

Robust Mean-Variance Portfolio Selection

Cédric Perret-Gentil and Maria-Pia Victoria-Feser

No 2003.02

Cahiers du département d'économétrie
Faculté des sciences économiques et sociales
Université de Genève

Février 2003

Département d'économétrie
Université de Genève, 40 Boulevard du Pont-d'Arve, CH -1211 Genève 4
<http://www.unige.ch/ses/metri/>

Robust Mean-Variance Portfolio Selection

Cédric Perret-Gentil ^a, Maria-Pia Victoria-Feser ^{b,*}

^a *Union Bancaire Privée, 96-98 rue du Rhône, CH-1204 Geneva, Switzerland*

^b *HEC, University of Geneva, Blv. du Pont d'Arve 40, CH-1211 Geneva, Switzerland*

Abstract

This paper investigates model risk issues in the context of mean-variance portfolio selection. We analytically and numerically show that, under model misspecification, the use of statistically robust estimates instead of the widely used classical sample mean and covariance is highly beneficial for the stability properties of the mean-variance optimal portfolios. Moreover, we perform simulations leading to the conclusion that, under classical estimation, model risk bias dominates estimation risk bias. Finally, we suggest a diagnostic tool to warn the analyst of the presence of extreme returns that have an abnormally large influence on the optimization results.

JEL classification: C13; C51; G11

Keywords: Mean-variance efficient frontier; Outliers; Model risk; Robust estimation

*Corresponding author. HEC, University of Geneva, Blv. du Pont d'Arve 40, CH-1211 Geneva, Switzerland; tel.: +41 22 705 8807; fax: +41 22 705 8104. *E-mail addresses:* cedric.perret-gentil@switzerland.org (C. Perret-Gentil), maria-pia.victoriafeser@hec.unige.ch (M.-P. Victoria-Feser).

1 Introduction

The seminal work by Markowitz (1952, 1959) opened the era of modern finance, and the mean-variance framework is the root of modern investment theory. As Britten-Jones (1999) notes: “Mean-variance analysis is important for both practitioners and researchers in finance. For practitioners, theory suggests that mean-variance efficient portfolios can play an important role in portfolio management applications. For researchers in finance, mean-variance analysis is central to many asset pricing theories as well as to empirical tests of those theories; however, practitioners have reported difficulties in implementing mean-variance analysis. For example, Black and Litterman (1992) note that, ‘when investors have tried to use quantitative models to help optimize the critical allocation decision, the unreasonable nature of the results has often thwarted their efforts’ (p.28)”.

To compute the mean-variance efficient frontier and use the information it provides to select the unique optimal portfolio for a given level of risk (or return), we have to know the stochastic mechanism generating the returns for a given set of securities. In its standard formulation, the mean-variance efficient frontier (MVEF) model makes the assumption that the securities’ returns are independently and identically distributed (i.i.d.) as a multivariate normal distribution $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\boldsymbol{\mu}$ is the vector of the securities’ mean returns and $\boldsymbol{\Sigma}$ is the covariance matrix of the securities’ returns.¹ However, the uniqueness of the MVEF solution (i.e. the unique-

¹ The stochastic process generating the data can also include serial dependence. In practice however, the multivariate normal i.i.d. model is widely used.

ness of the optimal weights of the securities in the portfolio given a specific level of risk or return) depends on the implicit assumption that the inputs $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are *known*, whereas they must be *estimated* and therefore are subject to *statistical error*. The MVEF is computed after estimating the model (estimating $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$), and the resulting optimal portfolios might then be heavily biased by this statistical error occurring during the estimation process. This type of error is called *estimation risk* (see for instance Michaud, 1989), and authors have long addressed estimation risk in the context of portfolio selection.²

There is however another type of statistical error, which is concerned with the distribution of the data generating the model. As it is assumed that the observed returns are realizations of a multivariate normal distribution, what happens if this assumption is slightly violated, for instance when one or more of the securities have unexpectedly (non-normal) high or low returns? This model deviation called *model risk* was already addressed in the context of the MVEF model by Victoria-Feser (2000). More recently Cavadini, Sbuelz, and Trojani (2002) have considered model risk and estimation risk simultaneously in the same context of portfolio choices. They suggest different risk aversion corrections attempting to take estimation risk and model risk into account, and quantify their relative importance in a simulated example.

In the statistical literature this problem of model risk is referred to as a problem

² See for instance Barry (1974), Bawa, Brown, and Klein (1979), Alexander and Resnick (1985), Chopra and Ziemba (1993), and more recently ter Horst, de Roon, and Werker (2002).

of *statistical robustness*. The theory of robust statistics is concerned with the construction of statistical procedures that are stable even when the underlying model is not perfectly satisfied by the available dataset. It can deal with a part of the data that is not fully compatible with the distribution implied by the assumed model, i.e. when model misspecification exists, and in particular in the presence of outlying observations.³ Robust estimators have been extensively developed in statistics since the pioneering work of Huber (1964) and Hampel (1968).

A fundamental tool used for studying statistical robustness is the *influence function* proposed by Hampel (1968, 1974). The influence function is useful in determining analytically and numerically the *stability properties* of a statistic in case of model misspecification. In this paper, the influence function is used for studying the behavior of the estimator of the optimal portfolio weights, as well as for building a diagnostic tool to detect outlying returns. We also investigate by means of real market data and simulations how violation of the multivariate normal assumption can seriously affect the optimality characteristics of the solution of the MVEF model when computed with sample mean and sample covariance estimators.

The contributions of this paper to the literature are therefore twofold. First, we show analytically that the necessary and sufficient condition for the mean-variance portfolio optimizer to be robust to local nonparametric departures from multivariate normality is that the estimators of the model's parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ be robust with

³ Outlying observations can be defined as data points that have an infinitesimal probability of being generated by the model generating the rest of the data.

bounded influence functions. We suggest such robust estimators and show their remarkable behavior in the presence of outlying observations in a simulation study. Second, we present a diagnostic tool based on the influence function for detecting the outlying data from the sample that have an abnormally large influence on the optimization process. This tool assesses the quality of the data before their use in the optimizer, and is of particular interest to the analyst.

The paper is organized as follows. In section 2, we present two examples of mean-variance optimization applied to real market data. We show that a robust optimizer can lead to optimal portfolios that are different from the ones given by a classical (non-robust) optimizer, and that this difference is often due to only a few outlying returns from the sample (i.e. under slight model deviations). Section 3 contains the theoretical aspects of this paper. We first present the basic concepts of robust statistics and then use them to study the robustness properties of a portfolio optimizer. In particular, we show analytically that if the model parameters are not robustly estimated, the resulting optimizer can be seriously biased, leading to sub-optimal portfolio choices. We also develop a diagnostic tool based on the influence function for detecting influential returns from the sample. In section 4, a simulation study is performed to investigate the stability properties of the robust optimizer when compared with its classical counterpart. Section 5 concludes.

2 Illustrative examples

2.1 Diversified portfolio

This example replicates the kinds of security that investors may hold in their portfolios. The data set is composed of series of monthly returns on bonds, stocks and alternatives (hedge funds) represented by the following indices.⁴ For the bonds, we use a Merrill Lynch index available on Bloomberg (IND H355 <GO>). This index includes government and corporate bonds with ratings ranging from AAA to A, focuses on maturities between 1 and 10 years and is diversified across the following markets: United States, Canada, United Kingdom, Euroland, Switzerland, Denmark, Sweden, Japan and Australia. The well-known Morgan Stanley Capital International (MSCI) world developed market index is used for the stocks, and its data are collected from Datastream. The two above-mentioned indices are total-return indices (i.e. including reinvestment of coupons/dividends).

The alternative investment part of the portfolio is represented by the Credit Suisse First Boston (CSFB) / Tremont hedge fund index, and data come from their web site (www.hedgeindex.com). For the three above-mentioned asset classes, monthly returns span January 1997 to December 2002, for a total of 72 observations.⁵ The currency of all indices is the US dollar (USD). The returns of these three series

⁴ It can be argued that even if series of returns from the financial reality may exhibit autocorrelation, this practical case is of interest and raises a few important questions that will be addressed in detail in the simulation study of Section 4 with i.i.d. data.

⁵ From this point on, logarithmic returns are used for calculation.

are presented in Fig. 1.




Fig. 1. USD logarithmic monthly returns of the three series used to build the diversified portfolio. Observations span January 1997 to December 2002 for a total of 72 observations. The 9 vertical lines correspond to the 9 most influential returns detected by the *data influence measure* (see later), and highlight data points 1, 7, 15, 20, 22, 35, 36, 38 and 40.

We notice that some returns of the stocks are extreme, for instance data point 20 (August 1998), with a very low return, and a few returns in 2001 and 2002. On the other hand, data points 7 (July 1997), 36 (December 1999) and 38 (February 2000) exhibit three rather high returns for the alternatives. These extreme returns heavily bias the sample mean, covariance and correlation estimators, and this bias can be seen by comparing the results of the above estimators, which we characterize as *classical* estimators, with those of *robust* estimators for the mean, covariance and correlation of the return series. We suggest the use of the translated-biweight *S*-estimator (TBS estimator) proposed by Rocke (1996) as robust estimator (for details, see section 3). All calculations are done with S-Plus and its numerical

optimizer NUOPT.

The comparison between the classically estimated correlation matrix and its robust counterpart is presented in Fig. 2, in which the shapes of the classical and robust ellipses correspond to the shapes of the scatter plot of the data generated by a bivariate normal distribution with such correlations. The corresponding values of the correlations between the securities are shown in the lower triangular part of the Figure. Classical correlations are shown below the robust correlations.

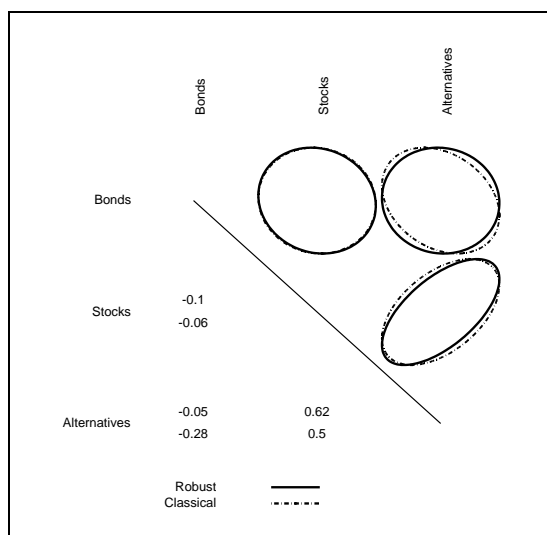


Fig. 2. Double comparison between the results of classical and robust correlation estimation. On the lower triangular part, classical correlations estimated by means of the sample correlation estimator (below) are compared with robust correlations estimated by means of the TBS estimator (above). On the upper triangular part, a comparison of the correlation structure is shown where the shapes of the classical and robust ellipses correspond to the shapes of the scatter plot of the data generated by a bivariate normal distribution with such correlations.

Correlations in terms of numbers are different according to the robust and classical estimates, especially in the case of bonds and alternatives. Moreover, when looking at monthly classical and robust mean and standard deviation of returns as

shown in Table 1, substantial differences can be seen. Such is the case with stocks in the mean estimate, and alternatives in the standard deviation estimate.

Table 1. Classical and robust estimation of mean and standard deviation of monthly returns.

Securities	Mean		Standard Deviation	
	Classical	Robust	Classical	Robust
Bonds	0.31%	0.30%	2.03%	2.07%
Stocks	0.09%	-0.12%	4.95%	5.25%
Alternatives	0.77%	0.68%	2.53%	1.70%

To assess the effect of these differences in estimation on the results of the portfolio optimizer, the classical (calculated with the classically estimated $\hat{\boldsymbol{\mu}}^{\text{cla}}$ and $\hat{\boldsymbol{\Sigma}}^{\text{cla}}$) and the robust (calculated with the robustly estimated $\hat{\boldsymbol{\mu}}^{\text{rob}}$ and $\hat{\boldsymbol{\Sigma}}^{\text{rob}}$) mean-variance efficient frontiers⁶ are presented in Fig. 3.

Classically and robustly estimated efficient frontiers give an important insight into the statistical properties of the data. The two efficient frontiers are distinct from each other, meaning that the influential data points that affect classical estimation are *treated* by robust estimation, so that their influence is kept *under control*. The robust efficient frontier, being higher and more to the left than the classical one, clearly indicates a reduction in the volatility structure of the inputs due to their robust estimation. Other interesting information is given by the composition of the so-called *minimum variance portfolios*, which are the portfolios characterized by the minimum risk on the two efficient frontiers. The classical and robust weights

⁶ Calculation in this example has been done without allowing for short selling, as is mainly the case in managed portfolios. This restriction is lifted in section 4.




Fig. 3. Classical mean-variance efficient frontier (MVEF) when the mean and covariance of returns are estimated using the sample mean and sample covariance estimators, and robust MVEF when the mean and covariance of returns are estimated using the TBS estimator. Short selling is not allowed.

of the three assets (bonds, stocks and alternatives respectively) in these particular portfolios are $\hat{\mathbf{p}}^{\text{cla}} = (0.59, 0, 0.41)$ and $\hat{\mathbf{p}}^{\text{rob}} = (0.41, 0, 0.59)$. The fact that the weights of bonds and alternatives are the exact reverse of each other is due to chance, but more worrying is that portfolio compositions that should represent the same reality (i.e. minimum risk) are quite different from each other. Stocks are obviously given a zero weight by the optimizer for such a low-risk portfolio, due to their high standard deviation of returns over the period.

A further step in the analysis is to check which observations are considered outliers and responsible for this shift of the efficient frontier. To do so, we use a diagnostic tool called *data influence measure*, which we present in section 3. Briefly, the data influence measure applied to our case is a tool for detecting outlying returns

that have an abnormally high influence on the estimator of the optimal portfolio weights. The data influence measure for the diversified portfolio data set is given in Fig. 4.




Fig. 4. Relative influence of each of the 72 returns on the estimator of the optimal portfolio weights as detected by means of the data influence measure diagnostic tool. As this measure relates to a specific portfolio on the classical mean-variance efficient frontier, the level of standard deviation of returns has been set to 1.5%.

The 9 returns that appear the most influential (compared with the majority) have also been highlighted in the time series plot of Fig. 1 by means of vertical lines. The striking feature is that these 9 returns do not always correspond to the most extreme returns of the different series, and that makes them difficult to find by simple visual inspection. For example, when looking at Fig. 1, data point 20 (August 1998) seems more statistically outlying than data point 36 (December 1999), which is not true when looking at Fig. 4. In fact, the power of the data influence measure is that it takes into account the multivariate structure of the model and highlights

the influential data points according to the specific estimator used.

In short, we have found that a few outlying observations in the data have a strong influence on the composition of the resulting optimal portfolios.⁷

2.2 Hedge fund portfolio

Let us now turn to another example in which we remain in the same asset class.⁸ We suggest building a portfolio composed only of alternatives (hedge funds) that belong to the arbitrage and relative value strategies. We use the full available return history in US dollars from January 1994 to December 2002, and consider the convertible arbitrage (CA), event-driven (ED) and fixed income arbitrage (FIA) strategies represented by their respective CSFB/Tremont sub-indices. We refer the reader to Lhabitant (2002) for specific details and characteristics of these strategies. The returns of these three hedge fund strategies are presented in Fig. 5. Compared with the preceding example, data seem more stable, with the notable exception of a few data which correspond to the Russian financial crisis of August 1998.

Presented in Fig. 6 is a comparison of the correlation structure under classical and robust estimation. Again, there are differences between classical and robust correlation estimates.

⁷ These 9 most outlying points represent 12.5% of the data.

⁸ As in the previous case, no direct conclusion can be drawn from this specific example as returns may exhibit serial dependence. However, these are the kinds of data that would be used in practice, and their statistical properties are of interest because of the influence a few data may have on the optimization results.




Fig. 5. USD logarithmic monthly returns of the three series used to build the hedge fund portfolio. Observations span January 1994 to December 2002 for a total of 108 observations. The 2 vertical lines correspond to the 2 most influential returns detected by the *data influence measure* (see later), and highlight data points 56 and 57.

The same is true for the classical and robust mean and standard deviation estimates presented in Table 2, where major differences exist between classical and robust estimation.

Table 2. Classical and robust estimation of mean and standard deviation of monthly returns.

Hedge fund strategies	Mean		Standard Deviation	
	Classical	Robust	Classical	Robust
Convertible arbitrage (CA)	0.81%	1.14%	1.41%	0.88%
Event-driven (ED)	0.83%	1.26%	1.86%	1.24%
Fixed income arbitrage (FIA)	0.54%	0.85%	1.20%	0.68%

Hence, even with series exhibiting low standard deviation of returns, classical and robust estimations do not give the same results, and this difference is further

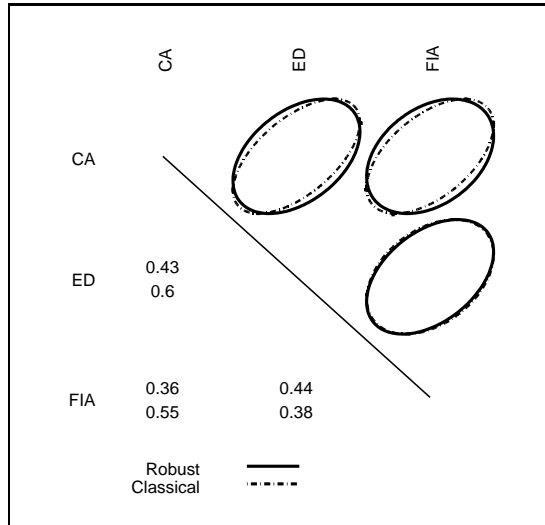


Fig. 6. Double comparison between the results of classical and robust correlation estimation. On the lower triangular part, classical correlations estimated by means of the sample correlation estimator (below) are compared with robust correlations estimated by means of the TBS estimator (above). On the upper triangular part, a comparison of the correlation structure is shown where the shapes of the classical and robust ellipses correspond to the shapes of the scatter plot of the data generated from a bivariate normal distribution with such correlations.

evidenced by classical and robust efficient frontiers⁹ shown in Fig. 7.

As in the previous example, the robust efficient frontier is located higher than and to the left of the classical one, indicating the presence of outlying data points. Looking at the composition of the minimum variance portfolios, the classical and robust weights (for convertible arbitrage, event-driven and fixed income strategies respectively) are $\hat{\mathbf{p}}^{\text{cla}} = (0.26, 0.09, 0.65)$ and $\hat{\mathbf{p}}^{\text{rob}} = (0.31, 0, 0.69)$. As can be seen, the difference in weights is mainly concentrated on the first two hedge fund strategies. Here again, it is interesting to identify the outlying returns responsible for such a change, and we make use of the data influence measure whose results are reported

⁹ Calculation has been done without allowing for short selling




Fig. 7. Classical mean-variance efficient frontier (MVEF) when the mean and covariance of returns are estimated using the sample mean and covariance estimators, and robust MVEF when the mean and covariance of returns are estimated using the TBS estimator. Short selling is not allowed.

in Fig. 8. The 2 most influential return points have also been highlighted in the serial plot of Fig. 5 by means of vertical lines. Data points 56 (August 1998) and 57 (September 1998) correspond to the very low returns recorded during the Russian financial crisis, and they are detected as having a very strong relative influence on the classical estimates of the optimal portfolio weights.

In short, just a few outlying data points from the sample can have a very strong influence on the classical estimates needed to calculate the efficient frontier and thus lead to different portfolio choices.¹⁰

At this point, we may still wonder what would be a strong enough reason for choosing a robust portfolio composition rather than a classical one, as both kinds of

¹⁰ These 2 strongly outlying observations represent 1.85% of the total 108 observations.




Fig. 8. Relative influence of each of the 108 returns on the estimator of the optimal portfolio weights, as detected by means of the data influence measure diagnostic tool. As this measure relates to a specific portfolio on the classical mean-variance efficient frontier, the level of standard deviation of returns has been set to 1.3%.

optimal portfolio could be considered acceptable. In fact, the reason for preferring a robust portfolio composition will become obvious in section 4, and will be strongly motivated by the sensitivity of the sample mean and covariance estimated MVEF model to the data as shown by a simulation study.

In the next section we review the basic concepts of robust statistics and apply them to the MVEF model.

3 Robustness properties of the MVEF model

3.1 Basic concepts of robust statistics

The pioneering work of Huber (1964) and Hampel (1968) has laid the ground for the theory of robust statistics. As a generalization of classical theory, robust statistics

takes into account the possibility of model misspecification (i.e. model deviation). This theory and its results are valid *at* the model as well as *in a neighborhood* of the model,¹¹ which is not the case for classical statistics, which is only valid *at* the model.

Define

$$\{G_\varepsilon | G_\varepsilon = (1 - \varepsilon)F_\theta + \varepsilon W\}, \quad (1)$$

with W an arbitrary distribution and $\varepsilon \in [0, 1]$, the set of all distributions defining a neighborhood of the parametric model F_θ . This neighborhood includes all possible misspecified distributions around F_θ , and G_ε can be considered a mixed distribution between F_θ and the contamination distribution W . An estimator is said *robust* if it remains stable in a neighborhood G_ε of F_θ . One particular case is when $W = \Delta_{\mathbf{y}_t}$, the distribution that gives a probability of one to a point \mathbf{y}_t chosen arbitrarily.¹² In this case, the neighborhood of the model featuring all local nonparametric departures from F_θ is given by

$$\{F_\varepsilon | F_\varepsilon = (1 - \varepsilon)F_\theta + \varepsilon \Delta_{\mathbf{y}_t}\}. \quad (2)$$

Hence F_ε generates observations from F_θ with probability $(1 - \varepsilon)$ and observations equal to an arbitrary point \mathbf{y}_t with probability ε .

One way of assessing the robustness properties of an estimator $\hat{\theta}$ of θ is to study

¹¹ In the presence of outlying observations acting as local nonparametric departures from the model, the distributional assumptions are violated and we therefore end up in a *neighborhood* of the model.

¹² \mathbf{y}_t can be a scalar or a vector.

its (asymptotic) stability properties in a neighborhood of the model considering a distribution of type G_ε , and there is no loss of generality in focusing on the particular case of F_ε since Hampel et al. (1986) showed that the maximal bias on $\hat{\theta}$ is obtained at $W = \Delta_{\mathbf{y}_t}$.

Considering the case when ε tends towards zero,¹³ we get the so-called *influence function* (IF) suggested by Hampel (1968, 1974) and further developed by Hampel et al. (1986). The IF gives the influence of an infinitesimal amount of contamination \mathbf{y}_t on the value of the estimator viewed as a function of the underlying distribution. The influence function is then defined as

$$\text{IF}(\mathbf{y}_t, \hat{\theta}, F_\theta) = \lim_{\varepsilon \downarrow 0} \left[\frac{\hat{\theta}(F_\varepsilon) - \hat{\theta}(F_\theta)}{\varepsilon} \right],$$

and, when the derivative exists, as

$$\text{IF}(\mathbf{y}_t, \hat{\theta}, F_\theta) = \left. \frac{\partial}{\partial \varepsilon} \hat{\theta}(F_\varepsilon) \right|_{\varepsilon=0}. \quad (3)$$

The IF is the directional derivative of the estimator $\hat{\theta}$ in a single point contamination direction $\Delta_{\mathbf{y}_t}$. Depending on the situation, this directional derivative can be scalar, vector or matrix valued.

The IF is a powerful tool for assessing the robustness properties of estimators. Indeed, Hampel et al. (1986) show that *only* the IF is needed to fully describe the asymptotic bias of an estimator caused by a contamination, implying that an estimator with a bounded IF automatically has a bounded asymptotic bias. Therefore, an

¹³ We consider an infinitesimal amount of contamination, and therefore remain strictly in a *close* neighborhood of the model.

estimator with a bounded IF is robust in a general neighborhood of the parametric model defined by Eq. (1).

The IF can also be used as a diagnostic tool to detect observations that have a large influence on the estimator $\hat{\theta}$. We suggest using the Euclidean norm of the influence function defined under Eq. (3) and call such a measure the *data influence measure* (DIM). It is given by

$$\text{DIM}(\mathbf{y}_t, \hat{\theta}) = \left(\text{IF}(\mathbf{y}_t, \hat{\theta}, F_{\hat{\theta}})' \text{IF}(\mathbf{y}_t, \hat{\theta}, F_{\hat{\theta}}) \right)^{1/2}. \quad (4)$$

An observation \mathbf{y}_t with a large DIM is then considered an influential observation. It should be noticed that the supremum in \mathbf{y}_t of the DIM gives the *gross error sensitivity* defined in Hampel et al. (1986) and used as a tool for controlling the robustness properties of statistical procedures. To compute the DIM, the true parameter value θ has to be known, which is seldom the case in practice. The value of θ has then to be estimated in a robust way so as to ensure that this diagnostic tool is not biased by the outlying observations it is supposed to detect.

3.2 Classical MVEF model estimation

Let us suppose that there are N securities to choose from and let $\mathbf{p} = (p_1, \dots, p_N)'$ be the vector of portfolio weights, so that $\sum_{i=1}^N p_i = 1$. Recalling that $\boldsymbol{\mu}$ is the vector of size N containing the mean returns of the securities, and that $\boldsymbol{\Sigma}$ is the $(N \times N)$ covariance matrix of the returns of the securities, the portfolio mean can be written as

$$R(\mathbf{p}) = \boldsymbol{\mu}' \mathbf{p},$$

and the portfolio variance as

$$S(\mathbf{p}) = \mathbf{p}' \boldsymbol{\Sigma} \mathbf{p}.$$

For a given value of the risk aversion parameter λ , the mean-variance optimization selects the portfolio \mathbf{p}^* which maximizes in \mathbf{p}

$$R(\mathbf{p}) - \frac{\lambda}{2} S(\mathbf{p}),$$

subject to

$$\mathbf{e}'_N \mathbf{p} = \mathbf{1},$$

where \mathbf{e}_N of size $(N \times 1)$ is a vector of ones. The set of optimal portfolios for all possible values of the risk aversion parameter λ defines the mean-variance efficient frontier. Depending on the situation, the constraint of no short selling ($\mathbf{p} \geq 0$) might be added as well as other constraints.

In the unconstrained case, the solution is explicit and it is well known that the optimal portfolio weights for a given value of λ are given by

$$\mathbf{p}^* = \frac{1}{\lambda} \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu} - \eta \mathbf{e}_N), \quad (5)$$

with

$$\eta = \frac{(\mathbf{e}'_N \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} - \lambda)}{\mathbf{e}'_N \boldsymbol{\Sigma}^{-1} \mathbf{e}_N}. \quad (6)$$

As can be seen, \mathbf{p}^* depends directly on $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, implying that the resulting estimated optimal portfolio weights are directly affected by potential estimation bias in the mean and covariance of returns. This is reflected in the following proposition, where we show that the influence function of the estimator of the optimal portfolio

weights depends directly on the influence functions of the estimators of both $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$.

Proposition 1 *The (asymptotic) bias of the resulting estimator $\hat{\mathbf{p}}^*$ of the optimal portfolio weights only depends on the (asymptotic) bias of the estimators of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$.*

Proof. We show that the IF of the estimator of the weights \mathbf{p}^* is proportional to the IF of the estimators of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ (see Appendix A).

Therefore, unless the mean vector and covariance matrix are robustly estimated, the mean-variance optimizer can lead to portfolio compositions heavily influenced by just a few outlying observations from the sample.

The mean and covariance of the returns are in practice often estimated by their sample counterpart, i.e. by the maximum likelihood (ML) estimator under the multivariate normal model. In this case, the estimator of $\boldsymbol{\mu}$ can be explicitly expressed as

$$\hat{\boldsymbol{\mu}}'_{\text{ML}} = \frac{1}{T} \mathbf{e}'_T \mathbf{Y},$$

where \mathbf{Y} is the $(T \times N)$ matrix containing the returns of each security (columnwise) for T periods, and given by

$$\mathbf{Y} = \begin{bmatrix} \mathbf{y}'_1 \\ \vdots \\ \mathbf{y}'_T \end{bmatrix} = [\mathbf{y}_1 \dots \mathbf{y}_N] = \begin{bmatrix} y_{11} & \dots & y_{1N} \\ \vdots & \ddots & \vdots \\ y_{T1} & \dots & y_{TN} \end{bmatrix}.$$

Moreover, the ML estimator of the covariance matrix Σ is calculated as

$$\widehat{\Sigma}_{\text{ML}} = \frac{1}{T} (\mathbf{Y} - \mathbf{e}_T \hat{\boldsymbol{\mu}}'_{\text{ML}})' (\mathbf{Y} - \mathbf{e}_T \hat{\boldsymbol{\mu}}'_{\text{ML}}),$$

where \mathbf{e}_T of size $(T \times 1)$ is a vector of ones. Under *normality*, the maximum likelihood estimators are the most efficient. However, in the presence of outlying observations, how can they cope with model deviation and what kind of influence have outlying data on them?

To answer this central question we make use of the influence function. As shown in Hampel et al. (1986), the respective IFs for $\hat{\boldsymbol{\mu}}_{\text{ML}}$ and $\widehat{\Sigma}_{\text{ML}}$ are given by

$$\text{IF}(\mathbf{y}_t, \hat{\boldsymbol{\mu}}_{\text{ML}}, F_\theta) = -\boldsymbol{\mu} + \mathbf{y}_t, \quad (7)$$

and

$$\text{IF}(\mathbf{y}_t, \widehat{\Sigma}_{\text{ML}}, F_\theta) = -\Sigma + (\mathbf{y}_t - \boldsymbol{\mu})(\mathbf{y}_t - \boldsymbol{\mu})'. \quad (8)$$

We can easily see that both IFs for the maximum likelihood estimators of $\boldsymbol{\mu}$ and Σ are unbounded, since for large values of some or all elements of \mathbf{y}_t , both IFs may become *arbitrarily large*.

Making use of the above results, the following proposition derives the explicit influence function of the estimator of the optimal portfolio weights when ML estimation is used for $\boldsymbol{\mu}$ and Σ .

Proposition 2 *When sample mean and covariance are used to estimate the mean vector $\boldsymbol{\mu}$ and covariance matrix Σ , the (asymptotic) bias on the resulting estimator $\hat{\mathbf{p}}^*$ of the optimal portfolio weights can be infinite under infinitesimal departures from multivariate normality.*

Proof. The explicit expression for the IF of $\hat{\mathbf{p}}^*$ is given by replacing in Eq. (11) (see Appendix A) the influence functions of the respective ML estimators reported in Eq. (7) and Eq. (8). It is then straightforward to show that the IF of $\hat{\mathbf{p}}^*$ is unbounded (see Appendix B).

In short, the ML-estimated mean-variance optimizer is *not robust to model risk*. It is therefore of interest to identify, before estimating the model, the returns from the sample that will have an abnormally large influence on the optimization results under maximum likelihood estimation. However, it is difficult to *visually* identify these specific returns, due to the multivariate structure of the model and the statistical properties of the ML estimator. To address this issue, we suggest using as a diagnostic tool the data influence measure given in Eq. (4) in conjunction with the influence function given in Eq. (12) (see Appendix B), recalling that $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ have to be estimated in a robust way for the reason already mentioned above. This produces graphs like those in Fig. 4 or 8, and influential data points are then easy to identify. However, it should be stressed that there are an infinite number of IFs (and thus of DIMs) since there are an infinite number of optimal portfolio weighting schemes. The value of the risk aversion parameter λ has to be fixed so as to characterize the portfolio weighting scheme of interest, and to be able to compute the DIM.¹⁴

¹⁴ We found however that the choice of the value for λ does not have a decisive impact on the detection of outlying returns.

3.3 Robust MVEF model estimation

As demonstrated above, the estimators used for the mean vector and the covariance matrix determines the robustness properties of the estimator of the optimal portfolio weights in the optimizer. We should thus choose a robust estimator with a bounded influence function to estimate $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$.¹⁵

We suggest using Rocke's (1996) *translated biweight S-estimator* (TBS estimator), which belongs to the class of *S*-estimators (see Rousseeuw and Leroy, 1987). The TBS estimator has the advantage of being quite efficient compared with other robust estimators. Moreover, its robustness against relatively large quantities of outlying data can be controlled (see below).

The TBS estimator is defined as the solution in $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ of

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T w_t^\mu (\mathbf{y}_t - \boldsymbol{\mu}) &= 0 \\ \frac{1}{T} \sum_{t=1}^T w_t^\delta \boldsymbol{\Sigma} - w_t^\eta (\mathbf{y}_t - \boldsymbol{\mu})(\mathbf{y}_t - \boldsymbol{\mu})' &= 0, \end{aligned}$$

where the weights w_t^μ , w_t^δ and w_t^η depend on two control parameters (see again Rocke, 1996).

The first control parameter is the *breakdown point* ε^* , and the second is the *asymptotic rejection probability* (ARP) α . The breakdown point of an estimator is the maximal amount of model deviation (or model risk) it can withstand before it breaks down, i.e. before it can take arbitrary values under model deviations such as

¹⁵ See for instance Victoria-Feser (2000) for a review and discussion of the desirable properties of such a robust estimator.

in Eq. (2) (see Hampel et al., 1986). The ARP can be interpreted as the probability that an estimator, in large samples and under a reference distribution, will give a null (or nearly null) weight to extreme observations.

The TBS estimator is a consistent estimator of mean and covariance. However, compared with the classical ML estimator, the TBS estimator is less efficient at the model. This loss of efficiency is the *price to pay* for its robust behavior in a neighborhood of the model and its property of safeguarding the estimators of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ against the influence of model deviations. A robust mean-variance portfolio optimizer is thus obtained by estimating $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ by means of the TBS estimator, and by using the resulting estimates in Eq. (5).

It should be stressed at this point that robust estimators cannot be simply reduced to estimators that “remove information from the sample”. Although they give small weights to specific observations according to their abnormal influence on the estimation or testing procedure (as it is the case for the TBS estimator), they actually consider the whole data set simultaneously. We would for instance agree that the median is an alternative estimator to the mean for the center of a (symmetric) distribution, but this estimator is definitely not seen as an estimator that “removes information from the sample”. However, the median is a robust estimator of location giving weights of 0 to all observations except the central one(s). Moreover, as can be seen from the two illustrative examples, robust estimation does not always mean a lower second moment compared with classical estimation. Robustness does not just “truncate” data, but is aimed at accurately revealing the underlying structure

of the model, often masked by outliers under classical estimation.

In the next section, a simulation study is performed to investigate the behavior of mean-variance optimal portfolios when classical and robust estimates of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are used in the presence of artificially contaminated data.

4 Simulation study

We here extend the example of the hedge fund portfolio by means of a simulation study. It should be stressed that the dimension of the problem (i.e. the number of asset classes under consideration) does not play a significant role in the performance of the robust optimizer. However, as the latter is built on robust estimators of mean and covariance that take into account the dimensionality of the data, only the computational speed would be affected by high dimension.

Let the population of returns be i.i.d. and generated by a multivariate normal distribution $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. To do so, we consider the robust TBS estimators ($\hat{\boldsymbol{\mu}}_{\text{TBS}}$ and $\hat{\boldsymbol{\Sigma}}_{\text{TBS}}$) computed on the *original* hedge fund portfolio data set as being the parameters of this new trivariate normal population of returns on the convertible arbitrage, event-driven and fixed income arbitrage strategies.¹⁶ We later refer to this new population as the *true* model. From this population, we simulate 20 samples

¹⁶ Taking as population parameters the classical $\hat{\boldsymbol{\mu}}_{\text{ML}}$ and $\hat{\boldsymbol{\Sigma}}_{\text{ML}}$ estimators is another option that leads to the same conclusion. However, by working with their robust counterparts we can highlight an interesting point related to model risk and estimation risk.

of size 120 representing 10 years of trivariate normal monthly returns.¹⁷ These 20 uncontaminated samples remain the same and are used throughout the following study.

As the aim of performing this simulation study is to assess the stability properties of the results of the MVEF model when data departs from normality, we work with uncontaminated and *contaminated* samples. The contaminated samples are constructed by applying transformation to the returns of the uncontaminated samples. Four types of return contaminations have been studied.

1. *Substitutive contamination*: random *replacement* of a given percentage of the 120 returns by a specific value. This value is the sum of the true mean of the series and of a positive or negative multiple of the true standard deviation of the series.
2. *Additive contamination*: random *addition* of a specific value to a given percentage of the 120 returns. This value is a positive or negative multiple of the true standard deviation of the series.
3. *Multiplicative contamination*: random *multiplication* of a given percentage of the 120 returns by a specific value.

¹⁷ 20 samples ensures clarity in subsequent graphics, as we found that a larger number of samples leads to similar results. The size of 120 has been chosen to reflect an average of what is commonly seen in practice. Owing to the lack of history on many indices or securities, it is indeed questionable whether portfolio optimization carried out in the financial reality could be based, on average, on much more data.

4. *Point mass multiplicative contamination*: random *multiplication* of a given percentage of the 120 returns by a specific value. The contamination occurs for each of the three series on the *same* data point(s).

The first three types of contamination concern each asset separately. This could correspond in practice to shocks in a specific sub-market not affecting all other sub-markets simultaneously. The fourth type of contamination, on the other hand, simulates simultaneous shocks in all markets. This fourth type of contamination tends to replicate shocks like the financial crisis in Asia in 1997 or in Russia in 1998.

The percentage of data contaminated and the multiplicative coefficients used in the four types of contamination above may vary from one series to another.¹⁸ It should be emphasized that these contaminations are carried over independently to each of the 3 series except the point mass multiplicative contamination case. We suggest contaminating 3% of the returns in each of the three return series, making the total contamination on each sample add up to 9%.

In what follows, classical estimation always refers to the use of the ML estimator, whereas robust estimation refers to the use of the TBS estimator (with breakdown point $\varepsilon^* = 0.35$ and asymptotic rejection probability $\alpha = 0.001$).¹⁹ Classical and

¹⁸ All above-mentioned contamination types are based on transformation of specific returns. Other types of contamination, including direct transformation on the parameters $\hat{\mu}$ and $\hat{\Sigma}$, have not been considered as they seem less relevant for replicating possible contamination in real market data.

¹⁹ A breakdown point of 35% and an asymptotic rejection probability of 0.1% strike a good balance between robustness and efficiency.

robust estimation of the parameters are performed to allow computation of the so-called classical and robust efficient frontiers (i.e. the frontiers calculated by means of the classically or robustly estimated parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$).

4.1 Behavior of efficient frontiers

First, let us compute the 20 classical and 20 robust efficient frontiers on the 20 normal uncontaminated simulated samples. The results are shown in Fig. 9.²⁰

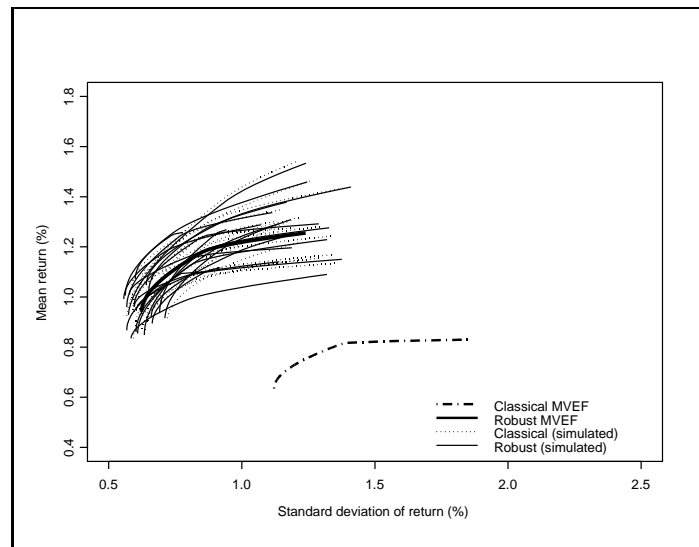


Fig. 9. Classical and robust efficient frontiers computed on the same 20 normal uncontaminated simulated samples. Short selling is not allowed. The size of each sample is 120. The classical and robust MVEFs computed on the original data set are reported as reference.

The classical and robust efficient frontiers computed on the original data set are the same as those in Fig. 7, and are reported as reference. As anticipated, the classical and robust simulation results are very similar, as no contamination is

²⁰ The efficient frontiers of Fig. 9 have been computed with the constraint of no short selling.

present in the data.²¹ The dispersion of both the 20 classical and 20 robust efficient frontiers indicates that even when the distributional normality assumption is not violated, sampling variability (even in series with rather low standard deviation of returns) is enough to make efficient frontiers move around the true one, making this situation an illustration of estimation risk. This implies that, for a given level of standard deviation of returns, the impact in terms of performance is far from being negligible when considering portfolios alternately located on each of the simulated efficient frontiers. However, this sampling variability is far smaller than the bias between classical and robust estimation on the original data represented by the classical and robust MVEFs in Fig. 9 .

Let us now turn to the results when classical and robust estimation is carried over to contaminated samples. The plot of Fig. 10 shows the specific results of the *substitutive contamination* type with independent contamination of 3% and a multiplicative coefficient of 5 applied to each series. As in Fig. 9, classical and robust efficient frontiers computed on the original data set are reported as reference.

The striking feature of Fig. 10 is that the 20 robust efficient frontiers are in the same position as in the previous (uncontaminated) case, while the 20 classical efficient frontiers have all clearly shifted higher and to the right. They no longer cluster around the true efficient frontier.

This type of contamination indeed implies a lateral shift due to variance increase

²¹ The population parameters being the robust mean and covariance estimates ($\hat{\mu}_{\text{TBS}}$ and $\hat{\Sigma}_{\text{TBS}}$), classical and robust efficient frontiers calculated on normal uncontaminated simulated samples cluster around the *true* robust mean-variance efficient frontier reported in bold.

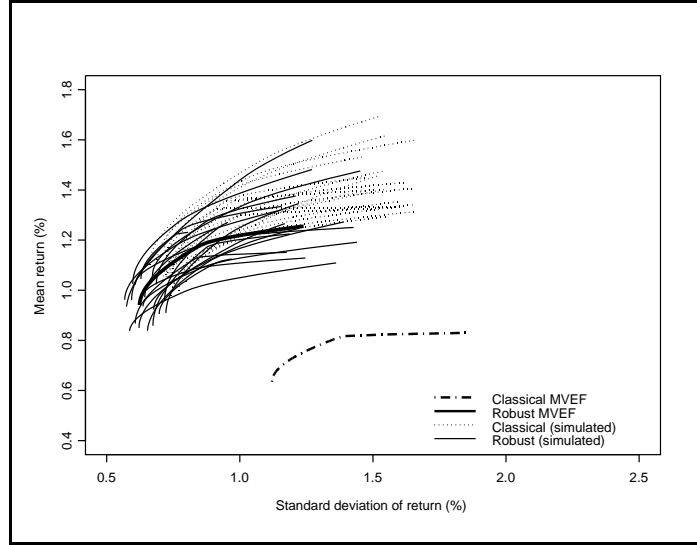


Fig. 10. Classical and robust efficient frontiers computed on the same 20 contaminated simulated samples. Short selling is not allowed. The size of each sample is 120. Substitutive contamination is applied with an independent contamination of 3% on each series and a multiplicative coefficient of 5. The classical and robust MVEFs computed on the original data set are reported as reference.

in the sample, and a vertical shift due to the positive multiplicative value of 5 increasing the probability of replacing returns by higher ones, as each of the three original return series exhibits a positive mean. Hence, for a fixed level of return, standard deviation is overestimated. This illustrates the possibly misleading results obtained when using classical estimation to compute efficient frontiers, as the position of the latter in Fig. 10 wrongly suggests a change in the true model. In fact, this bias is only caused by the bad influence of the 9% of outlying data on the estimation process.

Similar results are obtained with substitutive, additive and multiplicative contamination with the parameters mentioned above. When the multiplicative coefficient is negative, we found that efficient frontiers shift lower. And when a mix of positive

and negative coefficients is used, we found that the vertical shift almost disappears and only a horizontal shift remains. In all cases however, the horizontal shift occurs to the right as the variance of the data increases due to the contamination.

The above results have been obtained with independent contamination of 3% on each series and with a multiplicative coefficient of 5. However, we may argue that this type of contamination is too high or too low, or that the choice of the multiplicative coefficient is inadequate compared with data contamination encountered in financial reality. That's why we suggest looking again at the original data set and specifically at the plot of Fig. 8 showing the data influence measure for the hedge fund portfolio. As already mentioned, the data influence measure clearly identifies two (trivariate) returns as having a heavy influence on the estimates of the optimal portfolio weights, namely data points 56 and 57, corresponding to the Russian financial crisis of late summer 1998. These two data points represent 1.85% of the whole sample.

We compute the median of each of the three series of returns (i.e. convertible arbitrage, event-driven and fixed income arbitrage strategies) over the whole period.²² We also compute the median (in this case equal to the arithmetic mean) of each of the three series for these two specific outlying data points. The ratio of the medians (taken series by series) equals $(-4, -7, -3)$, for the convertible arbitrage, event-driven and fixed income arbitrage strategies respectively. This means for instance, in the case of event-driven strategies, that the median monthly return of the two specific outlying data points is seven times larger (in absolute terms) than the

²² The median is preferred to the mean for its robustness properties.

median monthly return on the whole period.²³

We are now able to artificially reproduce part of the contamination present in the original data set of the hedge fund portfolio, and we use the above-mentioned *point mass multiplicative contamination* to study the behavior of the efficient frontiers in that specific case. The contamination is set at 2.5% and the multiplicative coefficients are set at $(-4, -7, -3)$ for the respective strategies.²⁴ From now on, only this contaminated sample and the uncontaminated sample are used in the study.

The results displayed in Fig. 11 are disturbing: while, as in the previous (contaminated) case, the 20 robust efficient frontiers seem unaffected by contamination, the 20 classical efficient frontiers shift sharply lower and extend to the right, to the point that they even overlap the classical efficient frontier computed on the original data set.

In this case, when using classical estimation, model risk clearly dominates estimation risk, and the danger when making portfolio choices is of ending up on an efficient frontier strongly influenced by the characteristics of just a few outlying observations. Needless to say, portfolio choices made in this context may lead to sub-optimal decisions. On the other hand, robust efficient frontiers show no apparent bias, and thus more accurately represent the statistical properties and the structure of the non-outlying 97.5% of data of the 20 samples.

²³ This shows that the previously used value of 5 was (in absolute terms) a realistic choice for the multiplicative coefficient.

²⁴ The contamination of 2.5% is chosen to obtain exactly 3 contaminated data points among the 120 data points of each sample. Note that this contamination is smaller than that used before.




Fig. 11. Classical and robust efficient frontiers computed on the same 20 contaminated simulated samples. Short selling is not allowed. The size of each sample is 120. Point mass multiplicative contamination is applied with a contamination of 2.5% and respective multiplicative coefficients equal to $(-4, -7, -3)$. The classical and robust MVEFs computed on the original data set are reported as reference.

This case of point mass multiplicative contamination illustrates an interesting point, namely the non-continuous behavior of classical estimators in the presence of outlying data. Here, a mere 2.5% of outlying data has a far greater impact on the results than the previous 9% of contaminated data.

As shifts in efficient frontiers imply changes at the level of the underlying portfolios in terms of mean return, standard deviation and optimal weights of the different securities, a closer look at the portfolios themselves is also of interest.

4.2 Behavior of portfolios

Until now, we have considered portfolio optimization with a constraint of no short selling, for better visual identification of shifts of efficient frontiers, and to replicate

what is often done in practice. We suggest lifting this restriction to be fully in line with the results of section 3, derived in the unconstrained case, and where the solution for the optimal portfolio weights is an explicit expression. We now focus on the behavior of minimum variance portfolios, but we found similar results in the case of portfolios with a given level of standard deviation of returns. Fig. 12 shows the boxplots of standard deviation of returns for minimum variance portfolios.




Fig. 12. Boxplots of standard deviation of returns when classical and robust efficient frontiers are alternately computed on uncontaminated and contaminated samples. Short selling is allowed. The size of each sample is 120. Point mass multiplicative contamination is applied with a contamination of 2.5% and respective multiplicative coefficients equal to $(-4, -7, -3)$. Reading from the left, boxplots 1 and 2 show the variability in standard deviation of returns of the minimum variance portfolios when (classical) ML estimation, and (robust) TBS estimation is used on 20 uncontaminated samples. Boxplots 3 and 4 show the same information when estimation is carried over to 20 contaminated samples.

The black diamonds represent the standard deviation of returns of the *true* minimum variance portfolio located on the true robust MVEF. The two boxplots on the left show that the characteristics of the standard deviation of returns for minimum

variance portfolios resulting from 20 classical and 20 robust estimations on uncontaminated samples are very similar, with the exception of a little loss of efficiency for the robust estimation. On the contrary, under contamination, the 20 classical estimations of efficient frontiers lead to minimum variance portfolios with a large bias in standard deviation of returns (see boxplot 3), whereas the 20 robust estimations of efficient frontiers exhibit similar minimum variance portfolio characteristics (see boxplot 4) to those found in the uncontaminated case.²⁵ This sensitivity of results shows again the lack of robustness of the classical ML-estimated MVEF model. Just a few outlying data (here 2.5%) are enough to heavily bias the estimation and make the optimization process give misleading results.

Another way of looking at this problem is to focus on the *composition* of the minimum variance portfolios by examining the vector $\hat{\mathbf{p}}$ of estimated optimal portfolios weights. The boxplots of the minimum variance portfolio weights for the 3 hedge fund strategies under classical estimation are shown in Fig. 13.

While classical estimation on the 20 uncontaminated samples (from the left, boxplots 1, 3, and 5) leads to boxplots reflecting only sampling variability (i.e. estimation risk), classical estimation on the 20 contaminated samples (from the left, boxplots 2, 4 and 6) shows very significant variability (i.e. model risk in addition to estimation risk) of the optimal weights within these minimum variance portfolios. Moreover, in the case of event-driven (ED) and fixed income arbitrage (FIA) strate-

²⁵ We obtained similar results with the behavior of the *mean return* of minimum variance portfolios.




Fig. 13. Boxplots of the weights for minimum variance portfolios under classical estimation on uncontaminated and contaminated samples. Short selling is allowed. The size of each sample is 120. Point mass multiplicative contamination is applied with a contamination of 2.5% and respective multiplicative coefficients equal to $(-4, -7, -3)$. Reading from the left, boxplots 1 and 2 are each based on the 20 optimal weights of the convertible arbitrage (CA) strategy within the minimum variance portfolios located on classically estimated efficient frontiers computed on uncontaminated and contaminated samples. Boxplots 3 and 4 show the same kind of information for the event-driven (ED) strategy. Boxplots 5 and 6 do the same for the fixed income arbitrage (FIA) strategy.

gies, the central 50% of the boxplots does not even overlap with the true weights represented by the black diamonds. Portfolio weights are heavily biased, due to the extreme sensitivity of the model to only a few outlying data when estimated in a classical way. Once again, model risk dominates estimation risk. In the case of robust estimation, results are far more stable, as shown in Fig. 14.

Uncontaminated and contaminated boxplots look very similar and show the definitive advantage of using robust estimation in the case of the MVEF model. Even though both estimations may lead to portfolios behaving in a specific way according to given market conditions, the use of robust estimation makes the results of the




Fig. 14. Boxplots of the weights for minimum variance portfolios under robust estimation on uncontaminated and contaminated samples. Short selling is allowed. The size of each sample is 120. Point mass multiplicative contamination is applied with a contamination of 2.5% and respective multiplicative coefficients equal to $(-4, -7, -3)$. Reading from the left, boxplots 1 and 2 are each based on the 20 optimal weights of the convertible arbitrage (CA) strategy within the minimum variance portfolios located on robustly estimated efficient frontiers computed on uncontaminated and contaminated samples. Boxplots 3 and 4 show the same kind of information for the event-driven (ED) strategy. Boxplots 5 and 6 do the same for the fixed income arbitrage (FIA) strategy.

model less sensitive to a few outlying observations: model risk is kept under control. This feature is of particular importance when analyzing real data, since the presence (or absence) of just a small percentage of data from the sample should not have a decisive impact on optimal portfolio choices.

5 Conclusions

We investigated the properties of the maximum likelihood estimated mean-variance portfolio optimizer and found that this model is not robust to deviations from the

assumption of multivariate normality. We showed analytically that the influence function of the estimator of the optimal portfolio weights, when computed with the maximum likelihood estimator as is often the case, is unbounded, meaning that even a single outlier may take these weights beyond any predefined value.

We introduced the data influence measure as a powerful diagnostic tool for detecting the specific returns from the sample that bias estimation of the optimal portfolio weights. We also highlighted that outlying observations may be characterized otherwise than by extreme returns, making them difficult to find without using this diagnostic tool. Moreover, to address the problem of non-robustness of the classical maximum likelihood estimator, we suggested replacing it by the translated-biweight S -estimator. This estimator is robust to local departures from normality ensuring that the resulting mean-variance optimal portfolios truly reflect the statistical properties of the *majority* of the data, and that their characteristics do not suffer from the possibly abnormally strong influence of a few outliers.

Furthermore, the simulation study makes clear that the classically estimated mean-variance efficient frontier model suffers from model risk when data underlying its computation are not exactly generated by a multivariate normal distribution, and that model risk dominates estimation risk.

It is also expected that the robust portfolio will outperform the classical one if future returns are generated under the same process. This is simply because robustly estimated parameters will be closer to the true parameter values of the underlying generating process than their classical counterparts.

Finally, it should be stressed that although we have here considered the multivariate normal stochastic process as generating the (majority of the) independently and identically distributed data, the same concepts of statistical robustness can be applied to more sophisticated models.

As normality is the exception rather than the rule in financial reality, the use of robust statistics in quantitative portfolio management opens the way to fruitful research.

Acknowledgements

Cédric Perret-Gentil gratefully acknowledges Barbara Caffin from Union Bancaire Privée. Maria-Pia Victoria-Feser gratefully acknowledges the financial support of the Swiss National Science Foundation (grant 610-057883.99) and that of the NCCR FINRISK.

Appendix A

Proof of Proposition 1

We make use of the IF to describe the behavior of the asymptotic bias of $\hat{\mathbf{p}}^*$ in a neighborhood of the model. Let us first write $\hat{\mathbf{p}}^* = \hat{\mathbf{p}}^*(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}})$ as the estimated optimal portfolio weights given by Eq. (5) jointly with Eq. (6). Eq. (5) then becomes

$$\hat{\mathbf{p}}^*(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}) = \frac{1}{\lambda} \hat{\boldsymbol{\Sigma}}^{-1} \left[\hat{\boldsymbol{\mu}} - \frac{(\mathbf{e}'_N \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\mu}} - \lambda) \mathbf{e}_N}{\mathbf{e}'_N \hat{\boldsymbol{\Sigma}}^{-1} \mathbf{e}_N} \right], \quad (9)$$

and, under local nonparametric departures from the model as in Eq. (2), we have

$$\hat{\mathbf{p}}^*(\hat{\boldsymbol{\mu}}(F_\varepsilon), \hat{\boldsymbol{\Sigma}}(F_\varepsilon)) = \frac{1}{\lambda} \hat{\boldsymbol{\Sigma}}(F_\varepsilon)^{-1} \left[\hat{\boldsymbol{\mu}}(F_\varepsilon) - \frac{(\mathbf{e}'_N \hat{\boldsymbol{\Sigma}}(F_\varepsilon)^{-1} \hat{\boldsymbol{\mu}}(F_\varepsilon) - \lambda) \mathbf{e}_N}{\mathbf{e}'_N \hat{\boldsymbol{\Sigma}}(F_\varepsilon)^{-1} \mathbf{e}_N} \right]. \quad (10)$$

Taking the derivative of Eq. (10) with respect to ε in $\varepsilon = 0$ to derive the IF of $\hat{\mathbf{p}}^*$ for the mean-variance optimal portfolio for a given value of λ , we get

$$\begin{aligned} & \text{IF}(\mathbf{y}_t, \hat{\mathbf{p}}^*, F_\theta) \\ &= -\boldsymbol{\Sigma}^{-1} \text{IF}(\mathbf{y}_t, \hat{\boldsymbol{\Sigma}}, F_\theta) \mathbf{p}^* \\ &+ \frac{1}{\lambda} \boldsymbol{\Sigma}^{-1} \left\{ \text{IF}(\mathbf{y}_t, \hat{\boldsymbol{\mu}}, F_\theta) + \frac{[\mathbf{e}'_N \boldsymbol{\Sigma}^{-1} \text{IF}(\mathbf{y}_t, \hat{\boldsymbol{\Sigma}}, F_\theta) \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} - \mathbf{e}'_N \boldsymbol{\Sigma}^{-1} \text{IF}(\mathbf{y}_t, \hat{\boldsymbol{\mu}}, F_\theta)] \mathbf{e}_N}{\mathbf{e}'_N \boldsymbol{\Sigma}^{-1} \mathbf{e}_N} \right. \\ &\quad \left. - \frac{\mathbf{e}'_N \boldsymbol{\Sigma}^{-1} \text{IF}(\mathbf{y}_t, \hat{\boldsymbol{\Sigma}}, F_\theta) \boldsymbol{\Sigma}^{-1} \mathbf{e}_N (\mathbf{e}'_N \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} - \lambda) \mathbf{e}_N}{(\mathbf{e}'_N \boldsymbol{\Sigma}^{-1} \mathbf{e}_N)^2} \right\}. \quad (11) \end{aligned}$$

Hence, the IF of $\hat{\mathbf{p}}^*$ depends directly on the IFs of both $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$. This means that the estimator of the optimal portfolio weights directly inherits the stability properties of the estimators of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$.

Appendix B

Proof of Proposition 2

Starting from Eq. (11) and replacing the respective IFs with the expression of the IFs for the maximum likelihood estimator of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ (see Eq. (7) and Eq. (8)), we have

$$\begin{aligned} & \text{IF}(\mathbf{y}_t, \hat{\mathbf{p}}^*, F_\theta) \\ &= (\mathbf{I} - A(\mathbf{y}_t)\boldsymbol{\Sigma})\mathbf{p}^* \\ & \quad + \frac{1}{\lambda}\boldsymbol{\Sigma}^{-1}\left[\mathbf{y}_t - \boldsymbol{\mu} + \frac{1}{b}[c(\mathbf{y}_t) - d(\mathbf{y}_t)]\mathbf{e}_N - \frac{1}{b^2}[k(\mathbf{y}_t) - b](l - \lambda)\mathbf{e}_N\right], \quad (12) \end{aligned}$$

where

$$\begin{aligned} A(\mathbf{y}_t) &= \boldsymbol{\Sigma}^{-1}(\mathbf{y}_t - \boldsymbol{\mu})(\mathbf{y}_t - \boldsymbol{\mu})'\boldsymbol{\Sigma}^{-1} \\ b &= \mathbf{e}'_N\boldsymbol{\Sigma}^{-1}\mathbf{e}_N \\ c(\mathbf{y}_t) &= \mathbf{e}'_NA(\mathbf{y}_t)\boldsymbol{\mu} \\ d(\mathbf{y}_t) &= \mathbf{e}'_N\boldsymbol{\Sigma}^{-1}\mathbf{y}_t \\ k(\mathbf{y}_t) &= \mathbf{e}'_NA(\mathbf{y}_t)\mathbf{e}_N \\ l &= \mathbf{e}'_N\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}. \end{aligned}$$

The IF of $\hat{\mathbf{p}}^*$ under ML estimation is clearly unbounded in \mathbf{y}_t , which means that the asymptotic bias of the estimated optimal portfolio weights can become arbitrarily large under model deviation.

References

- Alexander, G.J., Resnick, B.G., 1985. More on estimation risk and simple rules for optimal portfolio selection. *Journal of Finance* 40, 125–133.
- Barry, C.B., 1974. Portfolio analysis under uncertain means, variances, and covariances. *Journal of Finance* 29, 515–522.
- Bawa, V.S., Brown, S.J., Klein, R.W., 1979. Estimation risk and optimal portfolio choice. North-Holland, Amsterdam.
- Black, F., Litterman, R., 1992. Global portfolio optimization. *Financial Analysts Journal* 48, 28–43.
- Britten-Jones, M., 1999. The sampling error in estimates of mean-variance efficient portfolio weights. *Journal of Finance* 54, 655–671.
- Cavadini, F., Sbuelz, A., Trojani, F., 2002. A simplified way of incorporating model risk, estimation risk and robustness in mean variance portfolio management. Working paper, University of Southern Switzerland.
- Chopra, V.K., Ziemba, W.T., 1993. The effect of errors in means, variances and covariances on optimal portfolio choice. *Journal of Portfolio Management* 19, 6–11.
- Hampel, F.R., 1968. Contribution to the theory of robust estimation. Ph.D. thesis, University of California, Berkeley, CA.
- Hampel, F.R., 1974. The influence curve and its role in robust estimation. *Journal of the American Statistical Association* 69, 383–393.

- Hampel, F.R., Ronchetti, E.M., Rousseeuw, P.J., Stahel, W.A., 1986. Robust statistics: The approach based on influence functions. John Wiley & Sons, New York, N.Y.
- Huber, P.J., 1964. Robust estimation of a location parameter. *Annals of Mathematical Statistics* 35, 73–101.
- Lhabitant, F.-S., 2002. Hedge funds, myths and limits. John Wiley & Sons, New York, N.Y.
- Markowitz, H.M., 1952. Portfolio selection. *Journal of Finance* 7, 77–91.
- Markowitz, H.M., 1959. Portfolio selection: Efficient diversification of investments. Yale University Press, New Haven, CT.
- Michaud, R.O., 1989. The Markowitz optimization enigma: Is ‘optimized’ optimal? *Financial Analysts Journal* 45, 31–42.
- Rocke, D.M., 1996. Robustness properties of S -estimators of multivariate location and shape in high dimension. *The Annals of Statistics* 24, 1327–1345.
- Rousseeuw, P.J., Leroy, A.M., 1987. Robust regression and outlier detection. John Wiley & Sons, New York, N.Y.
- ter Horst, J.R., de Roon F.A., Werker, B.J., 2002. Incorporating estimation risk in portfolio choice. Working paper, Tilburg University, The Netherlands.
- Victoria-Feser, M.-P., 2000. Robust portfolio selection. Cahiers de recherche HEC 2000.14, University of Geneva, Switzerland.