Handling Parameter Uncertainty in Portfolio Risk Minimization

The robust optimization approach.

Dessislava Pachamanova

Markowitz [1952] was the first to model the important trade-off between risk and return in portfolio selection as an optimization problem. He suggested choosing an asset mix so that the portfolio variance is minimal for any target level of expected return.

60 years after Markowitz's seminal idea, despite substantial contributions to portfolio risk management theory, important practical issues remain unresolved. Portfolio allocation decisions are frequently made on the basis of solutions of optimization algorithms that treat parameters in the models such as means, variances, and covariances of returns, as given. These parameters, however, are estimated through error-prone procedures, e.g., statistical modeling or subjective evaluation. Since optimization results are very sensitive to perturbations in parameter values, the computed optimal portfolio strategies are unreliable.

In the case of standard mean-variance portfolio optimization, practitioners frequently resolve the issue by sampling the mean returns and the covariance matrix from a confidence interval around a nominal set of parameters, and then aggregating the portfolios obtained by solving a Markowitz problem for each sample. Unfortunately, this technique does not provide any guarantees, and may become quite inefficient as the number of assets grows.

Dessislava Pachamanova is an assistant professor of operations research at Babson College in Wellesley, MA.
dpachamanova@babson.edu

ROBUST OPTIMIZATION

This article explains how to use a recently developed field in optimization, *robust optimization*, to address uncertainty in parameter estimation problems. The advantage of the technique is that it takes uncertainty into consideration as the optimization problem is solved—the optimal solution remains feasible for any values of the uncertain parameters that fall within prespecified uncertainty sets. The most elementary uncertainty set is an interval for the value of each uncertain parameter (e.g., a 95% confidence interval based on historical data). More sophisticated uncertainty sets can be defined so that they depend on an investor's tolerance to overall estimation error, as well as on the covariance structure of the uncertain parameters.

Goldfarb and Iyengar [2003] develop a framework for applying robust optimization methods to some portfolio optimization problems, such as the Markowitz mean-variance model, when future asset returns are normally distributed and are given by a factor model. Their computational experiments suggest that the worst case performance of robust portfolios is significantly better than the worst case performance of portfolios optimized using classic methods.

Because of the fact that the minimum portfolio variance formulation is a quadratic optimization problem, and because of the particular definitions of the uncertainty sets that Goldfarb and Iyengar use, their robust optimization problem formulations are so-called second-order cone problems. Second-order cone problems are computationally tractable problems, but they require specialized nonlinear optimization software. Adding complex constraints, e.g., integrality constraints that restrict the portfolio structure, quickly turns such problems into very hard problems to solve.

There is also substantial evidence that real-world asset returns are not normally distributed. When returns are not symmetrically distributed, or when a downside risk is weighted more than an upside risk, the variance is not an accurate measure of investor risk preferences, and can lead to inferior allocation strategies. Such considerations and the theory of stochastic dominance (see Levy [1992]) have spurred interest in quantile-based risk measures, e.g., Value at Risk (VaR) (Jorion [2000]), and Conditional Value at Risk (CVaR) (Artzner et al. [1997], Rockafellar and Uryasev [2002]).

A detailed review of portfolio risk measures and estimation and optimization techniques, as well as examples of robust optimization applications to portfolio management, appears in Fabozzi et al. [2006]. Here I focus on a particular quantile-based measure, portfolio *shortfall* (Bertsimas, Lauprete, and Samarov [2004]).

Shortfall is proportional to standard deviation when returns follow symmetric elliptic distributions (*e.g., normal distribution*). Thus, the problem of minimizing portfolio shortfall is equivalent to the classic problem of minimizing portfolio standard deviation for an important class of asset returns. It has been shown also that shortfall is always at least as large as mean-adjusted VaR; that is, minimization of shortfall results also in optimization of the portfolio VaR.

Furthermore, the sample shortfall problem has a linear optimization problem formulation that makes it particularly appealing from a computational standpoint. Uncertainty sets for a robust formulation of the portfolio shortfall minimization problem can be defined in a way that preserves the linear structure of the optimization problem, so the problem can be handled by any optimization software.

Computational experiments with simulated and real market data indicate that robust formulations have a number of advantages in terms of portfolio return performance, and are therefore both an efficient and an effective way to incorporate uncertainty in portfolio optimization problems.

PORTFOLIO SHORTFALL MINIMIZATION PROBLEM

Shortfall derives from CVaR, and measures how great an expected loss will be if portfolio return drops below the α-quantile of its distribution. Mathematically, shortfall is defined as:

$$s_\alpha(\mathrm{x}) \doteq E[\tilde{\mathbf{r}}'\mathbf{x}] - E[\tilde{\mathbf{r}}'\mathbf{x} \mid \tilde{\mathbf{r}}'\mathbf{x} \leq q_\alpha(\tilde{\mathbf{r}}'\mathbf{x})],\ \alpha \in (0,1)$$

where $\tilde{\mathbf{r}}$ is a vector of random future asset returns, $\mathbf{x}$ is a vector of asset weights in the portfolio, q_α is the α-quantile of the distribution of random portfolio returns, and $E[.]$ denotes the standard expectations operator. The choice of α depends on the investor—by analogy with VaR, one could select $\alpha = 5\%$ or $\alpha = 1\%$.

In practice, quantile-based portfolio risk minimization is frequently implemented through sample optimization. Suppose we are given observations of vectors $\mathbf{r}_1, \ldots, \mathbf{r}_T$ that include the returns of N assets in T scenarios. These T scenarios can be obtained from historical data or generated in a more sophisticated way. A possible estimate for the vector of expected asset returns can be obtained by averaging the returns over the T scenarios, although a portfolio manager can input any other values deemed appropriate without changing the rest of the analysis.

The expected portfolio return for a given vector of asset weights $\mathbf{x}$ is $\hat{\mathbf{r}}'\mathbf{x}$. An estimator for the expected return if the portfolio return drops below the α-quantile of the distribution is obtained by averaging the K lowest portfolio returns $r_{(1)}(\mathbf{x}),\ldots,r_{(K)}(\mathbf{x})$. $K = [\alpha T]$ counts how many observations should be considered to be in the tail of the distribution. The portfolio shortfall is then estimated by subtracting the expected return in the tail from the expected return:

$$\hat{s}_\alpha(\mathbf{x}) = \hat{\mathbf{r}}'\mathbf{x} - \frac{1}{K}\sum_{t=1}^{K} r_{(t)}(\mathbf{x})$$

One cannot minimize the expression for the portfolio shortfall directly because of the presence of order statistics $r_{(1)}(\mathbf{x}),\ldots,r_{(K)}(\mathbf{x})$. Note that we cannot know their values in advance since they depend on the optimal vector of asset weights $\mathbf{x}$, which is an *output* of the optimization problem. Bertsimas, Lauprete, and Samarov [2004] use optimization duality theory to show that the portfolio shortfall minimization problem for a given level of desired return r_{target} can be written as:

$$\begin{aligned}
\min_{\mathbf{x},\nu,\mathbf{y}}\ & \hat{\mathbf{r}}'\mathbf{x} - \nu - \frac{1}{K}\sum_{t=1}^{T} y_t \\
\text{subject to}\ & \hat{\mathbf{r}}'\mathbf{x} \geq r_{target} \\
& \mathbf{x}'\mathbf{e} = 1 \\
& \nu + y_t \leq \mathbf{x}'\mathbf{r}_t,\ t = 1,\ldots,T \\
& y_t \leq 0,\ t = 1,\ldots,T
\end{aligned}$$

The unknowns in this problem are the asset weights $\mathbf{x}$ and the auxiliary variables ν and y_t, $t = 1,\ldots,T$. The variables ν and y_t are introduced so that the order statistics in the expression for shortfall can be written algebraically in a form (represented by the last two constraints) that can be entered into an optimization solver. The first two constraints are the traditional constraints in portfolio optimization—the expected return should be higher than the desired return, and the sum of the weights of all stocks in the portfolio should add up to 1.

Note that the problem structure is linear, which is an advantage both in terms of computational tractability and in terms of availability of commercial software.

INTUITION BEHIND THE ROBUST OPTIMIZATION METHOD

Suppose we think our estimates $\hat{\mathbf{r}}$ of expected asset returns in the shortfall problem formulation are inaccurate. We want to make portfolio shortfall minimization robust with respect to these inaccuracies, i.e., to modify the optimization problem formulation so that the portfolio allocation strategy it selects remains optimal even when the actual realizations of returns depart from their nominal (forecasted) values.

In practice, of course, we cannot protect ourselves against *all* possible events that could occur. In an uncertain world, it may be best to think of robustness in terms of *probability* that nature will act against us, and set the level of the probability we are willing to tolerate. We can specify an uncertainty set $\mathcal{U}^{\hat{\mathbf{r}}}$ for the vector of uncertain parameters $\hat{\mathbf{r}}$, try to relate it to probabilistic guarantees, and protect ourselves against the worst possible scenario when $\hat{\mathbf{r}}$ varies in that uncertainty set. The larger the uncertainty set we specify, the greater the level of protection we require.

Unfortunately, there is no free lunch: as we require more protection against inaccuracies, we restrict the number of acceptable portfolio allocation strategies, and sacrifice more of the optimality of the solution.

We can write a robust counterpart of the shortfall optimization problem as:

$$\begin{aligned}
\min_{\mathbf{x},\nu,\mathbf{y}}\ & \max_{\hat{\mathbf{r}}} \left\{\hat{\mathbf{r}}'\mathbf{x}\right\} - \nu - \frac{1}{K}\sum_{t=1}^{T} y_t \\
\text{subject to}\ & \min_{\hat{\mathbf{r}}} \left\{\hat{\mathbf{r}}'\mathbf{x}\right\} \geq r_{\mathbf{target}} \\
& \mathbf{x}'\mathbf{e} = 1 \\
& \nu + y_t \leq \mathbf{x}'\mathbf{r}_t,\ t = 1,\ldots,T \\
& y_t \leq 0,\ t = 1,\ldots,T \\
& \text{for every } \hat{\mathbf{r}} \text{ in } \mathcal{U}^{\hat{\mathbf{r}}}
\end{aligned}$$

The robust formulation looks very similar to the original nominal formulation, except that the inaccurate terms $\hat{\mathbf{r}}'\mathbf{x}$ are replaced with their worst case values for the particular expression in which they appear.

For example, the worst case scenario for the original constraint $\hat{\mathbf{r}}'\mathbf{x} \geq r_{target}$ happens when, because of inaccuracies in the estimate of $\hat{\mathbf{r}}$, the optimal asset allocation $\mathbf{x}$ for the nominal shortfall minimization problem results in an average portfolio return that is lower than the target return. Writing the constraint as $\min_{\hat{\mathbf{r}}}\{\hat{\mathbf{r}}'\mathbf{x}\} \geq r_{target}$ ensures that the optimal portfolio weights $\mathbf{x}$ that will be returned by the solver are chosen in such a way that even if the inaccuracies in the estimate of $\hat{\mathbf{r}}$ align in the worst possible way within the uncertainty set we have specified, the average portfolio return remains higher than the target return.

The selection of the uncertainty set for the inaccurate parameters is very important. Depending on the structure of the uncertainty set, the resulting robust optimization problem can be easy or virtually impossible to handle with commercial optimization solvers. For example, we might assume one is rarely so unlucky that all of one's estimates are wrong and differ by their worst case amount. In practice, it may make more sense to presume that the divergence in some parameters will be offset to some extent by divergence in other parameters, and will be influenced by the covariance structure of the uncertainty in the estimates.

Instead of individual worst case bounds for each parameter, one can therefore restrict the *total* deviation of the realized expected returns $\tilde{\mathbf{r}}$ from their nominal estimates $\hat{\mathbf{r}}$ to be less than or equal to some robustness budget Δ. The magnitude of Δ would correspond to one's tolerance for the total estimation error. Δ could be set, for example, to a value equal to 10% of the sum of the nominal (point) estimates for expected returns. Then the investor would be protected against movements in expected returns of up to 10% of the original estimate. An example of such an uncertainty set is

$$u^{\hat{\mathbf{r}}} = \{\tilde{\mathbf{r}} \mid \|\Sigma^{-\frac{1}{2}}\left(\tilde{\mathbf{r}} - \hat{\mathbf{r}}\right)\| \leq \Delta\} \tag{1}$$

where Σ is the covariance matrix of the uncertain estimates. In the simplest case, when the uncertain parameters are not correlated, each deviation of the particular uncertain parameter from its nominal value is divided by that parameter estimate's standard error. Intuitively, if there is little variability of an uncertain coefficient, it is unlikely that the coefficient will change by a lot. If it does change by a lot (an event of low probability), then the robustness budget Δ is affected more.

The norm $\|.\|$ can be any norm by which the distance is measured. For example, it can be the "natural distance" Euclidean norm. The d-norm described in Bertsimas, Pachamanova, and Sim [2004]), whose parameter can be chosen so that it approximates the Euclidean norm, can help preserve the linear structure of the optimization problem. A linear problem structure is preferable when solution time is an issue or when there are complex constraints.

Such constraints might include control of numerous positions and trades, maintaining a desired degree of liquidity, minimization of total transaction costs, and limiting portfolio exposure to different sectors or industries. The formulation of these constraints may require integer programming (see, for example, Bertsimas, Darnell, and Soucy [1999]). Adding such constraints to a non-linear problem makes the problem very difficult to handle computationally, but mixed-integer *linear* solvers are well developed and widely available.

Determining values for the parameters in the robust optimization model—Δ and standard deviations and covariances of nominal return estimates—is an important practical consideration. The appropriate values depend to a great extent on the procedure used to obtain the estimators of expected returns. If the expected returns $\hat{\mathbf{r}}$ are calculated by averaging historical data or returns from a simulation, one can estimate the variability in a particular expected asset return by computing the standard deviation of the returns of that asset, and dividing it by the square root of the number of observations.

More sophisticated forecasting techniques are also used to obtain expected return estimates. Many of these techniques are based on regression modeling. Then one possibility is to use the standard error of the regression parameter estimate as an approximation of the standard deviation of the expected return estimate.

The choice of value for the robustness budget Δ can be related to general probabilistic guarantees. The theoretical results in Bertsimas, Pachamanova, and Sim [2004], for example, can be used to make statements such as:

> The probability that the average portfolio return will be less than the target return given the number of stocks in the portfolio and the parameters of the norm used in the definition of the uncertainty set is at most 75%.

Unfortunately, because of the generality of these results (they hold for any distribution of returns), the bounds they provide are sometimes too high to be useful. It is recommended instead that the value for Δ be selected according to the particular investor's tolerance to error, and after backtesting the performance of the robust model for different possible values of Δ.

ROBUST PORTFOLIO SHORTFALL MINIMIZATION PROBLEM

The robust counterpart of the portfolio shortfall minimization problem for uncertainty sets for expected returns of the kind (1) can be derived using optimization duality theory (see Ben-Tal and Nemirovski [1998], El Ghaoui, Oustry, and Lebret [1998], and Bertsimas, Pachamanova, and Sim [2004]):

$$\min_{\mathbf{x},\nu,\mathbf{y}} \ \hat{\mathbf{r}}'\mathbf{x} + \Delta \| \Sigma^{\frac{1}{2}} \mathbf{x} \|^* \ \nu - \frac{1}{K}\sum_{t=1}^{T} y_t$$

$$\text{subject to} \ \ \hat{\mathbf{r}}'\mathbf{x} - \Delta \| \Sigma^{\frac{1}{2}} \mathbf{x} \|^* \geq r_{\mathbf{target}}$$

$$\nu + y_t \leq \mathbf{x}'\mathbf{r}_t, y_t \leq 0, t = 1,\ldots,T$$

$$\mathbf{x}'\mathbf{e} = 1$$

The robust formulation includes a buffer term $\Delta \| \Sigma^{\frac{1}{2}} \mathbf{x} \|^*$ in each expression with uncertain parameters, so that, in order to ensure robustness, the constraints are more restrictive than in the nominal optimization problem, and the objective function is more conservative. Note that the buffer term involves the degree of tolerance Δ the investor has for inaccuracy in estimates.

The symbol $\|.\|^*$ denotes the so-called *dual norm* of the norm used to define how the distance between the forecasted and the realized values of the uncertain return estimates is measured in the uncertainty set definition (for a discussion of dual norms, see Bertsimas, Pachamanova, and Sim [2004]). The dual norm of the "natural distance" Euclidean norm is the Euclidean norm itself. If the Euclidean norm is used, the robust formulation is a non-linear (second-order cone) problem. If the *d*-norm of Bertsimas, Pachamanova, and Sim [2004] is used, its dual norm can be expressed in such a way that the robust formulation is a linear optimization problem.

For the sake of simplicity, in the formulation above we ignore the uncertainty in the order statistics $r_{(i)}(\mathbf{x})$, and assume only that our estimates of expected returns are inaccurate. If one is concerned about it, a simple way to input it into the problem would be to specify an interval for each sample asset return (or a joint uncertainty set for all asset returns at a particular period $\mathbf{r}_t$), and treat the constraints $\nu + y_t \leq \mathbf{x}'\mathbf{r}_t, t = 1,\ldots,T$ the way we treat the expressions including $\hat{\mathbf{r}}$.

In our experience, however, the uncertainty in the estimates of the expected returns $\hat{\mathbf{r}}$ is the major factor in portfolio performance. In addition, incorporating too much variability makes the problem too conservative, and ultimately results in poorer portfolio allocation strategies than the strategies resulting from an optimization problem that does not take uncertainty into consideration at all.

EMPIRICAL COMPARISON

Is incorporating uncertainty about the accuracy of estimation of parameters in the portfolio risk minimization problem worth the effort? We perform a number of experiments with simulated and real-world data to compare the performance of robust and non-robust portfolio shortfall minimization strategies. To obtain some additional perspective, we also look at how the performance of the robust and non-robust minimum shortfall portfolios compares to the performance of portfolios obtained using classic mean-variance optimization.

Our general observations are that robust shortfall minimization results in portfolios with higher average and median returns, as well as higher worst case and best case returns than non-robust optimization. Both robust and non-robust minimum shortfall policies dominate classic mean-variance portfolios when asset returns follow non-normal distributions.

We provide examples of our computational results below. Although these experiments are not exhaustive, they suggest that robust optimization methods merit further attention in practice.

Simulated Data

The first set of simulation experiments looks at multivariate-normal asset returns. When future asset returns are drawn from a normal distribution, mean-variance portfolio allocation is optimal for any set of investor preferences (Chamberlain [1983]). Knowing the theoretically optimal solution allows us not only to compare the robust and non-robust strategies, but also to measure the difference between these strategies and the best possible portfolio strategy, classic mean-variance optimization. The second set of simulations tests how the different risk measures perform when future returns follow asymmetric (lognormal) distributions.

Consider a portfolio of $N = 150$ stocks with increasing expected (nominal) returns $\hat{r}_1 < \ldots < \hat{r}_i < \ldots < \hat{r}_{150}$ and increasing standard deviations $\sigma_1 < \ldots < \sigma_i < \ldots < \sigma_{150}$ generated by using the formulas: $\hat{r}_i = 1.15 + i\dfrac{0.05}{150}$, and $\sigma_i = (0.05/450)\sqrt{2iN(N+1)}$. This results in $\hat{r}_1 = 1.15,\ldots,\hat{r}_{150} = 1.2$, and $\sigma_1 = 0.0236,\ldots,\sigma_{150} = 0.2896$. Note that the stocks with the highest expected returns also have the highest variability. The stock returns are assumed to be uncorrelated.

We set the target portfolio return r_{target} to 17.5%, and run the following experiment 300 times:

1. We generate a sample of $T = 1000$ observations from the multivariate distribution of 150 stocks described above.

2. We estimate the sample covariance matrix for the asset returns, and solve for the optimal weights using the classic mean-variance approach.
3. We solve for the optimal weights of the non-robust portfolio shortfall minimization problem using the 20th percentile ($\alpha = 20\%$). According to Bertsimas, Lauprete, and Samarov [2004], for $\alpha > 15\%$, the weights of the minimum variance and minimum shortfall sample portfolio optimization are virtually identical, so we believe that $\alpha = 20\%$ gives us a good basis for comparison with the optimal mean-variance portfolio.
4. We solve for the optimal portfolio weights of the corresponding robust minimum shortfall optimization problem with parameters estimated from the sample data.
5. We simulate 1,000 realizations from the same multivariate-normal distribution of returns for a test set.
6. We compute the descriptive statistics of the realized returns of the optimal portfolio allocations for each approach for the test set data.

Exhibit 1 displays the *average* expected return, standard deviation, worst case return, best case return, and Sharpe ratio of the sample minimum variance (Sample MV) and the non-robust and robust shortfall approaches over the 300 runs with resampled training and testing sets of 1,000 multivariate-normal returns. The MV Perfect Information (MVPI) portfolio is computed using the Markowitz minimum variance formulation with correct values of the expected returns and standard deviations of all assets (i.e., the latter are not estimated from the training set, as the Sample MV, non-robust shortfall, and robust shortfall parameters are). Since the MVPI optimal portfolio computation is not contaminated by parameter estimation error, its performance can be considered theoretically optimal.

The descriptive statistics indicate that non-robust shortfall minimization generally produces portfolio strategies that result in expected returns that are comparable to the expected returns of the MVPI and the Sample MV portfolios. However, the nonrobust shortfall portfolios are characterized by a higher portfolio standard deviation, which leads to a lower Sharpe ratio. The Sharpe ratio is a good indicator of risk/return performance when the returns are multivariate-normal. The non-robust shortfall portfolios also have a worse worst case performance (more precisely, a wider worst case-best case spread).

Making shortfall minimization robust improves the performance significantly. As the protection level Δ increases, the average expected return and the Sharpe ratio increase as well. The standard deviation of the portfolio declines with Δ, and eventually starts increasing again. The worst case performance improves to surpass the worst case performance of both the MVPI and the sample mean-variance portfolios, while at the same time the best case performance remains better than the best case performance of the mean-variance portfolios.

EXHIBIT 1
Average Performance of Optimal Portfolios (Normal Returns)

Approach	Parameter Δ (% of Total Mean)	Mean	StdDev	Min	Max	Sharpe Ratio
MV Perfect Information	-	17.489	1.103	13.813	20.856	15.887
Sample MV	-	17.490	1.309	13.194	21.468	13.384
Nonrobust Shortfall	-	17.484	1.499	12.654	22.038	11.695
Robust Shortfall	0.00 (0.0)	17.484	1.499	12.654	22.038	11.695
	0.13 (0.5)	17.501	1.489	12.717	21.995	11.786
	0.26 (1.0)	17.519	1.481	12.782	21.989	11.860
	1.31 (5.0)	17.644	1.435	13.086	21.983	12.324
	2.62 (10.0)	17.768	1.408	13.284	22.073	12.639
	3.93 (15.0)	17.879	1.394	13.399	22.145	12.854
	5.24 (20.0)	17.980	1.390	13.520	22.220	12.963
	6.54 (25.0)	18.075	1.389	13.606	22.294	13.040
	7.85 (30.0)	18.167	1.390	13.686	22.409	13.098
	13.09 (50.0)	18.502	1.408	13.989	22.777	13.161
	20.94 (80.0)	18.981	1.443	14.385	23.315	13.177
	26.18 (100.0)	19.315	1.458	14.612	23.702	13.264

Obtained by running (1)-(6) 300 times. Average value for Δ over 300 iterations is given in terms of percentage total return. Terms in parentheses denote the percentage Δ represents of the total sample mean return. Training and test set sizes in each iteration are 1,000 observations from a multivariate-normal distribution.

Quantile-based risk measures have an advantage over standard deviation and variance, especially in the case of asymmetrically distributed asset returns, because they do not penalize for large positive deviations from the expected values. Exhibit 2 provides descriptive statistics for the optimal portfolio performance when returns are assumed to follow independent skewed (lognormal) distributions with the same means and standard deviations as the multivariate-normal data set. The quantile-based risk measures optimization results in portfolios with higher expected return and better worst case and best case performance.

In this case the average standard deviations and the Sharpe ratios are not as informative as in the case of multivariate-normal returns; a high standard deviation may simply mean that a larger number of high returns are realized by strategies resulting from quantile-based optimization than by strategies resulting from mean-variance optimization. Again, robust optimization helps to raise the expected portfolio return, and improves the worst case and best case performance significantly.

Real Market Data

We also test whether the observations from the controlled simulated experiments can be confirmed by real-world data, considering a portfolio of 31 stocks in different Dow Jones industry categories. We use monthly historical returns from April 1986 through December 2001, and record the realized portfolio returns as well as the cumulative portfolio performance over multiple investment periods.

Determining the inputs to the sample minimum variance and the robust and non-robust shortfall minimization problems—e.g., expected returns, standard deviations, and covariances of estimates—is a challenging task by itself. We take a simplified approach: We use a moving average of 60 past returns to determine the expected values and the covariance matrices of the next-period returns, as well as the standard deviations of the estimates of expected returns $\tilde{\mathbf{r}}$ for the next investment period.

There are 129 investment periods (from April 2001 through December 2001). We set the target portfolio return r_{target} for each investment period to the average return of a portfolio of the 31 stocks, equally weighted, over the previous 60 months. Again, the choice of values for the target returns is arbitrary; for example, we could tie the target return to the return on a major market index instead. Since the choice of input parameters affects both the robust and the non-robust portfolio shortfall minimization problems, we can still have a fair comparison of the performance of those two approaches. The performance of the minimum variance portfolios is used mostly as a benchmark.

We use $\alpha = 20\%$ and $K = 12$ in the portfolio shortfall optimization problem formulations. d, the d-norm parameter, is set equal to $\sqrt{31} = 5.57$ (when the d-norm parameter equals the square root of the number of stocks, the d-norm approximates the Euclidean norm most closely). The robustness budget Δ is set to a fixed percentage of the target return for each time period.

The summary statistics in Exhibit 3 indicate that the robust strategies on average result in higher returns than the non-robust strategies, and the standard deviations for the robust and the non-robust shortfall strategies are comparable. The non-robust shortfall and the robust shortfall strategies strongly dominate the sample minimum variance strategies in mean, standard deviation, median, and worst and best case return over the 129 investment periods. This is confirmed by the cumulative performance of the three types of approaches (Exhibit 4).

EXHIBIT 2
Average Performance of Optimal Portfolios (Lognormal Returns)

Approach	Parameter Δ (% of Total Mean)	Mean	StdDev	Min	Max	Sharpe Ratio
MV Perfect Information	-	17.492	1.101	14.281	21.531	15.891
Sample MV	-	17.493	1.202	14.045	22.216	14.558
Nonrobust Shortfall	-	17.528	1.297	13.848	22.827	13.528
Robust Shortfall	0.00 (0.0)	17.528	1.297	13.848	22.827	13.528
	0.13 (0.5)	17.540	1.295	13.868	22.840	13.559
	0.26 (1.0)	17.552	1.293	13.883	22.831	13.587
	1.31 (5.0)	17.635	1.284	14.008	22.841	13.747
	2.63 (10.0)	17.725	1.283	14.101	22.934	13.831
	3.94 (15.0)	17.804	1.286	14.176	23.045	13.852
	5.26 (20.0)	17.876	1.292	14.234	23.159	13.844
	6.57 (25.0)	17.944	1.298	14.297	23.278	13.838
	7.89 (30.0)	18.008	1.305	14.349	23.384	13.813
	13.15 (50.0)	18.242	1.336	14.518	23.800	13.669
	21.04 (80.0)	18.513	1.363	14.730	24.208	13.594
	26.30 (100.0)	18.735	1.393	14.888	24.561	13.462

EXHIBIT 3
Summary Statistics for Realized Returns—Real Market Data

	Sample MV	Non-robust ($\Delta = 0$)	Robust ($\Delta = 10\%$)	Robust ($\Delta = 20\%$)	Robust ($\Delta = 30\%$)	Robust ($\Delta = 40\%$)	Robust ($\Delta = 50\%$)
Mean	0.433	1.122	1.135	1.126	1.123	1.189	1.184
Std Dev	5.080	4.070	4.038	4.114	4.070	4.041	4.081
Median	0.266	1.115	1.116	1.254	1.423	0.924	1.390
Minimum	-14.773	-10.962	-11.377	-12.012	-11.116	-10.317	-11.366
Maximim	13.042	13.342	10.655	11.541	10.781	11.038	10.595

Exhibit 4

Cumulative Performance Comparison for Different Values of Δ

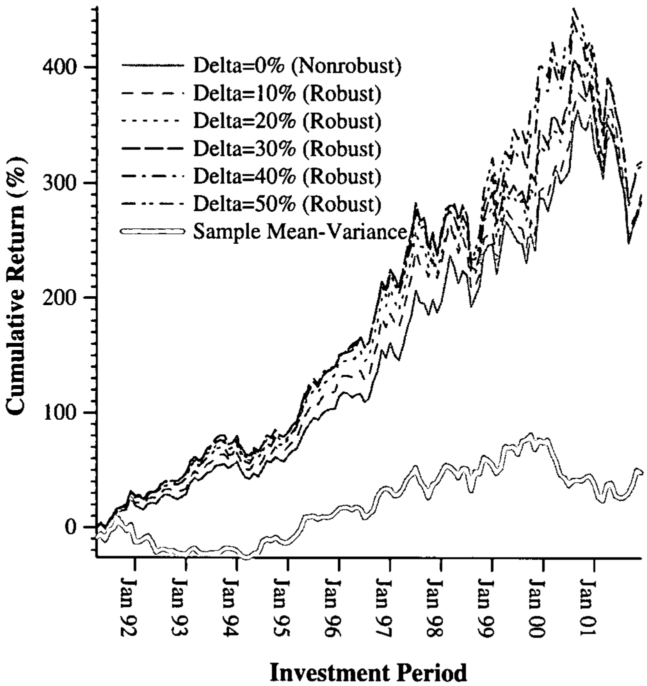

Over the 129 investment periods, the sample minimum variance strategies realize a 48.57% cumulative return versus 284.07% for the non-robust shortfall and 291.13%, 285.11%, 284.56%, 319.36%, and 315.42% for the robust shortfall with Δ equal to 10%, 20%, 30%, 40%, and 50% of the average return over the previous period, respectively. Even though we use simplified methods to estimate the parameters that go into the optimization problems, these results highlight the sensitivity of mean-variance optimal portfolios to volatility in parameter estimates.

We should note that the robust shortfall portfolios are not better than the non-robust shortfall one in all investment periods. However, whenever the robust portfolios underperform the non-robust portfolio, the realized returns fall outside the uncertainty set we have prespecified at the beginning. In other words, requiring protection against adverse movements in stock returns makes the robust portfolios more conservative, but we do not require a sufficient level of protection to account for the extreme movements in stock returns that happen in those particular instances.

On the whole, the robust shortfall strategies result in a portfolio return that is higher than the non-robust shortfall realized return in more than 50% of all investment periods (57% of all investment periods for Δ = 10%, 30%, or 40%; 55% for Δ = 20%; and 53% for Δ = 50%).

CONCLUDING REMARKS

We have described the intuition behind the robust optimization technique, and shown how to apply it in order to handle parameter uncertainty in the estimates of expected asset returns in portfolio shortfall minimization. Experiments with simulated and real-world market data suggest that robust optimization techniques may lead to superior investment strategies.

This approach can be extended to incorporate more complex uncertainty sets and different risk management problems. One primary issue for further study is how best to calibrate the robust model parameters for investor preferences and market conditions.

ENDNOTE

The author is grateful to Petter Kolm for valuable suggestions.

REFERENCES

Artzner, P., F. Delbaen, J.M. Eber, and D. Heath. "Thinking Coherently." *Risk*, Vol. 10 (1997), pp. 68-71.

Ben-Tal, A., and A. Nemirovski. "Robust Convex Optimization." *Mathematics of Operations Research*, Vol. 23, 4 (1998), pp. 769-805.

Bertsimas, D., C. Darnell, and R. Soucy. "Portfolio Construction Through Mixed-Integer Programming at Grantham, Mayo, Van Otterloo and Company." *Interfaces*, Vol. 29, 1 (1999), pp. 49-66.

Bertsimas, D., G. Lauprete, and A. Samarov. "Shortfall as a Risk Measure: Properties, Optimization, and Applications." *Journal of Economic Dynamics and Control*, Vol. 88, 7 (2004), p. 1353.

Bertsimas, D., D. Pachamanova, and M. Sim. "Robust Linear Optimization Under General Norms." *Operations Research Letters*, Vol. 32 (2004), pp. 510-516.

Chamberlain, G. "A Characterization of the Distributions That Imply Mean-Variance Utility Functions." *Journal of Economic Theory*, Vol. 29 (1983), pp. 185-201.

El Ghaoui, L., F. Oustry, and H. Lebret. "Robust Solutions to Uncertain Semidefinite Programs." *SIAM Journal on Optimization*, Vol. 9, 1 (1998), pp. 33-52.

Fabozzi, F., S. Focardi, P. Kolm, and D. Pachamanova. *Robust Portfolio Optimization and Estimation.* New York: John Wiley & Sons, 2006 (forthcoming).

Goldfarb, D., and G. Iyengar. "Robust Portfolio Selection Problems." *Mathematics of Operations Research*, Vol. 28, 1 (2003), pp. 1-39.

Jorion, P. *Value at Risk: The New Benchmark for Managing Financial Risk.* New York: McGraw-Hill, 2000.

Levy, H. "Stochastic Dominance and Expected Utility: Survey and Analysis." *Management Science*, Vol. 38, 4 (1992), pp. 555-593.

Markowitz, H.M. "Portfolio Selection." *Journal of Finance*, Vol. 7 (1952), pp. 77-91.

Rockafellar, R.T., and S. Uryasev. "Conditional Value-at-Risk for General Loss Distributions." *Journal of Banking and Finance*, Vol. 26, 7 (2002), pp. 1443-1471.