computational Econometrics*. Wiley, 2009.

Manfred Gilli, Enrico Schumann, Giacomo di Tollo, and Gerda Cabej. Constructing Long/Short Portfolios with the Omega Ratio. *Swiss Finance Institute Research Paper No. 08-34*, 2008.

Rob J. Hyndman and Yanan Fan. Sample quantiles in statistical packages. *The American Statistician*, 50(4):361–365, 1996.

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

Philippe Jorion. International portfolio diversification with estimation risk. *Journal of Business*, 58(3):259–278, 1985.

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

Paul D. Kaplan and James A. Knowles. Kappa: A Generalized Downside Risk-Adjusted Performance Measure. *Journal of Performance Measurement*, 8(3):42–54, 2004.

Con Keating and Bill Shadwick. An Introduction to Omega. *AIMA Newsletter*, April 2002.

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

Helmut Mausser, David Saunders, and Luis Seco. Optimising Omega. *Risk Magazine*, (11): 88–92, November 2006.

Christian S. Pedersen and Stephen E. Satchell. An Extended Family of Financial-Risk Measures. *Geneva Papers on Risk and Insurance Theory*, 23:89–117, 1998.

Christian S. Pedersen and Stephen E. Satchell. On the foundation of performance measures under asymmetric returns. *Quantitative Finance*, 2:217–223, 2002.

Svetlozar T. Rachev, Frank J. Fabozzi, and Christian Menn. *Fat-Tailed and Skewed Asset Return Distributions: Implications for Risk Management, Portfolio Selection, and Option Pricing*, chapter 13: Risk Measures in Portfolio Optimization and Performance Measures, pages 199–212. Wiley, 2005.

Borjana Racheva-Iotova and Stoyan Stoyanov. Post-Modern Approaches for Portfolio Optimization. In *Handbook on Information Technology in Finance*. Springer, 2008.

William F. Sharpe. Mutual fund performance. *Journal of Business*, 39(1):119–138, 1966.

William F. Sharpe. The Sharpe Ratio. *Journal of Portfolio Management*, 21(1):49–58, Fall 1994.

Frank Sortino, Robert van der Meer, and Auke Plantinga. The Dutch Triangle. *Journal of Portfolio Management*, 26(1):50–58, 1999.

Stoyan V. Stoyanov, Svetlozar T. Rachev, and Frank J. Fabozzi. Optimal Financial Portfolios. *Applied Mathematical Finance*, 14(5):401–436, 2007.

Peter Winker. *Optimization Heuristics in Econometrics: Applications of Threshold Accepting*. Wiley, 2001.

A Minimum variance portfolios with scenarios and historical data

Our benchmark MV-portfolios were computed from resampled scenarios. We tested the change in performance when using historical data instead, but results were roughly similar. Data sensitivity was higher for historical data, annual returns spanned a range of more than five percentage points (11 to 16% instead of 12 to 15%, see Figure 1). While median returns were slightly higher for historical data (13.7% compared with 13.2% for scenarios), a main difference concerned the skewness of the distribution: on historical data, the selected portfolios exhibited a markedly lower, often even negative, skewness.

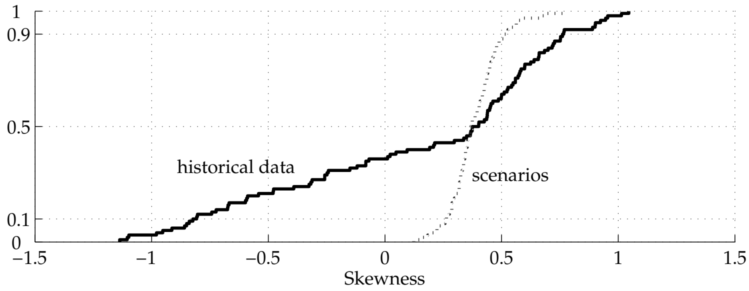

Figure 3: Skewness of MV-portfolios.

Hence, to have a fair comparison, we compared results for portfolio optimised on resampled data.

B Tables

The following tables give various statistics for the different selection criteria. All statistics are computed from monthly return data and are, if applicable, annualised. For the Sharpe and Sortino ratio, the riskfree rate and minimum acceptable return, respectively, are set to zero. Kurtosis is defined such that a Gaussian variate has a kurtosis of three.

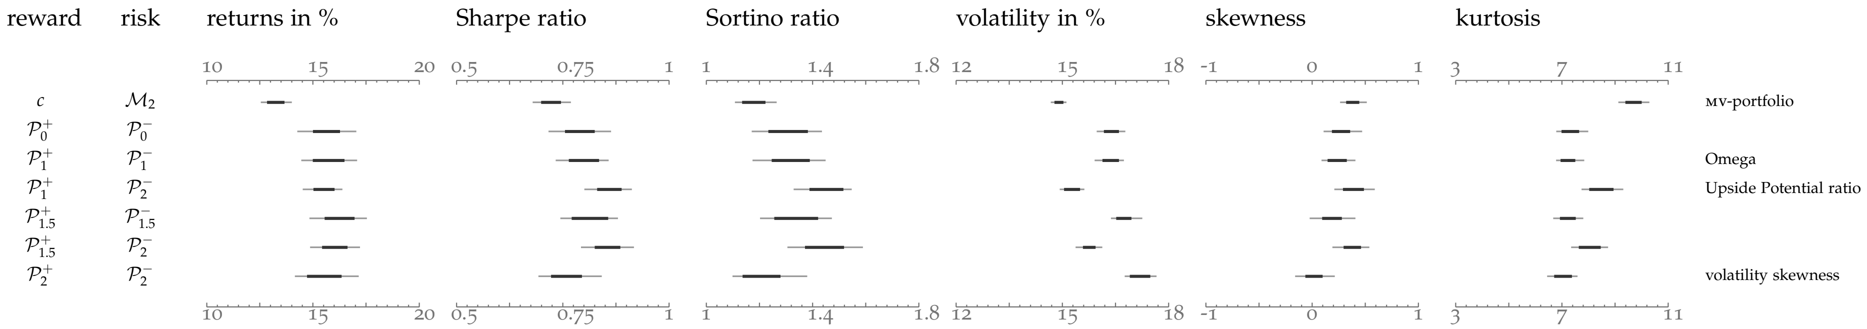

Table 1: Minimising ratios of partial moments.

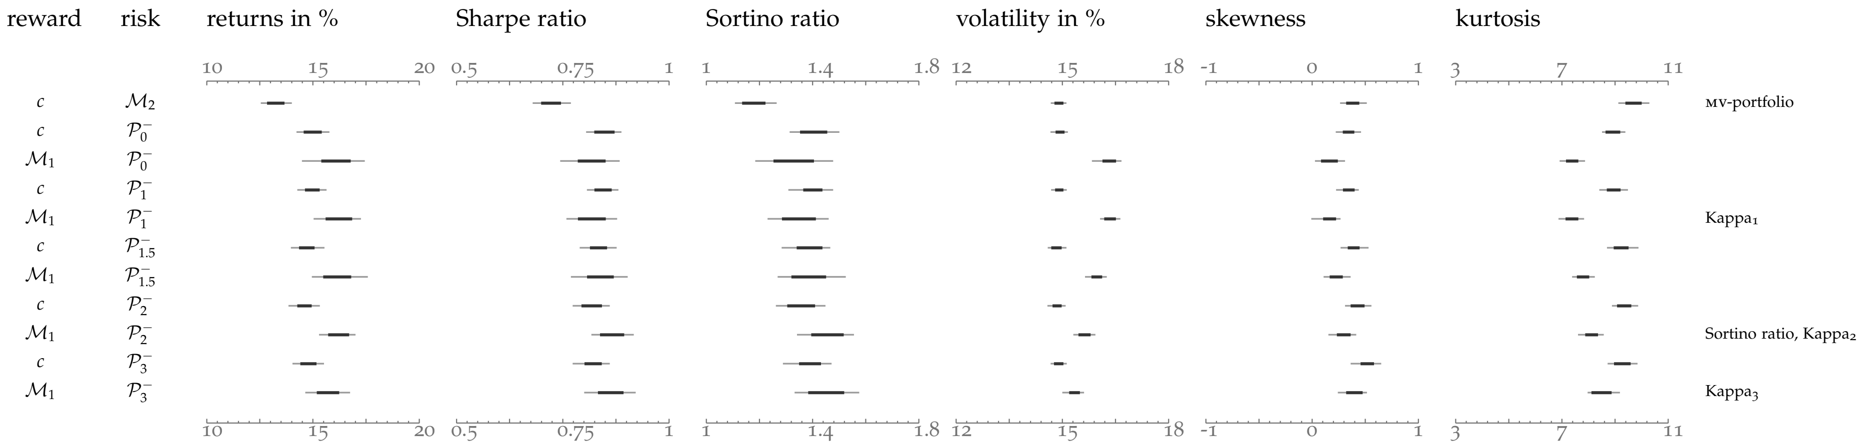

Table 2: Minimising partial moments.

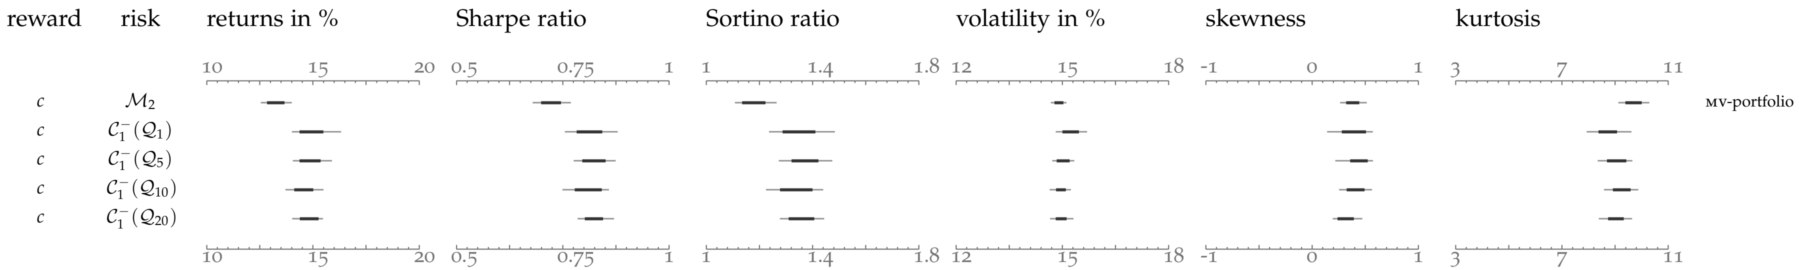

Table 3: Minimising conditional moments: Expected Shortfall.

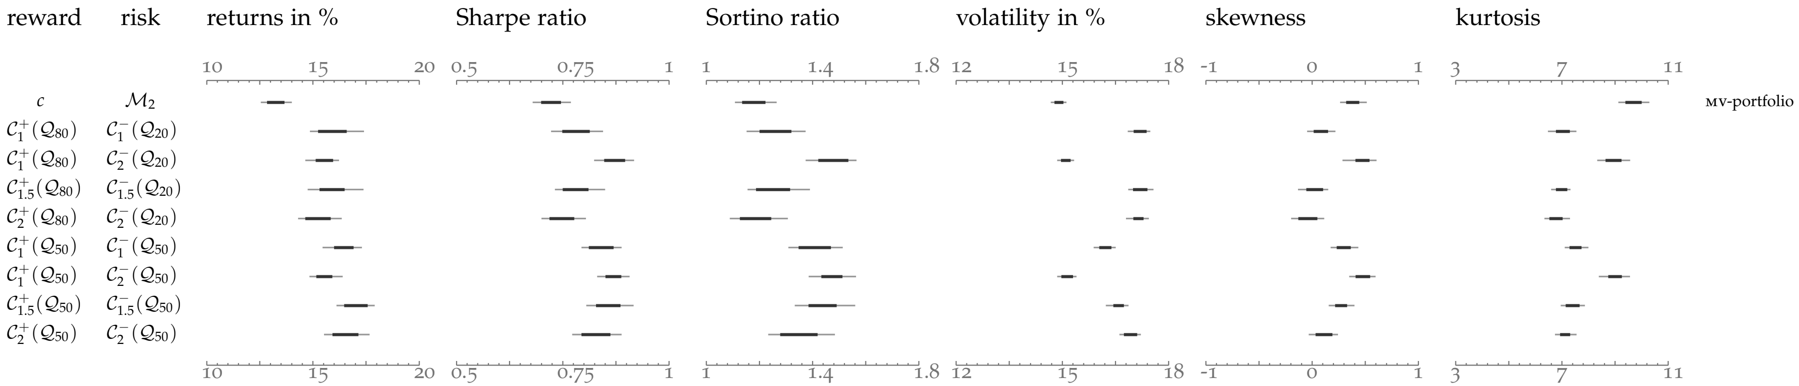

Table 4: Minimising ratios of conditional moments.

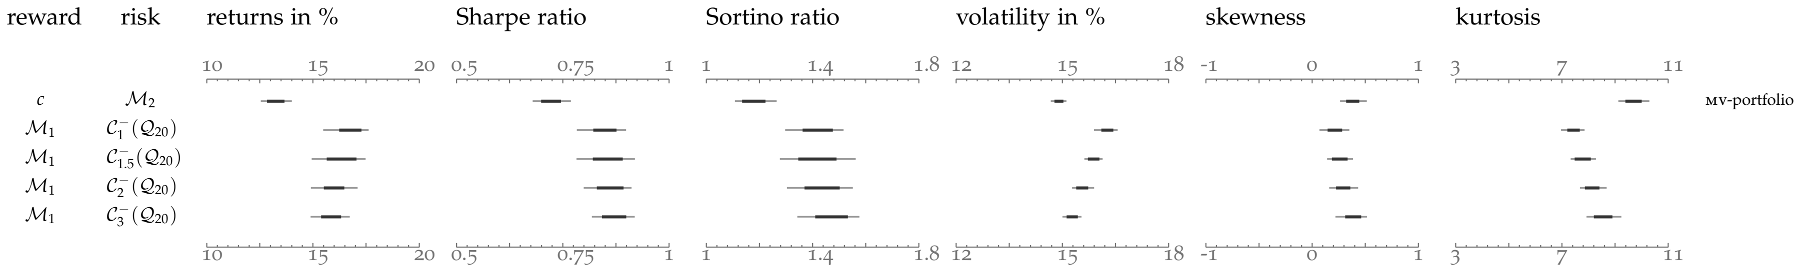

Table 5: Minimising ratios of conditional moments: mean as the reward function.

reward	risk	returns in %	Sharpe ratio	Sortino ratio	volatility in %	skewness	kurtosis	
c	$\mathcal{M}_2$							MV-portfolio
c	$\mathcal{Q}_1$							VAR
c	$\mathcal{Q}_5$							VAR
c	$\mathcal{Q}_{10}$							VAR
c	$\mathcal{Q}_{20}$							VAR

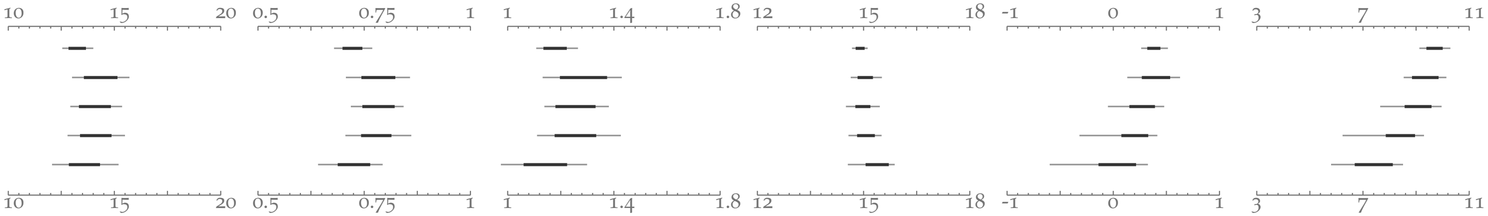

Table 6: Minimising VAR.