

10

Incorporating Estimation Errors into Portfolio Selection: Robust Portfolio Construction

*Sebastián Ceria**

is President and CEO of Axioma, Inc., a New York-based company that produces portfolio construction analytics. Before founding Axioma, he was an associate professor of decision, risk and operations at Columbia Business School from 1993 to 1998. He completed his PhD in Operations Research at Carnegie Mellon University's Graduate School of Industrial Administration.

Robert A. Stubbs

is Director of Research and Development, for Axioma, Inc., where he develops specialised optimisation solution techniques for financial problems and improved portfolio management methodologies and software. He earned a PhD in Industrial Engineering and Management Science and a Master of Science degree in Industrial Engineering and Management Science from Northwestern University, and a Bachelor of Science degree in Applied Mathematics from Auburn University.

**Axioma, Inc., 100 5th Avenue, Suite 901, New York, NY 10011, USA*

Tel: +1 (212) 901 1000, Fax: +1 (212) 901 1911; e-mail: sceria@axiomainc.com

Introduction

More than 50 years have elapsed since Markowitz (1952) first introduced his Nobel Prize-winning work on mean-variance portfolio optimisation. His work led to the creation of the field now known as Modern Portfolio Theory (MPT). Throughout this time, MPT has had many followers but has also been challenged by sceptics at academic and financial institutions alike. Today, even though MPT is still widely accepted as the primary theoretical framework for portfolio construction, its employment by investment professionals is not as ubiquitous as one might expect. There are several reasons for the lack of acceptance

Reprinted from 'Incorporating estimation errors into portfolio selection: Robust portfolio construction' by Sebastián Ceria and Robert A. Stubbs in *Journal of Asset Management*, 7, 2006, pp. 109–127, DOI:10.1057/palgrave.jam.2240207. With kind permission from Palgrave Macmillan Ltd. All rights reserved.

of MPT among practitioners, but perhaps the most significant is the argument that 'optimal' portfolios obtained through the mean-variance approach are often 'counterintuitive', 'inexplicable' and 'overly sensitive to the input parameters'.

The fact that mean variance 'optimal' portfolios are sensitive to small changes in input data is well documented in the literature. Chopra (1993) shows that even slight changes to the estimates of expected returns or risk can produce vastly different mean-variance optimised portfolios. Best and Grauer (1991) analyse the sensitivity of optimal portfolios to changes in expected return estimates. Instead of focusing on the weights of the assets in optimal portfolios, others have focused on the financial impact of mean-variance efficient portfolios computed from estimates. Jobson and Korkie (1981) show that even an equal-weighted portfolio can have a greater Sharpe ratio than an optimal mean-variance portfolio computed using estimated inputs. Broadie (1993) shows how the estimated efficient frontier overestimates the expected returns of portfolios for varying levels of estimation errors. Because of the ill-effects of estimation errors on optimal portfolios, portfolio optimisation has been called 'error maximisation' (See Michaud, 1989). Michaud argues that mean-variance optimisation overweights those assets with a large estimated return to estimated variance ratio (under weights those with a low ratio) and that these are precisely the assets likely to have large estimation errors.

It is widely believed that most of the estimation risk in optimal portfolios is due to errors in estimates of expected returns, and not in the estimates of risk. Chopra and Ziemba (1993) argue that cash-equivalent losses due to errors in estimates of expected returns are an order of magnitude greater than those for errors in estimates of variances or covariances. Many portfolio managers concur, saying that their confidence in risk estimates is much greater than their confidence in expected return estimates.

In order to cope with the effect of estimation errors in the estimates of expected returns, attempts have been made to create better and more stable mean-variance optimal portfolios by using expected return estimators that have a better behaviour when used in the context of the mean-variance framework. One of the more common techniques is the use of James-Stein estimators (see Jobson and Korkie, 1981). These estimators shrink the expected returns towards the average expected return based on the volatility of the asset and the distance of its expected return from the average. Jorion (1985) developed a similar technique that shrinks the expected return estimate towards the minimum variance portfolio. More recently, Black and Litterman (1990)

have developed a new Bayesian approach for producing stable expected return estimates that combines equilibrium expected returns and investors' views on specific assets or weighted groups of assets. The area of robust statistics (see Cavadini *et al.*, 2002) has recently been employed to create stable expected return estimates as well.

While strides have been made to improve estimates of expected returns, there will always be errors in these estimates because of the inherent stochastic nature of the asset return process. Even a portfolio manager employing Bayesian estimators such as James–Stein or Black–Litterman will admit that estimation error remains a factor in the 'corrected' estimates of expected returns, even if it is significantly less than that obtained without the use of these methods. In fact, the main premise of Bayesian statistics is that estimates do have distributions. This has led some authors to consider ways in which to account for estimation errors directly in the portfolio construction process.

One possible strategy for considering estimation error is to increase the risk-aversion parameter or modify the risk estimates by increasing the overall volatility. Since the estimated efficient frontier is an overestimate of the true efficient frontier because of the error-maximisation property, it can be argued that by increasing the risk-aversion parameter, the resulting portfolio on the actual frontier will be closer to the true frontier. Horst *et al.* (2001) show how to create an optimal pseudo risk-aversion parameter to use in a mean-variance optimisation problem rather than using the actual risk-aversion parameter. One problem with this approach is that it assumes that the covariance matrix of the estimation error is a constant multiple of the covariance matrix of returns, which is rarely the case in practice. Since expected return estimates are typically generated independent of the factor risk model, the distribution of the estimation error is likely to be quite different from that of the risk model.

Another development that has recently received much attention is the portfolio resampling methodology of Michaud (1999). Michaud introduces a statistical resampling technique that indirectly considers estimation error by averaging the individual optimal portfolios that result from optimising with respect to many randomly generated expected-return and risk estimates. Portfolio resampling, however, is a somewhat *ad hoc* methodology that has many pitfalls (see Scherer, 2002). Because portfolio resampling is a simulation procedure in which each iteration involves a resampling of a time-series, creating mean-variance input estimators and determining the optimal portfolio, it is overly time-consuming to compute. Like the modified risk-aversion parameter approach, portfolio resampling does not actually consider

the portfolio manager's estimation error. It only considers the error of estimating a mean and covariance matrix from a simulated time-series from a stationary return process using the expected returns and covariance matrix to generate the time-series. Additionally, the resulting optimal portfolio does not necessarily satisfy all constraints. If a non-convex constraint, eg a limit on the number of assets in the portfolio, is present, the average of portfolios that satisfy the constraints individually will not necessarily satisfy the constraints.

Others have proposed adding constraints to control the ill-effects of estimation error on optimisation generated portfolios. While some constraints can reduce the sensitivity of optimal portfolios to changes in inputs, this paper shows that constraints can actually exacerbate the problem. Furthermore, it shows that the overestimate of expected return of an optimal portfolio can also be exacerbated by the presence of constraints.

This paper discusses an optimisation methodology known as robust optimisation, which considers uncertainty in unknown parameters directly and explicitly in the optimisation problem. It is generally concerned with ensuring that decisions are 'adequate' even if estimates of the input parameters are incorrect. Robust optimisation was introduced by Ben-Tal and Nemirovski (1997) for robust truss topology design. In a paper that describes several applications of robust optimisation, Lobo *et al.* (1998) introduced the concept of considering the distribution of estimation errors of expected returns explicitly in a portfolio optimization problem. Since then, Goldfarb and Iyengar (2003) consider uncertainties in the factor exposure matrix of a factor risk model directly in the portfolio optimisation problem.

Robust portfolio optimisation is a fundamentally different way of handling estimation error in the portfolio construction process. Unlike the previously mentioned approaches, robust optimisation considers the estimation error directly in the optimisation problem itself. Here, a financial motivation is given for using robust portfolio optimisation as a means of considering errors in the expected return estimates directly in the portfolio construction process. This motivation allows one to see that the 'standard' formulation is only applicable in certain cases. The fourth section introduces modified forms of robust mean-variance optimisation that are applicable in other commonly used portfolio management strategies. Many of the results in this paper were first presented at a practitioners conference in April of 2003 by Ceria (see Ceria and Stubbs, 2003). Since then, a number of other authors have independently proposed an approach similar to the present authors'. Of particular relevance is the paper of Garlappi *et al.* (2004).

This paper introduces an optimisation methodology that significantly reduces some of the ill-effects of portfolio optimisation that are caused by estimation error in expected return estimates. It shows that errors in expected return estimates can lead to optimal portfolios whose weights are significantly different from those in the true optimal portfolio and whose expected return is significantly overestimated. It shows that this can be particularly true in the presence of commonly found types of constraints. We discuss a 'standard' robust optimisation methodology that alleviates some of these ill-effects and introduces new variants that more effectively handle the difficulties caused by estimation error in commonly used portfolio management strategies. Finally, some computational results are discussed that demonstrate the potentially significant economic benefits of investing in portfolios computed using standard robust optimisation and the variants introduced here.

Estimation errors and classical mean-variance optimisation

It is a well-documented fact in the investment management literature that mean-variance optimisers are very sensitive to small variations in expected returns. Slightly different expected return vectors can lead to drastically different portfolios. The seemingly unexplainable changes in asset weights due to small perturbations in expected returns are not the only pitfall of classical mean-variance optimisation. Because of the error-maximisation effect, it is typically the case that the expected return is significantly overestimated.

In order to understand better the effect of estimation error in expected returns on optimal portfolios, consider the following example. Suppose there are two assets where the objective is to maximise expected return subject to a budget constraint that forces full investment between the two assets, and a constraint that limits the total active risk to be no more than 10 per cent with respect to the benchmark portfolio (shown as point 'M' in Figure 10.1). The estimates of expected returns and standard deviations of the two assets are given in Table 10.1. It is assumed that the correlation between the two assets is 0.7. The feasible region of this example is illustrated in Figure 10.1 as the intersection of the shaded ellipsoidal region and the budget constraint, ie the feasible region of this example is simply the line segment between points A and B.

Using column 'Alpha 1' from Table 10.1 as the estimates of expected returns, the optimal portfolio is at point A in Figure 10.1. Using the slightly different expected returns given in column 'Alpha 2', the

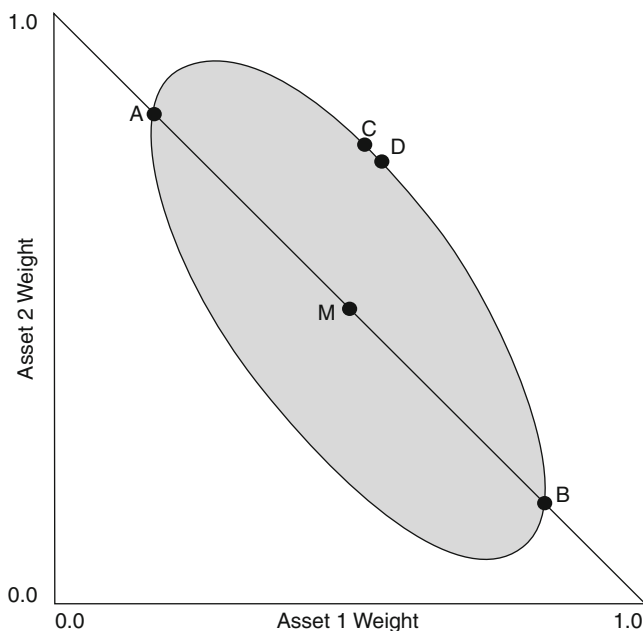


Figure 10.1 Feasible region for example 1

Table 10.1 Expected returns and standard deviations for example 1

Asset	Benchmark weights	Alpha 1 (%)	Alpha 2 (%)	True Alpha (%)	Return Std. Dev. (%)	Alpha Std. Dev. (%)
Asset 1	0.5	2.4	2.5	2.48	0.42	0.5
Asset 2	0.5	2.5	2.4	2.42	0.33	0.5

optimal portfolio is at point B. (The values of the portfolio weights are given in Table 10.2.) This example shows that with only a very small change in the estimates of expected returns of the assets, the weights of the assets in the optimal portfolios changed dramatically. The true optimal solution is at point B with an expected return of 2.46986 per cent. The estimated expected return of points A and B using 'Alpha 1' and 'Alpha 2', respectively, are both 2.4831 per cent. In this example, the expected returns of both optimal portfolios evaluated with respect to their expected return estimates overestimate the true expected return.

Table 10.2 Optimal portfolios for example 1

Attribute	Folio A	Folio B	Folio C	Folio D
Alpha	1	2	1	2
Budget	✓	✓		
Asset 1 weight	0.169	0.831	0.5253	0.5546
Asset 2 weight	0.831	0.169	0.7796	0.7503

From this example, it is clear why portfolio managers find ‘optimised’ portfolios to be counter intuitive and impractical.

Some changes in optimal weights should be expected when using different estimates of expected returns. Most of the variations in asset weights, however, arise due to optimisers exacerbating the estimation error problem by significantly overweighting assets with an error to the upside and underweighting assets with an error to the downside. Though this behaviour has been described before, the authors are not aware of any studies that have given a precise and intuitive explanation of the error exacerbating effect. The present authors claim that the cause of the ‘error maximisation’ property of mean-variance optimisers is not only the presence of estimation error, but also the interaction of the estimation error in expected returns with the constraints present in the portfolio optimisation problem.

If we reconsider our example, but drop the budget constraint, then the optimal portfolios with respect to ‘Alpha 1’ and ‘Alpha 2’ are points C and D, respectively. Figure 10.1 and Table 10.2 show that the change in optimal portfolio weights with respect to their expected return estimates is much smaller when the budget constraint is dropped. The true optimal solution is at a point on the boundary of the ellipsoid between points C and D with an expected return of 3.191258 per cent. The estimated expected return of points C and D using ‘Alpha 1’ and ‘Alpha 2’, respectively, are 3.20972 per cent and 3.18722 per cent. In this example, the expected return of portfolio C evaluated with respect to its expected return estimate overestimates the true expected return, but portfolio D evaluated with respect to its expected return estimate actually underestimates the true expected return. This situation is extremely rare when constraints are present, particularly in higher dimensions.

This simple example was created to illustrate geometrically how slightly different expected return estimates can lead to very different portfolios and how this phenomenon can be exacerbated by the

introduction of constraints. It also shows how the error in expected returns is optimised so that the estimated expected return of a portfolio typically overestimates the true expected return. In this small example, the change in expected returns of the portfolios was small, but this was only a two-asset example. To illustrate the error-maximisation effect better, efficient frontiers are considered in a more realistic investment scenario.

As defined by Broadie (1993), the terms ‘true frontier’, ‘estimated frontier’, and ‘actual frontier’ are used to refer to the efficient frontiers computed using the true expected returns (unobservable), estimated expected returns and true expected returns of the portfolios on the estimated frontier, respectively. Specifically, the frontier computed using the true, but unknown, expected returns is referred to as the true frontier. Similarly, the frontier computed using estimates of the expected returns and the true covariance matrix is referred to as the estimated frontier. Finally, the actual frontier is defined as follows. We take the portfolios on the estimated frontier and then calculate their expected returns using the true expected returns. Since we are using the true covariance matrix, the variance of a portfolio on the estimated frontier is the same as the variance on the actual frontier. By definition, the actual frontier will always lie below the true frontier. The estimated frontier can lie anywhere with respect to the other frontiers. If the errors in the expected return estimates have a mean of zero, however, the estimated frontier will lie above the true frontier with extremely high probability, particularly when the investment universe is large.

Using the covariance matrix and expected return vector from Idzorek (2002), we randomly generated a time-series of normally distributed returns and computed the average to use as estimates of expected returns. Using this computed expected-return estimate and the true covariance matrix, we generated an estimated efficient frontier of active risk versus active return where the portfolios were subject to no-shorting constraints and a budget constraint that forces the sum of the weights to be one. Similarly, the true efficient frontier was generated using the original covariance matrix and expected return vector. Finally, the actual ‘frontier’ was generated by computing the expected return and risk of the portfolios on the estimated frontier with the true covariance and expected return values. The actual ‘frontier’ is not necessarily concave, as it is not computed as the result of any optimisation, but rather by applying the true expected returns and true covariance to the efficient portfolios in the estimated efficiency frontier. These three frontiers are illustrated in Figure 10.2. Using the same estimate of expected

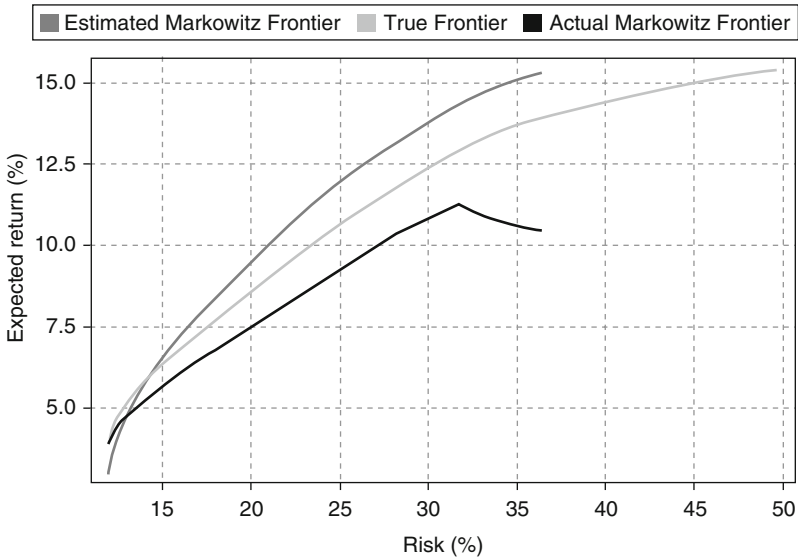


Figure 10.2 Markowitz Efficient Frontiers

returns, we also generated active risk versus active return where we also constrained the active holdings of the assets to be ± 3 per cent of the benchmark holding of each asset. These frontiers are illustrated in Figure 10.3. Note how the estimated frontiers significantly over estimate the expected return for most risk levels in both types of frontiers. More importantly, note that the actual frontier lies far below the true frontier in both cases.

This shows that the ‘optimal’ mean-variance portfolio is not necessarily a good portfolio, ie it is not ‘mean-variance efficient’.

In general, it is not known how far the actual expected return may be from the expected return of the mean-variance optimal portfolio. Returning to the example, suppose that the true expected return estimate is some convex combination¹ of the expected return estimates, $(\alpha_1, \alpha_2) = (2.5, 2.4)$ and $(\alpha_1, \alpha_2) = (2.2, 2.7)$ and that one value is no more likely to occur than another. Depending on the point estimate of expected returns used in the mean-variance optimisation problem, the optimal portfolio will be either portfolio A or portfolio B. The actual expected returns of these portfolios for the two extreme expected return estimates are given in Table 10.3. Suppose that the estimate of

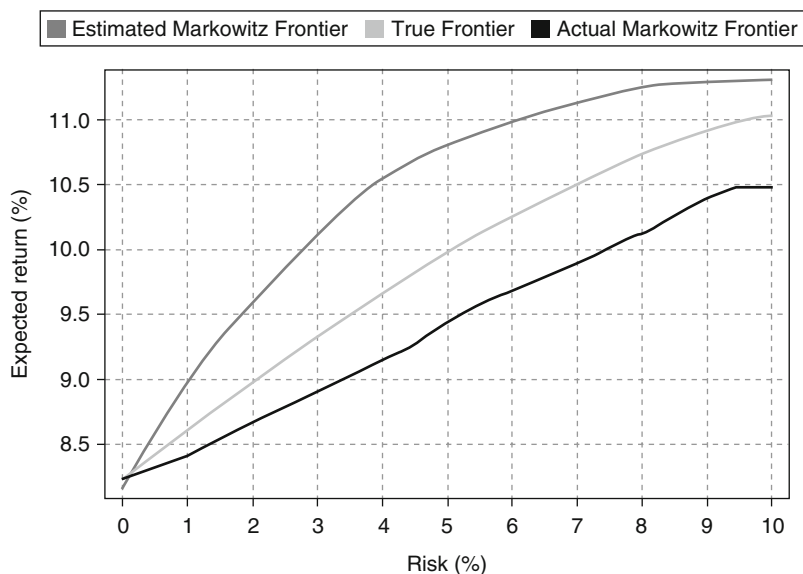


Figure 10.3 Markowitz benchmark-relative efficient frontiers

Table 10.3 Extreme expected return for optimal portfolios

Portfolio	Expected return (2.5, 2.4) (%)	Expected return (2.2, 2.7) (%)
A	2.4169	2.6155
B	2.4831	2.2845

the expected returns leads to optimal portfolio B. In the best scenario, portfolio B will have an expected return that is 0.0662 percentage points greater than that of portfolio A. In the worst-case, however, portfolio B will have an expected return that is 0.331 percentage points less than that of portfolio A. So, by investing in a portfolio that may have an expected return that is 0.0662 percentages points greater than an alternative, there is the risk that the expected return may be as much as 0.331 percentage points less. Since a uniform distribution of expected returns between the two extreme values is assumed, it could be argued that portfolio A is a better, more robust portfolio. That is, the portfolio performs better under more situations within the range of uncertainty of expected returns.

Robust portfolio optimisation

The actual frontier resulting from classical mean-variance optimisation can be far away from both the true and estimated frontiers because of estimation error. The estimated frontier generally lies well above the actual frontier. This study will now analyse just how far apart the estimated and actual frontiers can be over a specified confidence region of the true expected return. It is assumed that the n -dimensional vector of true expected returns α , is normally distributed. Given an estimate of expected return $\bar{\alpha}$ and a covariance matrix Σ of the estimates of expected returns,² it is assumed the true expected returns lies inside the confidence region

$$(\alpha - \bar{\alpha})^T \Sigma^{-1} (\alpha - \bar{\alpha}) \leq \kappa^2 \quad (1)$$

with probability 100η per cent where $\kappa^2 = \chi_n^2(1 - \eta)$ and χ_n^2 is the inverse cumulative distribution function of the chi-squared distribution with n degrees of freedom.³

If the covariance matrix of returns Q is full rank, then one can compute points on the efficient frontier by solving the maximum expected return problem,

$$\begin{aligned} &\text{maximise} && \bar{\alpha}^T w \\ &\text{subject to} && w^T Q w \leq \nu \end{aligned} \quad (2)$$

for varying values of ν , where α is the expected return estimate, Q is the covariance matrix of returns, and ν is the target portfolio variance. It is easy to show that the optimal holdings to this maximum expected return problem are given by

$$w = \sqrt{\frac{\nu}{\bar{\alpha}^T Q^{-1} \bar{\alpha}}} Q^{-1} \bar{\alpha}$$

Let α^* be the true, but unknown, expected return vector and $\bar{\alpha}$ be an expected return estimate. Recall that the actual frontier is constructed using the true expected return α^* . That is, the true expected return of a portfolio on the estimated from tier is computed as

$$\sqrt{\frac{\nu}{\bar{\alpha}^T Q^{-1} \bar{\alpha}}} \alpha^{*T} Q^{-1} \bar{\alpha}$$

Let \tilde{w} be the optimal portfolio on the estimated frontier for a given target risk level. Generally, the estimated expected return is greater than the actual expected return because of the ‘error maximisation’ effect of the optimiser. The question addressed is ‘How great can the difference be?’ To answer this, consider the maximum difference between the estimated expected return and the actual expected return of \tilde{w} . This difference can be written as

$$\bar{\alpha}^T \tilde{w} - \alpha^{*T} \tilde{w}$$

For a 100η per cent-confidence region of α , the maximum difference between the expected returns on the estimated efficient frontier and the actual efficient frontier is computed by solving

$$\begin{aligned} &\text{maximise} && \bar{\alpha}^T \tilde{w} - \alpha^T \tilde{w} \\ &\text{subject to} && (\alpha - \bar{\alpha})^T \Sigma^{-1} (\alpha - \bar{\alpha}) \leq \kappa^2. \end{aligned} \tag{3}$$

Note that \tilde{w} is fixed in problem (3). We are optimising over the variable α . The optimal solution to (3) can be shown to be

$$\alpha = \bar{\alpha} - \sqrt{\frac{\kappa^2}{\tilde{w}^T \Sigma \tilde{w}}} \Sigma \tilde{w} \tag{4}$$

Therefore, the lowest possible value of the actual expected return of the portfolio over the given confidence region of true expected returns is computed as

$$\alpha^T \tilde{w} = \bar{\alpha}^T \tilde{w} - \kappa \|\Sigma^{1/2} \tilde{w}\| \tag{5}$$

and the maximum difference between the estimated frontier and the actual frontier is

$$\bar{\alpha}^T \tilde{w} - (\bar{\alpha}^T \tilde{w} - \kappa \|\Sigma^{1/2} \tilde{w}\|) = \kappa \|\Sigma^{1/2} \tilde{w}\| \tag{6}$$

(Throughout this paper, $\|\cdot\|$ refers to the 2-norm.)

Naturally, one would like this difference to be as small as possible. This would reduce the error-maximisation effect, bring the estimated and actual frontiers closer together, and thus create portfolios that are closer to the true efficient frontier. Simply minimising the distance between the two frontiers, however, will drive the optimal portfolio towards a portfolio that minimises the estimation risk. Clearly, this is

not what we want to do. There is no point in considering estimation error if one does not consider the estimates. Instead, we simultaneously want to continue to maximise the expected return of the portfolio so that we are minimising the estimation risk for a given level of estimated expected return. In order to do this, we solve an optimisation problem where we maximise an objective of the form of (5). With this optimisation problem, we will bring the actual and estimated frontiers closer together. We will not be able to guarantee, however, that these frontiers are actually closer to the true frontier.

In this problem, w , the vector of optimal holdings, is not fixed. We optimise over w to find the optimal asset weights. Additionally, any set of portfolio constraints can be added. For instance, a long-only robust portfolio satisfying a budget constraint and a variance constraint can be written as

$$\begin{aligned} & \text{maximise} && \bar{\alpha}^T w - \kappa \|\Sigma^{1/2} w\| \\ & \text{subject to} && e^T w = 1 \\ & && w^T Q w \leq \nu \\ & && w \geq 0 \end{aligned} \tag{7}$$

where ν is a variance target. Note that this problem is exactly the same as a classical mean-variance optimisation problem except for the $\kappa \|\Sigma^{1/2} w\|$ term in the objective. This term is related to the estimation error and its inclusion in the objective function reduces the effect of estimation error on the optimal portfolio.

There is an important distinction between Q and Σ . Q is the covariance matrix of returns, while Σ is the covariance matrix of estimated expected returns, which is related to the estimation error arising from the process of estimating α , the vector of expected returns. This distinction is even more relevant in practice, where typically Q is obtained from a risk model provider, and is completely independent from Σ which is the result of a proprietary estimation process for α of which the risk model provider is not even aware.

Let us consider just how this additional objective term affects an optimal solution. If one considers equation (4), it can be seen that the expected returns of those assets with positive weights will be effectively⁴ adjusted downwards.⁵ Similarly, the expected returns of those assets with negative weights, ie short holdings, will be adjusted upwards. The size of the adjustment is controlled by the size of κ , ie the size of the confidence region. Note that the alpha correction term in Equation (4)

is constant κ multiplied by the marginal contribution to estimation risk of the assets. Therefore, for a portfolio with a lot of estimation risk in a single asset, the expected return for that asset is effectively adjusted so as to reduce the marginal contribution to the estimation risk of that asset. The purpose of adjusting the expected returns estimates in this way is to counter the error-maximisation effects of portfolio optimisation.

We have just described what we refer to as a robust objective problem. The other forms of classical mean-variance optimisation can also be modelled using robust optimisation. For instance, the maximum utility form of the problem can be written as

$$\begin{aligned} \text{maximise} \quad & \bar{\alpha}^T w - \kappa \|\Sigma^{1/2} w\| \\ & - p / 2w^T Q w \end{aligned} \tag{8}$$

Similarly, the minimum volatility form of the problem can be written as

$$\begin{aligned} \text{maximise} \quad & w^T Q w \\ \text{subject to} \quad & \bar{\alpha}^T w - \kappa \|\Sigma^{1/2} w\| \geq r \end{aligned} \tag{9}$$

Problem (7) and its variants cannot be solved by a standard mean-variance optimiser or even a general-purpose quadratic optimiser because the estimation-error term is a 2-norm which contains a square root and cannot be reformulated as a pure quadratic problem. This robust optimisation problem must be solved by either an optimiser capable of handling general convex expressions or a symmetric second-order cone optimiser. Second-order cone optimisation is a relatively new branch of optimisation and special purpose optimisers have been created to solve problems of this type. These specialised solvers can optimize robust optimisation problems in roughly the same amount of time that a mean-variance optimiser can solve the classical problem.

Alternative forms of robust portfolio optimisation

The robust optimisation problem introduced in the previous section three will only adjust the estimates of expected returns downwards if long-only constraints are present. Assuming that each expected return estimate overestimates the true expected return and adjusting all estimates downwards is too pessimistic. Even though there are errors in an expected return estimate, it is not likely that the expected return

estimate of each asset is an overestimate of the actual expected return. Similarly, in management of an active fund, the expected returns will be adjusted downwards for any asset with a positive weight. This really does not make sense because this study is interested in active returns. One would not expect our alpha to be adjusted downwards for an asset that already has a negative active weight.

This section introduces new variants of robust optimisation that deal with these issues. It should be noted that the two variants introduced here do not necessarily cover all real-world situations. These variants, along with the standard formulation, do provide ways of handling most commonly found portfolio management strategies. At the end of this section, a more general framework is described under which to view these alternative forms of robust optimisation. This framework can be used to develop other extensions for applicable circumstances.

Zero net alpha-adjustment frontiers

The standard robust optimisation problem discussed in the previous section considered the maximum possible difference between the estimated frontier and the actual frontier. This maximum difference was then minimised. Depending on the goals of the portfolio manager, this approach can potentially be too conservative as the net adjustment to the estimated expected return of a portfolio will always be downwards. If the manager's expected returns are symmetrically distributed around the point estimate, however, one would expect that there are approximately as many expected returns above their estimated values as there are below the true values. It may be more natural and less conservative to build this expectation into the model.

In order to incorporate a zero net alpha-adjustment into the robust problem, (6) is modified by the addition of the linear constraint

$$e^T D(\alpha - \bar{\alpha}) = 0 \tag{10}$$

for some symmetric invertible matrix D to obtain the following

$$\begin{aligned} & \text{maximise} && \bar{\alpha}^T \tilde{w} - \alpha^T \tilde{w} \\ & \text{subject to} && (\alpha - \bar{\alpha})^T \Sigma^{-1} (\alpha - \bar{\alpha}) \leq \kappa^2 \\ & && e^T D(\alpha - \bar{\alpha}) = 0 \end{aligned} \tag{11}$$

For now, assume that $D = I$, in which case (10) forces the total net adjustment to the expected returns to be zero. That is, for every basis point decrease in an expected return of an asset, there must be a

corresponding gross basis point increase in the expected return of other assets.

$$\alpha = \bar{\alpha} - \frac{\sqrt{\frac{\kappa^2}{\left(\Sigma\tilde{w} - \frac{e^T D \Sigma \tilde{w}}{e^T D \Sigma D^T e} \Sigma D^T e\right)^T \Sigma^{-1} \left(\Sigma\tilde{w} - \frac{e^T D \Sigma \tilde{w}}{e^T D \Sigma D^T e} \Sigma D^T e\right)}}}{\left(\Sigma\tilde{w} - \frac{e^T D \Sigma \tilde{w}}{e^T D \Sigma D^T e} \Sigma D^T e\right)^T \Sigma^{-1} \left(\Sigma\tilde{w} - \frac{e^T D \Sigma \tilde{w}}{e^T D \Sigma D^T e} \Sigma D^T e\right)} \left(\Sigma\tilde{w} - \frac{e^T D \Sigma \tilde{w}}{e^T D \Sigma D^T e} \Sigma D^T e \right) \quad (12)$$

It can be shown that the optimal solution to problem (11) is, (see Equation 12 above).

Therefore,

$$\alpha^T \tilde{w} = \bar{\alpha}^T \tilde{w} - \kappa \left\| \left\| \Sigma - \frac{1}{e^T D \Sigma D^T e} \Sigma D^T e e^T D \Sigma \right\|^{1/2} \tilde{w} \right\| \quad (13)$$

Instead of having a zero net adjustment of the alphas, one could restrict the alpha region to have a zero net adjustment in standard deviations of the alphas. To do this, one sets $D = L^{-1}$, where $\Sigma = LL^T$ is the Cholesky decomposition of Σ . This forces every standard deviation of upward adjustment in the alphas to be offset by an equal downward adjustment of one standard deviation. Similarly, we could restrict the alpha region to have a zero net adjustment in the variance of alphas in which case one sets $D = \Sigma^{-1}$.

Now, let us consider how this objective is effectively adjusting alphas when $D = \Sigma^{-1}$. In this case, the adjustment term becomes, (see Equation 14 below).

For a problem with a dollar-neutral constraint, ie $e^T \tilde{w} = 0$, the zero-net alpha adjustment form of robust optimisation is equivalent to the standard form. In a fully invested problem, however, there

$$\alpha = \bar{\alpha} - \frac{\sqrt{\frac{\kappa^2}{\left(\tilde{w} - \frac{e^T \tilde{w}}{e^T \Sigma^{-1} e} \Sigma^{-1} e\right)^T \Sigma \left(\tilde{w} - \frac{e^T \tilde{w}}{e^T \Sigma^{-1} e} \Sigma^{-1} e\right)}}}{\left(\tilde{w} - \frac{e^T \tilde{w}}{e^T \Sigma^{-1} e} \Sigma^{-1} e\right)^T \Sigma \left(\tilde{w} - \frac{e^T \tilde{w}}{e^T \Sigma^{-1} e} \Sigma^{-1} e\right)} \Sigma \left(\tilde{w} - \frac{e^T \tilde{w}}{e^T \Sigma^{-1} e} \Sigma^{-1} e \right) \quad (14)$$

will be a budget constraint of the form $e^T \tilde{w} = 1$. In this case, the term

$$\frac{e^T \tilde{w}}{e^T \Sigma^{-1} e} \Sigma^{-1} e$$

is exactly the portfolio that minimises estimation error subject to being fully invested. In this case, if a portfolio weight is above that which minimises estimation error, the effective alpha is adjusted downwards. Similarly, if the weight of an asset is below that which minimises estimation error, the effective alpha is adjusted upwards.

Robust active return/active risk frontiers

Thus far, this paper has discussed the classical efficient frontier that demonstrates the trade off between the expected values of total return and total risk. Active managers are more interested in an efficient frontier comparing the expected values of active return and active risk. For a 100η per cent-confidence region of α , the most that the difference between the expected active returns on the estimated efficient frontier and the actual frontier can be is computed by

$$\begin{aligned} &\text{maximise} && \bar{\alpha}^T(\tilde{w} - b) - \alpha^T(\tilde{w} - b) \\ &\text{subject to} && (\alpha - \bar{\alpha})^T \Sigma^{-1}(\alpha - \bar{\alpha}) \leq \kappa^2 \end{aligned} \tag{15}$$

where b is the benchmark holdings. The optimal solution to this problem is

$$\alpha - \bar{\alpha} = \sqrt{\frac{\kappa^2}{(\tilde{w} - b)^T \Sigma (\tilde{w} - b)}} \Sigma (\tilde{w} - b) \tag{16}$$

which implies that

$$\begin{aligned} \alpha^T(\tilde{w} - b) &= \bar{\alpha}^T(\tilde{w} - b) \\ &\quad - \kappa \left\| \Sigma^{1/2}(\tilde{w} - b) \right\| \end{aligned} \tag{17}$$

This gives the following robust optimisation problem for long-only active funds

$$\begin{aligned} &\text{maximise} && \bar{\alpha}^T w - \kappa \left\| \Sigma^{1/2}(\tilde{w} - b) \right\| \\ &\text{subject to} && e^T w = 1 \\ &&& (w - b)^T Q(w - b) \leq v \\ &&& w \geq 0 \end{aligned} \tag{18}$$

Now, let us see how this variant of robust objective function effectively adjusts expected return estimates. If the holding in an asset is below the benchmark weight, the α is adjusted upwards. Similarly, if the holding in an asset is above the benchmark weight, the α for that particular

asset is adjusted downwards. This behaviour is much more intuitive and performs much better in practice for active strategies.

General robust optimisation framework

All three forms of robust portfolio optimisation discussed thus far can all be cast in a single generalised form. Note that the only difference between Equations (4), (12) and (16) is the model portfolio that is compared with the vector of portfolio holdings \tilde{w} , in constructing the expected return adjustments. Let z be the generic ‘model’ portfolio. Then the generic expected return adjustment can be written as

$$\alpha - \bar{\alpha} - \sqrt{\frac{\kappa^2}{(\tilde{w} - z)^T \Sigma (\tilde{w} - z)}} \Sigma (\tilde{w} - z) \tag{19}$$

In Equations (4), (12) and (16), z is

$$0, \frac{e^T D \Sigma \tilde{w}}{e^T D \Sigma D^T e} D^T e \text{ and } b, \text{ respectively.}$$

Note that z can be dependent on \tilde{w} as it is for the zero-net alpha case.

This generic framework allows for the construction of other alternative forms of robust portfolio optimisation. Both of the alternatives introduced were created to prevent robust optimisation from adjusting alphas based on anything other than estimation error. For example, in the case of the active manager that measures performance relative to a benchmark, the portfolio weights are compared with the benchmark to expected returns from being adjusted because of the active manager’s constraints. Similarly, in the case of a fully invested fund, the zero-net alpha adjustment that compares the portfolio weights with the fully invested minimum estimation error portfolio was introduced. The adjustment prevents the expected returns from always being adjusted downwards because of the fully invested constraint.

For different investment strategies, other constraints may force the expected returns to be adjusted in a particular way, even if it is not suggested by estimation error. In these cases, the general form can be used to create an effective robust portfolio construction strategy.

Numerical experiments

In order to measure the effect of the proposed methodology on the efficient frontiers, the experiments used to produce Figure 10.2 were

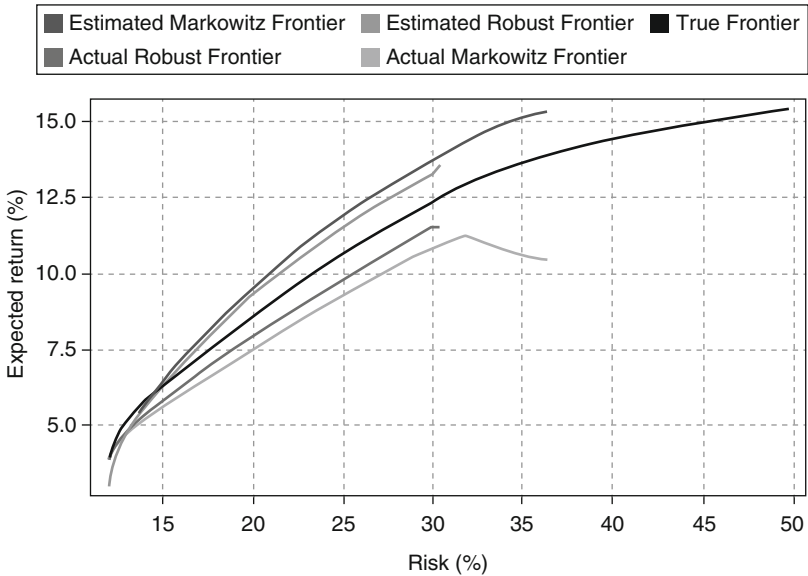


Figure 10.4 Robust efficient frontiers

re-run using robust optimisation. With $D = I$, efficient frontiers were generated using both the standard mean-variance problem and the equivalent robust optimisation problem and compared them with the true frontier as in Figure 10.2. These frontiers are illustrated in Figure 10.4. Similarly, the efficient frontier of active risk versus active return was generated using both classical mean-variance optimisation and the equivalent robust counterpart. These frontiers are illustrated in Figure 10.5. Incorporating the estimation error into the portfolio construction process significantly reduced its effect on the optimal portfolio. In both cases, the predicted return for any given risk level was not exaggerated nearly as much. More importantly, the actual robust frontiers are much closer to the true frontiers than are the actual mean-variance frontiers.

As expected, the computational experiments show that when using robust optimisation, the actual and estimated frontiers lie closer to each other. This is due to the objective function in the robust optimisation problem being based on reducing the distance between the predicted and actual frontiers. The real goal, however, is to get these frontiers not only closer together, but also closer to the true efficient frontier. It is believed that this result will be very difficult to establish theoretically,




Figure 10.5 Robust Active Return Efficient Frontiers

and for this reason we demonstrated it empirically by running a very large number of computational experiments, which are outlined below.

While frontiers help illustrate the effect of robust optimisation, they only represent one rebalancing period. One cannot say that portfolios constructed using robust optimisation will outperform those constructed using classical mean-variance optimisation each month with certainty. It is argued, however, that portfolios constructed using robust optimisation do outperform those constructed using classical mean-variance optimisation the majority of the time. To demonstrate this, simulated backtests were run using the various forms of robust optimisation described in this paper.

For each simulated backtest, a time-series of monthly returns was generated using the excess expected returns and covariance matrix from Idzorek (2002) for 30 US equities. For each month, a mean vector of returns μ is computed using the previous number of historical periods, T , specified in the backtest. The sample covariance matrix of returns, S , is computed over the same time horizon. For each month, an expected return estimate $\alpha = (1 - \lambda)\mu + \lambda r$ is computed, where r is that month's realised returns, and $0 \leq \lambda \leq 1$ is a parameter specified in the backtest.

Table 10.4 Backtest results for long-short dollar-neutral strategy

Lambda	Kappa	Markowitz Ann. Ret. (%)	Robust Ann. Ret. (%)	Robust Win (%)
0.025	1	2.06	2.90	82
0.025	2	2.06	3.62	80
0.025	3	2.06	4.22	83
0.025	4	2.06	4.71	83
0.050	1	10.53	11.3	79
0.050	2	10.53	11.92	78
0.050	3	10.53	12.38	76
0.050	4	10.53	12.66	76

The value of α used in the backtests intentionally contains some look-ahead bias that is designed to simulate portfolio managers' information.

The estimation error matrix $\Sigma = (1 - \lambda) S$, is used in the robust objective term for each backtest. The value of κ used in each backtest is specified in the results tables. All results are based on 100 different runs of the backtests using different seeds for the random number generation. All backtests cover 120 periods, or 10 years of monthly rebalancings.

The first set of backtests simulate a long-short dollar-neutral strategy with a limit on the total risk of 10 per cent. Asset weights were constrained to be within ± 25 per cent of the amount invested. The portfolio is also restricted so that the maximum total value of the long positions is equal to the amount invested in order to restrict leverage. The value of T in each of the backtests was 120. The results are shown in Table 10.4. The columns labeled 'Ann. Ret.' give the average annualised return over all 100 simulations. The column 'Robust win (%)' gives the percentage of the simulations in which the total return using robust optimisation was greater than the total return using classical mean-variance optimisation. In these tests, the total excess returns for the robust backtests were greater than the total excess returns for the classical tests between 76 and 83 per cent of the time. Also note that the average annualised return of the robust portfolios is between 84 and 265 basis points greater than the average annualised return of the classical portfolios.

The second set of backtests simulate a long-only maximum return strategy. Here, expected returns in a fully invested long-only portfolio are maximised with a limit of 20 per cent expected total risk. In these backtests, the monthly round-trip turnover is also limited to be at most 15 per cent by imposing a linear constraint in the portfolio construction

Table 10.5 Backtest results for long-only maximum total return strategy

Lambda	Kappa	Markowitz Ann. Ret. (%)	Robust Ann. Ret. (%)	Robust Win (%)
0.075	1	11.93	12.30	87
0.075	3	11.93	12.59	81
0.075	5	11.93	12.78	75
0.075	7	11.93	12.77	68
0.100	1	14.04	14.46	84
0.100	3	14.04	14.82	81
0.100	5	14.04	14.97	74
0.100	7	14.04	14.97	68

Table 10.6 Backtest results for long-only active strategy

Lambda	Kappa	Markowitz Ann. Ret. (%)	Robust Ann. Ret. (%)	Robust Win (%)
0.025	1	1.59	1.69	69
0.025	3	1.59	1.85	68
0.025	5	1.59	1.98	70
0.025	7	1.59	2.02	65
0.050	1	3.35	3.48	76
0.050	3	3.35	3.68	78
0.050	5	3.35	3.80	75
0.050	7	3.35	3.80	65

problem. The zero-net alpha adjustment version of robust optimisation introduced earlier was used to construct the robust portfolios. The value of T in each of the backtests was 120. The results are shown in Table 10.5. In these tests, the total excess returns for the robust backtests were greater than the total excess returns for the classical tests between 68 and 87 per cent of the time. Also note that the average annualised excess return of the robust portfolios is between 37 and 93 basis points greater than the average annualised excess return of the classical portfolios.

The last set of backtests simulate a long-only active strategy. The portfolios are constrained to be fully invested, have at most a 3 per cent active risk, and the asset weights must be within ± 10 per cent of the investment size of the benchmark weights. In these backtests, we also limited the monthly roundtrip turn over to be at most 15 per cent. The active return/active risk version of robust optimisation introduced in Section 4.2 was used to construct the robust portfolios. Again, the value

of T in each of the backtests was 120. The results of these backtests are given in Table 10.6. The results are based on active returns rather than excess returns, but otherwise show the same type of information as did the previous tables. Again, the total returns for the robust backtests were greater than the total returns for the classical tests. This time, the robust portfolios were superior between 65 and 78 per cent of the time. The average annualised active return of the robust portfolios is between 10 and 45 basis points greater than the average annualised active return of the classical portfolios.

Conclusions

The authors believe that one of the main reasons why modern portfolio theory is not being fully used in practical portfolio management is the fact that the result from a classical mean-variance framework are unstable and too sensitive to expected return estimates. It is argued that these ill-effects of classical portfolio optimisation are caused by the error-maximisation property. The robust optimisation technology described in this paper directly addresses these issues.

The frontier illustrations show that portfolios generated using robust optimisation may be closer to the true efficient frontier. The backtesting results indicate that portfolios constructed using robust optimisation outperformed those created using traditional mean-variance optimisation in the majority of cases. The realised returns were greater when using robust optimisation. The authors believe that the reason for this is that more information is transferred to the portfolios when constructing them using robust optimisation. Classical optimisation will tend to overweight assets with positive estimation error in the expected returns. Because of this, portfolios constructed using mean-variance optimisation typically represent less information from the true expected returns. That is, the portfolios constructed using robust optimisation usually have a higher correlation between the true expected returns and the alphas implied from the portfolio than do those portfolios constructed using robust optimisation.

Robust optimisation is a fairly new optimisation methodology that has not yet found widespread use in the financial community. Robust portfolio optimisation problem is indeed a more complex optimisation problem, but one that can be efficiently handled by a class of interior-point optimisers that are capable of handling second-order cone constraints. Therefore, based on the computational results in this paper, the authors believe that robust portfolio optimisation is a practical and effective portfolio construction methodology.

Notes

1. A convex combination of two vectors a and b is defined to be $\lambda a + (1 - \lambda)b$ where $0 \leq \lambda \leq 1$.
2. Σ is a symmetric positive definite matrix.
3. We do not need to assume normality. It is only required that the distribution is elliptical. An elliptic distribution is a symmetric distribution such that any minimum volume confidence region of the distribution is defined by an n -dimensional ellipsoid of the form of Equation (1).
4. The effective alpha, or effectively adjusted expected return, is defined as the value of α determined by Equation (4) or the related equation for a variant of (6) for the optimal solution to the robust optimisation problem.
5. When considering adjustments of expected returns, we assume that Σ is a diagonal matrix so that we can easily conceptualise the effective expected-return adjustments without worrying about any interactions between the adjustments. That is, assuming a diagonal matrix, Σ , means that the adjustment to the expected return for asset i is dependent upon the weight of asset i , but no others. This is not to say that the adjustments are truly independent, though. Constraints may force one weight to go up if another goes down, which implicitly creates an interrelationship between the alpha adjustments. Because Σ is positive definite, a diagonal Σ will have all positive elements.

References

- Ben-Tal, A. and Nemirovski, A. S. (1997) 'Robust Truss Topology Design via Semidefinite Programming', *SIAM Journal on Optimization*, 7(4): 991–1016.
- Best, M.J. and Grauer, R. R. (1991) '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.
- Black, F. and Litterman, R. (1990) 'Asset Allocation: Combining Investors' Views with Market Equilibrium', Technical report, Goldman, Sachs, *Fixed Income Research*, September.
- Broadie, M. (1993) 'Computing Efficient Frontiers using Estimated Parameters', *Annals of Operations Research: Special Issue on Financial Engineering*, 45; 21–58.
- Cavadini, H., Sbuelz, A. and Trojani, F. (2002) 'A Simplified Way of Incorporating Model Risk, Estimation Risk, and Robustness in Mean Variance Portfolio Management', Technical Report, University of Southern Switzerland.
- Ceria, S. and Stubbs, R. A. (2003) 'Overcoming Estimation Error in Portfolio Construction: Update in the Latest Developments in Optimisation', paper presented at CQA Conference in Las Vegas, NV, April.
- Chopra, V. and Ziemba, W. T. (1993) 'The Effects of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice', *Journal of Portfolio Management*, Winter, 6–11.
- Chopra, V. K. (1993) 'Mean-variance revisited: Near-optimal Portfolios and Sensitivity to Input Variations', *Russell Research Commentaries*, December.
- Garlappi, L., Uppal, R. and Wang, T. (2004) 'Portfolio Selection with Parameter and Model Uncertainty: A Multi-prior Approach', Technical Report, McCombs School of Business, The University of Texas at Austin, Austin, TX 78712, April.

- Goldfarb, D. and Iyengar, G. (2003) 'Robust Portfolio Selection Problems', *Mathematics of Operations Research*, 28(1), 1–38.
- Horst, J. R. ter, de Roon, F. A. and Werker, B. J. M. (2001) 'Incorporating Estimation Risk in Portfolio Choice', Technical Report, Department of Finance and Center, Tilburg University, Netherlands, January.
- Idzorek, T. (2002) 'A Step-by-step Guide to the Black–Litterman Model', Technical Report, Duke University.
- Jobson, J. D. and Korki, R. (1981) 'Putting Markowitz Theory to Work', *Journal of Portfolio Management*, 7, 70–74.
- Jorion, P. (1985) 'International Portfolio Diversification with Estimation Risk', *The Journal of Business*, 58(3), 259–278.
- Lobo, M., Vandenberghe, L., Boyd, S. and Lebet, H. (1985) 'Applications of Second-order Cone Programming', *Linear Algebra and its Application*, 284, 193–228.
- Markowitz, H. (1952) Portfolio selection', *Journal of Finance*, VII(1),77–91.
- Michaud, R. (1989) 'The Markowitz Optimisation Enigma: Is "Optimised" Optimal?' *Financial Analysts Journal*, 45(1), 31–42.
- Michaud, R. 'Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation', Harvard Business School Press, Boston, MA.
- Scherer, B. (2002) 'Portfolio Resampling: Review and Critique', *Financial Analysts Journal*, 58(6), 98–109.