Robust Portfolio Optimization

Recent trends and future directions.

Frank J. Fabozzi, Petter N. Kolm,
Dessislava A. Pachamanova, and Sergio M. Focardi

Frank J. Fabozzi
is a professor in the practice of finance in the School of Management at Yale University in New Haven, CT.
frank.fabozzi@yale.edu

Petter N. Kolm
is a doctoral student in finance at Yale University, and a financial consultant in New York City.
petter.kolm@gmail.com

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

Sergio M. Focardi
is a partner of the Intertek Group in Paris.
interteksf@aol.com

The concepts of portfolio optimization and diversification have been instrumental in the understanding of financial markets and the development of financial decision-making. The major breakthrough came in 1952 with the publication of Harry Markowitz's theory of portfolio selection. Markowitz suggested that sound financial decision-making was a quantitative trade-off between risk and return. His work spurred a vast amount of research on quantifying market behavior—one of the main practical consequences was the acceptance of the notion that diversification reduces portfolio risk.

More than 50 years after Markowitz's seminal work, substantial advances have been made in both the theory and the practice of portfolio management. Today, quantitative techniques for forecasting asset returns, and for portfolio allocation, risk measurement, trading, and rebalancing, to mention a few, have a major presence in the financial industry. Their proliferation has been facilitated by the reduced cost of computing power and the increased availability of sophisticated and specialized software, allowing investors to incorporate their forecasts about the future direction of markets into disciplined analytical frameworks.

As the use of quantitative techniques has become widespread in the investment industry, the consideration of estimation risk and model risk has grown in importance. Bayesian techniques and robust estimation of model parameters, for example, are now common in financial applications. Most recently, practitioners have begun incorporating the uncertainty introduced by estimation errors directly into the portfolio optimization process by mathematical techniques referred to as *robust optimization.*

Unlike the traditional approach, where inputs to the portfolio allocation framework are treated as deterministic, robust portfolio optimization incorporates the notion that inputs have been estimated with errors. In this case, the inputs are not the traditional forecasts, such as expected returns and asset covariances, but rather *uncertainty sets* including these point estimates (e.g., confidence intervals around the forecasts).

We review major trends in robust optimization and its applications in portfolio management. We begin by explaining the main ideas behind the robust optimization approach, and discuss the relation between robust optimization and other robust methods for portfolio management. We describe some of the latest developments in robust optimization applications, and discuss future directions in robust portfolio management.

THE ROBUST OPTIMIZATION APPROACH

Introduced in the operations research literature by Ben-Tal and Nemirovski [1998] and El Ghaoui and Lebret [1997], modern robust optimization allows a portfolio manager to solve a robust formulation of the portfolio optimization problem with one single call to an optimization solver in about the same time as the classical portfolio optimization problem. The resulting optimal portfolio allocations tend to be more stable and less sensitive to changes in model parameters.

Consider the classical mean-variance portfolio allocation problem:

$$\begin{aligned} \max_{\mathbf{w}} \quad & \boldsymbol{\mu}'\mathbf{w} - \lambda\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w} \\ s.t. \quad & \mathbf{w}'\boldsymbol{\iota} = 1 \end{aligned}$$

where $\boldsymbol{\mu}$ is the vector of expected returns (alphas) for N assets in the investment universe, $\mathbf{w}$ is the N-dimensional vector of portfolio weights, λ is the risk aversion coefficient, $\boldsymbol{\Sigma}$ is the asset-asset covariance matrix, and $\boldsymbol{\iota}$ is a vector of ones.

This optimization problem simply states that the optimal portfolio weights should be chosen so that the expected portfolio return less the portfolio risk (scaled by the risk aversion coefficient) is as high as possible. The equality constraint ensures that the portfolio weights add up to one.

As demonstrated by Black and Litterman [1992], a small change in the expected asset returns can result in large changes in the optimal portfolio allocation. In other words, the classical portfolio optimization problem is *not* robust with respect to small changes in its inputs. As in practice expected returns and asset covariances cannot be measured exactly but have to be estimated—sometimes with large errors—it is important that we take uncertainty resulting from estimation errors into account.

One way to make the optimization problem robust with respect to estimation errors is to require that the optimal solution remain optimal for all values of the expected returns that are close to the estimates of the expected returns $\hat{\boldsymbol{\mu}}$. We can express this requirement in the optimization problem as follows: Instead of using the estimate $\hat{\boldsymbol{\mu}}$ of $\boldsymbol{\mu}$, we consider a set of vectors that are close to the estimate $\hat{\boldsymbol{\mu}}$, and solve the optimization problem for all vectors in this set.

The idea here is that the expected returns may have been estimated with some error, but the estimates are not too far away from the true expected returns. Mathematically, this idea is incorporated in the definition of an *uncertainty set* for $\hat{\boldsymbol{\mu}}$:

$$U_{\delta}(\hat{\boldsymbol{\mu}}) = \{\boldsymbol{\mu} \,|\, |\mu_i - \hat{\mu}_i| \leq \delta_i, \quad i = 1, \ldots, N\} \tag{1}$$

In words, the set $U_{\delta}(\hat{\boldsymbol{\mu}})$ includes all vectors $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_N)$ such that each component μ_i is in the interval $[\hat{\mu}_i - \delta_i, \hat{\mu}_i + \delta_i]$, and is often referred to as a "box" uncertainty set. From a statistical point of view, these intervals can be chosen to be certain confidence intervals around each point estimate μ_i.

We solve a modification of the original portfolio optimization problem such that even if $\boldsymbol{\mu}$ takes its worst possible value within the uncertainty set, the allocation remains optimal. Namely, we solve the max-min portfolio optimization problem:

$$\begin{aligned} \max_{\mathbf{w}} \quad & \left\{ \min_{\boldsymbol{\mu} \in U_{\delta}(\hat{\boldsymbol{\mu}})} \{\boldsymbol{\mu}'\mathbf{w}\} - \lambda\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w} \right\} \\ s.t. \quad & \mathbf{w}'\boldsymbol{\iota} = 1 \end{aligned}$$

At first sight, this optimization problem looks complicated, as we have to minimize the objective function with respect to $\boldsymbol{\mu}$ over the specified uncertainty set and simultaneously maximize the objective function with respect to $\mathbf{w}$ to find the optimal allocation. As we will see shortly, this problem can be reformulated into an equivalent maximization problem with respect to only $\mathbf{w}$. First, let us understand what this model implies from an intuitive perspective.

Observe that this model incorporates the notion of aversion to estimation error in the following sense. When the interval $[\hat{\mu}_i - \delta_i, \hat{\mu}_i + \delta_i]$ for the expected return of the i-th asset is wide, meaning that the expected return has been estimated with large estimation error, the minimization problem over $\boldsymbol{\mu}$ is less constrained. Consequently, the minimum will be lower than it would be when the interval for $\hat{\mu}_i$ is narrower. Obviously when the interval is narrow enough, the minimization problem will be so tightly constrained that it would deliver a solution that is close to the optimal solution of the classical portfolio optimization problem that ignores estimation errors. In other words, it is the *size* of the intervals (in general, the size of the uncertainty set) that controls the aversion to the uncertainty that comes from estimation errors.

The robust version of the classical portfolio optimization problem is obtained by solving the max-min problem above, and is in this case easy to derive without any involved mathematics. Namely, it is:

$$\begin{array}{ll} \max\limits_{\mathbf{w}} & \hat{\boldsymbol{\mu}}'\mathbf{w} - \boldsymbol{\delta}'|\mathbf{w}| - \lambda \mathbf{w}'\boldsymbol{\Sigma}\mathbf{w} \\ s.t. & \mathbf{w}'\boldsymbol{\iota} = 1 \end{array}$$

where $|\mathbf{w}|$ denotes the absolute value of the entries of the vector of weights $\mathbf{w}$.

To gain some intuition, notice that if the weight of asset i in the portfolio is negative, the worst case expected return for asset i is $\hat{\mu}_i + \delta_i$ (we lose the highest amount possible). If the weight of asset i in the portfolio is positive, then the worst case expected return for asset i is $\hat{\mu}_i - \delta_i$ (we gain the lowest amount possible). Observe that $\hat{\mu}_i w_i - \delta_i |w_i|$ equals $(\hat{\mu}_i - \delta_i)w_i$ if the weight w_i is positive and $(\hat{\mu}_i + \delta_i)w_i$ if the weight w_i is negative.

Hence, the mathematical expression in the objective agrees with our intuition: It minimizes the worst case expected portfolio return. In this robust version of the mean-variance formulation, assets whose mean return estimates are less accurate (have a higher estimation error δ_i) are therefore penalized in the objective function, and will tend to have a lower weight in the optimal portfolio allocation.

This optimization problem has the same computational complexity as the non-robust mean-variance formulation—namely, it can be stated as a quadratic optimization problem. The latter can be achieved by using a standard trick that allows us to get rid of the absolute values for the weights. The idea is to introduce an N-dimensional vector of additional variables $\boldsymbol{\psi}$ to replace the absolute values, and to write an equivalent version of the optimization problem:

$$\begin{array}{ll} \max\limits_{\mathbf{w},\boldsymbol{\psi}} & \hat{\boldsymbol{\mu}}'\mathbf{w} - \boldsymbol{\delta}'\boldsymbol{\psi} - \lambda \mathbf{w}'\boldsymbol{\Sigma}\mathbf{w} \\ s.t. & \mathbf{w}'\boldsymbol{\iota} = 1 \\ & \psi_i \geq w_i;\ \psi_i \geq -w_i,\quad i = 1,\ldots,N \end{array}$$

Therefore, incorporating considerations about uncertainty in the estimates of the expected returns in this example has virtually no additional computational cost.

We can view the effect of this particular "robustification" of the mean-variance portfolio optimization formulation in two different ways. On the one hand, we can see that the values of the expected returns for the different assets have been adjusted downward in the objective function of the optimization problem. That is, the robust optimization model "shrinks" the expected return of assets with large estimation error. On the other hand, we can interpret the additional term in the objective function as a "risk-like" term that represents a penalty for estimation error. The size of the penalty is determined by the investor's aversion to estimation risk, and is reflected in the size of the deltas.

More complicated specifications for uncertainty sets have more involved mathematical representations, but can still be selected so that they preserve an easy computational structure for the robust optimization problem. For example, a frequently used uncertainty set is

$$U_\delta(\hat{\boldsymbol{\mu}}) = \{\boldsymbol{\mu} \mid (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}})'\boldsymbol{\Sigma}_\mu^{-1}(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}) \leq \delta^2\} \qquad (2)$$

where $\boldsymbol{\Sigma}_\mu$ is the covariance matrix of estimation errors for the vector of expected returns $\boldsymbol{\mu}$. This uncertainty set represents the requirement that the scaled sum of squares (scaled by the inverse of the covariance matrix of estimation errors) between all elements in the set and the point estimates $\hat{\mu}_1, \hat{\mu}_2, \ldots, \hat{\mu}_N$ can be no higher than δ^2.

Note that this uncertainty set cannot be interpreted as individual confidence intervals around each point estimate. Instead, those familiar with statistics will notice that this uncertainty set captures the idea of a joint confidence region used, for example, in Wald tests. In practice, the covariance matrix of estimation errors is often assumed to be diagonal. In this particular case, the set includes all vectors of expected returns that are within a certain number of standard deviations from the point estimate of the vector of expected returns, and the resulting robust portfolio optimization problem would protect the investor if the vector of expected returns is indeed within that range.

Selecting Uncertainty Sets from Statistical Procedures

How do we select uncertainty sets for a particular application? In practice, their shape and size are usually based on statistical estimates and probabilistic guarantees. For example, uncertainty set (1) defines an N-dimensional box: It considers possible deviations of the N uncertain parameters from their expected values, and the resulting robust portfolio optimization problem protects against the worst possible realization of each individual parameter *separately*.

Uncertainty set (2) defines an N-dimensional ellipsoid (in two dimensions, an ellipsoid is an ellipse), and is not as conservative as (1). The resulting robust portfolio optimization offers protection from the worst possible *joint* deviation of the actual expected returns from the forecasts, by considering the correlations between the estimation errors of the uncertain parameters through the covariance matrix $\boldsymbol{\Sigma}_{\mu}$.

The calibration of the parameters that enter the definition of uncertainty sets is very important. For example, the intervals for $\hat{\mu}$ that define the uncertainty set (1) above can be matched to 95% or 99% confidence intervals for the estimates of the expected returns. The value of the parameter δ in the uncertainty set (2) can be related to probabilistic guarantees, as we will explain below.

The covariance matrix $\boldsymbol{\Sigma}_{\mu}$ of the errors in the estimated expected asset returns in uncertainty set (2) can be obtained using several different techniques. Its estimation can be problematic all the same, because it is hard to separate the estimation error in expected returns from the inherent variability in actual realized returns (Lee, Stefek, and Zhelenyak [2006]). Specifically, if a portfolio manager forecasts a 5% active return over the next time period, but achieves 1%, he cannot argue that there was a 4% error in his *expected* return, so evaluating $\boldsymbol{\Sigma}_{\mu}$ from historical data can be tricky.

In theory, if returns in a given sample of size T are assumed to come from a normal distribution, $\boldsymbol{\Sigma}_{\mu}$ equals $(1/T)\cdot \boldsymbol{\Sigma}$, where $\boldsymbol{\Sigma}$ is the covariance matrix of asset returns as before. Yet experience seems to suggest this may not be the best method in practice. One issue is that this approach applies only in a world where returns are stationary.

Another important issue is whether the estimate of the asset covariance matrix $\boldsymbol{\Sigma}$ itself is reliable if it is estimated from a sample of historical data. It is well-known that computing a meaningful asset return covariance matrix requires a large number of observations—many more observations than the number of assets in the portfolio—and even then the sample covariance matrix may involve large estimation errors that produce poor results in the mean-variance optimization. A fix when sufficient data are not available is to compute the estimation errors in expected returns at a factor level (e.g., industry, country, sector), and use their variances and covariances in the estimation error covariance matrix for the individual asset returns in a manner similar to standard factor models.

Several approximate methods for estimating $\boldsymbol{\Sigma}_{\mu}$ have also been found to work well in practice (see Stubbs and Vance [2005]). For example, it has been observed that simpler estimation approaches, such as computing the diagonal matrix of the variances of the estimates (as opposed to the complete error covariance matrix), often provide most of the benefit in robust portfolio optimization. In addition, standard approaches for estimating expected returns, such as Bayesian statistics and regression-based methods, generate estimates for the estimation error covariance matrix in the process of generating the estimates themselves.

Uncertainty sets (1) and (2) are both *symmetric*; that is, the sets are symmetric around the vector of estimates of the uncertain parameters $\hat{\boldsymbol{\mu}}$. One can also consider asymmetric uncertainty sets that better reflect information about the probability distributions of the uncertain parameters when the probability distributions are skewed (see Natarajan, Pachamanova, and Sim [2005]).

Recently, there has also been substantial interest in developing "structured" uncertainty sets, that is, uncertainty sets that are constructed for a specific purpose. Frequently, structured uncertainty sets based on simple intersections of elementary uncertainty sets are used to minimize the conservatism in traditional ellipsoidal or box uncertainty sets.

Clarifying a Misconception about Robust Optimization

Among practitioners, the notion of robust portfolio optimization is often equated with the robust mean-variance model, with uncertainty set (1) or (2) for the expected asset returns. While robust optimization applications frequently involve one form or another of this model, the actual scope of robust optimization can be much broader.

Note that the term, *robust optimization,* refers to the *technique* of incorporating information about uncertainty sets for the parameters in the optimization model, and not to the specific definitions of uncertainty sets or the choice of the parameters to model as uncertain. For example, we can use the robust optimization methodology to incorporate considerations for uncertainty in the estimate of the covariance matrix in addition to the uncertainty in expected returns, and obtain a different robust portfolio allocation formulation.

Robust optimization can be applied also to portfolio allocation models that are different from the mean-variance framework, e.g., the Sharpe ratio and value at risk optimization (see, for example, Goldfarb and Iyengar [2003] and Natarajan, Pachamanova, and Sim [2005]). There are numerous useful and reasonable robust formulations, and a complete review is beyond the scope of this article. See Fabozzi et al. [2007] for further details.

BAYESIAN METHODS AND ECONOMIC THEORY

Critics have argued that robust optimization is not really different from shrinkage estimators that combine the minimum-variance portfolio with a speculative investment portfolio. In the case of a particular uncertainty set for the expected returns (assuming all other parameters in the mean-variance problem are certain), it can be shown that the optimal mean-variance portfolio weights using robust optimization are a linear combination of the weights of the minimum-variance portfolio (which is independent of investor preferences or expected returns) and a mean-variance efficient portfolio. These portfolio weights can also be obtained by solving a standard mean-variance problem with expected return estimates derived from a standard shrinkage estimator with specific shrinkage parameters (see, for example, Garlappi, Uppal, and Wang [2007] and Scherer [2005]).

Robust optimization thus appears to offer a less transparent way to express investor preferences and tolerance for uncertainty than other approaches, such as Bayesian methods, where the shrinkage parameters can be defined explicitly.

In general, however, robust optimization is not necessarily equivalent to shrinkage estimation. For instance, differences are apparent in the presence of additional portfolio constraints. Furthermore, as we mentioned earlier, the robust optimization methodology can be used to account for uncertainty in parameters other than expected asset returns (like covariances of asset returns), making its relation to Bayesian estimation even less obvious.

The concept of robust optimization has been criticized from the point of view of classical economic theory (see, for example, Sims [2001]). From a behavioral and decision-making point of view, few individuals have max-min preferences. Indeed, max-min preferences describe the behavior of decision-makers who face great ambiguity and thus make choices consistent with the belief that the worst possible outcomes are highly likely. This kind of conservative behavior is not typical of the average investor. The problem of overconservativeness in applying robust optimization, however, can be controlled by modifying the specification of uncertainty sets for the parameters.

ROBUST PORTFOLIO OPTIMIZATION IN PRACTICE

One of the main problems in assessing the practical benefits of the robust optimization approach is that its performance depends on the choice (or calibration) of the model parameters, such as the coefficient of aversion to estimation error δ. In a sense, however, this issue is no different from the calibration of standard parameters in the classical portfolio optimization framework, such as the length of the estimation period to use for forecast generation and the choice of the risk aversion coefficient. These and other parameters need to be determined empirically or subjectively.

Note also that other robust modeling devices such as Bayesian estimation (James-Stein shrinkage estimators and the Black-Litterman model) have similar issues.[1] In particular, for shrinkage estimators, the portfolio manager needs to determine which shrinkage target to use and the size of the shrinkage parameter. In the Black-Litterman model, she needs to provide her confidence in equilibrium as well as her confidence in her views.

The values of the robust formulation parameters can sometimes be matched to probabilistic guarantees. For example, if the estimates of the expected asset returns are assumed to be normally distributed, there is an ω% chance that the estimates will fall in the ellipsoidal set (2) around the manager's estimates $\hat{\mu}$:

$$U_\delta(\hat{\mu}) = \{\mu \mid (\mu - \hat{\mu})'\Sigma_\mu^{-1}(\mu - \hat{\mu}) \le \delta^2\}$$

if δ^2 is assigned the value of the ωth percentile of a χ^2 distribution with degrees of freedom equal to the number of assets in the portfolio.

As an example, suppose there are 15 assets in the asset universe and that all returns are normally distributed. If we choose $\delta^2 = 25$, then 95% of all expected returns will be in the set $U_\delta(\hat{\mu})$.

More generally, if the expected returns can belong to any possible probability distribution, then assigning

$$\delta = \sqrt{\frac{1-\omega}{\omega}}$$

guarantees that the estimates will fall in the uncertainty set $U_\delta(\hat{\mu})$ with probability at least φ% (El Ghaoui, Oks, and Oustry [2003]).

It has been observed that in practice the standard robust mean-variance formulation with the uncertainty set specification above for estimated expected returns may result in portfolio allocations that are too pessimistic. Recall that the traditional robust counterpart tries to find the optimal solution so that constraints with uncertain coefficients are satisfied for the worst case realizations of the uncertain parameters.

Naturally, the larger the uncertainty set, the greater the chance that the optimal portfolio allocation will be conservative. Therefore, especially when the worst case expected returns can be far away from the estimated expected returns, some performance may be sacrificed.

Of course, we can always make a formulation less pessimistic by considering a smaller uncertainty set. For the uncertainty set above, we can achieve this by reducing the parameter δ, but, there is a recent trend among practitioners to apply more structured restrictions. We provide an example of a structured uncertainty set next.

When we first formulated the robust portfolio optimization problem, we assumed that all the actual realizations of expected returns could be worse than their expected values. Thus, the net adjustment in the expected portfolio return will always be downward.

While this leads to a more robust problem than the original one, in many instances it may be too pessimistic to assume that *all* estimation errors go against us. Instead, it may be more reasonable to believe that at least some of the true realizations will be above their expected values.

For example, we may make the assumption that there are approximately as many realizations above the estimated values as there are realizations below them. This condition can be incorporated in the portfolio optimization problem by adding an additional restriction to the uncertainty set (2). Ceria and Stubbs [2006] refer to this adjustment as a "zero net alpha adjustment."

Instead of adjusting the alphas of the estimates, we can perform this kind of adjustment also on their standard deviations or variances. It can be shown that the effect of the zero net adjustment is equivalent to modifying the covariance matrix $\mathbf{\Sigma}_\mu$ of estimation errors for the expected returns. Tests with real data indicate that robust mean-variance optimization with this kind of adjustment for expected return estimates outperforms classical mean-variance optimization 70% to 80% of the time.

Other structured uncertainty sets include tiered uncertainty sets where some of the uncertain parameters are modeled as *well-behaved*, while others are modeled as *misbehaving*. The modeler can require protection against a prespecified number of parameters that he believes will misbehave, that is, that will diverge significantly from their expected values (see, for example, Bienstock [2006]). In the context of portfolio optimization, we would specify misbehaving parameters as those realizations of expected asset returns that are likely to be lower than their estimates.

EFFECT OF ROBUST PORTFOLIO OPTIMIZATION FORMULATIONS ON PERFORMANCE

As we have noted, some tests with simulated and real market data indicate that robust optimization, when inaccuracy is assumed in the expected return estimates, outperforms classical mean-variance optimization in terms of total excess return a high percentage of the time. Other tests have not been as conclusive (Lee, Stefek, and Zhelenyak [2006]). The factor that accounts for much of the difference is how the uncertainty in parameters is modeled. Therefore, finding a suitable degree of robustness and appropriate definitions of uncertainty sets can have a significant impact on portfolio performance.

Independent tests by practitioners and academics using both simulated and market data appear to confirm that robust optimization generally results in more stable portfolio weights, i.e., that it eliminates the extreme corner solutions resulting from traditional mean-variance optimization. Robust mean-variance optimization also appears to improve worst case portfolio performance, and results in smoother and more consistent portfolio returns. Finally, by preventing large swings in positions, robust optimization frequently makes better use of the portfolio turnover budget and risk constraints.

Robust optimization, however, is not a panacea. By using robust portfolio optimization, investors are likely to trade off the optimality of their portfolio allocation in cases in which nature behaves as they predicted for protection against the risk of inaccurate estimation. Therefore, investors using the technique should not expect to do better than classical portfolio optimization when estimation errors have little impact, or when typical scenarios occur. They should, however, expect insurance when their estimates differ from the actual realized values by up to the amount they have prespecified in the modeling process.

PRACTICAL CONSIDERATIONS FOR ROBUST PORTFOLIO ALLOCATION

Which type of robust models is best for modeling financial portfolios? The short answer is: it depends. Among others, it depends on the size of the portfolio, the type of assets and their distributional characteristics, the portfolio strategies and trading styles involved, and the current infrastructure. Sometimes it makes sense to consider a combination of several techniques, such as a blend of Bayesian estimation and robust portfolio optimization. This is an empirical question—indeed, the only way to find out which strategy performs best is through thorough research and testing. A simple step-by-step checklist for robust quantitative portfolio management could include:[2]

1. Risk forecasting: Develop an accurate risk model.
2. Return forecasting: Construct robust expected return estimates.
3. Classical portfolio optimization: Start with a simple framework.
4. Model risk mitigation:
   a. Minimize estimation risk through the use of robust estimators.
   b. Improve the stability of the optimization framework through robust optimization.
5. Extensions.

In general, the most difficult item in this list is the calculation of robust expected return estimates. Developing profitable trading strategies ("α generation") is notoriously hard, but not impossible. It is important to remember that modern portfolio optimization techniques and fancy mathematics are not going to help at all if the underlying trading strategies are poor.

Implicit in this list is that for *each step* one needs to perform thorough testing in order to understand the effect of changes in and new additions to the model. It is not unusual that quantitative analysts and portfolio managers will have to revisit previous steps as part of the research and development process.

For example, it is important to understand the interplay between forecast generation and the reliability of optimized portfolio weights. Introducing a robust optimizer may lead to more reliable, and often more stable, portfolio weights. Yet how to make the optimization framework more robust depends on how expected return and risk forecasts are produced. Therefore, sometimes one has to refine or modify basic forecast generation. Identifying the individual and the combined contribution of different techniques is crucial in the development of a successful quantitative framework.

Minimizing estimation risk and improving the reliability of the optimization framework can proceed in either order, or sometimes at the same time. The goal of both approaches is of course to improve the overall reliability and performance of the portfolio allocation framework. Some important questions to consider here are: When and why does the framework perform well (or poorly)? How sensitive is it to changes in inputs? How does it behave when constraints change? Are portfolio weights intuitive—do they make sense? How high is the turnover of the portfolio over time?

Starting from the simple framework of classical portfolio optimization, many extensions are possible. Typical examples include the introduction of transaction cost models; more complex constraints (for example, integer constraints such as round-lotting or cardinality constraints); different risk measures (e.g, downside risk measures, higher moments); and dynamic and stochastic programming methods for incorporating intertemporal dependencies.

FUTURE DIRECTIONS

We conclude this article with some thoughts on the future direction of what can be described as *robust portfolio management.*

Advances in the mathematical and physical sciences have always had a major impact on finance. Probability theory, statistics, econometrics, and operations research have provided the necessary tools and discipline for the development of modern financial economics and large-scale portfolio management. The substantial advances in the areas of robust estimation and robust optimization during the 1990s have proven to be of significant importance for the practical applicability and reliability of portfolio management and optimization.

From a theoretical perspective, robust optimization is quite mature. By contrast, there are many unanswered questions in the practice of robust portfolio optimization. There is a need for more empirical research in order to provide better guidelines for applying robust optimization in a way that guarantees superior portfolio performance.

Practitioners need particularly to understand better: 1) the implications of using different types of uncertainty sets; 2) the interaction between different forecast generation methods (estimation techniques) and robust optimization; 3) how to calibrate model parameters in the optimization model; and 4) how to deal with the over-conservatism inherent in many robust models.

The robust optimization framework offers great flexibility and many new interesting possibilities in portfolio management. For instance, robust portfolio optimization can exploit the notion of *statistically equivalent portfolios*. Specifically, with robust optimization, a manager can find the best portfolio that 1) minimizes trading costs with respect to the current holdings, and 2) has an expected portfolio return and variance that are statistically equivalent to those of the classical mean-variance portfolio. Common portfolio constraints, such as transaction cost considerations and tax implications, can be handled efficiently in the robust optimization framework.

Robust optimization has also shown promise as a computationally attractive alternative to classical optimization methods when it comes to multi-period portfolio management. There are numerous benefits to taking a long-term view of investment management. Treating portfolio allocation as a multi-period problem provides a framework for robust overall portfolio management that takes into consideration the effects of rebalancing, transaction costs, future liabilities, and taxes.

By incorporating multi-period views on asset behavior in rebalancing models, portfolio managers may be able to reduce their transaction costs, as the portfolio will not be rebalanced unnecessarily often. As a simple example, if a portfolio manager expects asset returns to dip in the next time period but then recover, she may choose to hold onto the assets in her portfolio in order to minimize transaction costs. Or, if the net gain from realizing the tax loss is higher than the expense of the transactions, she may choose to trade for short-term benefit despite believing that the portfolio value will recover after two trading periods.

These trade-offs are complex to evaluate and model, and traditional optimization techniques for multi-stage optimization, such as dynamic programming (see, for example, Bertsekas [1995]) and stochastic programming (see, for example, Wallace and Ziemba [2005]), have not been very successful, as they result in computationally intractable problems due to the "curse of dimensionality." If future asset returns are treated as uncertain parameters, and the uncertainty in their estimates is modeled through appropriately chosen uncertainty sets, however, the resulting portfolio optimization formulations are computationally tractable.

We emphasize that while our focus has been on the application of robust optimization to portfolio construction, robust optimization is a powerful and general tool with financial applications that extend well beyond portfolio optimization. The robust optimization technique appears promising in enhancing current models for optimal trading, the computation of hedge ratios, the estimation of econometric models, and quantitative model selection—just to mention a few. The future may bring many more.

ENDNOTES

[1]See Fabozzi, Focardi, and Kolm [2006].

[2]By no means do we claim that this list is complete or that it has to be followed religiously. It is simply provided as a starting point for the quantitative portfolio manager.

REFERENCES

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

Bertsekas, D. *Dynamic Programming and Optimal Control*, Vols. I and II. Belmont, MA: Athena Scientific, 1995.

Bienstock, D. "Experiments with Robust Optimization." Presented at ISMP 2006, Rio de Janeiro, Brazil.

Black, F., and R. Litterman. "Global Portfolio Optimization." *Financial Analysts Journal*, 48 (1992), pp. 28-43.

Ceria, S., and R. Stubbs. "Incorporating Estimation Errors into Portfolio Selection: Robust Portfolio Construction." *Journal of Asset Management*, 7, 2 (July 2006), pp. 109-127.

El Ghaoui, L., and H. Lebret. "Robust Solutions to Least-Squares Problems with Uncertain Data." *SIAM Journal on Matrix Analysis and Applications,* 18 (1997), pp. 1035-1064.

El Ghaoui, L., M. Oks, and F. Oustry. "Worst-Case Value-at-Risk and Robust Portfolio Optimization: A Conic Optimization Approach." *Operations Research*, 51, 4 (2003), pp. 543-556.

Fabozzi, F., S. Focardi, and P. Kolm. *Financial Modeling of the Equity Market: From CAPM to Cointegration.* Hoboken, NJ: John Wiley & Sons, 2006.

Fabozzi, F., P. Kolm, D. Pachamanova, and S. Focardi. *Robust Portfolio Optimization and Management.* Hoboken: John Wiley & Sons, 2007.

Garlappi, L., R. Uppal, and T. Wang. "Portfolio Selection with Parameter and Model Uncertainty: A Multi-Prior Approach." *Review of Financial Studies,* 20, 1 (2007), pp. 41-81.

Goldfarb, D., and G. Iyengar. "Robust Portfolio Selection Problems." *Mathematics of Operations Research,* 28, 1 (2003), pp. 1-38.

Lee, J., D. Stefek, and A. Zhelenyak. "Robust Portfolio Optimization—A Closer Look." MSCI Barra Research Insights, June 2006.

Markowitz, H. "Portfolio Selection." *Journal of Finance,* March 1952, pp. 77-91.

Natarajan, K., D. Pachamanova, and M. Sim. "A Tractable Parametric Approach to Value-at-Risk Optimization." Working paper, 2005.

Scherer, B. "How Different is Robust Optimization Really?" Deutsche Asset Management, 2005.

Sims, C. A. "Pitfalls in a Minimax Approach to Model Uncertainty." *American Economic Review,* 91, 2 (2001), pp. 51-54.

Stubbs, R., and P. Vance. "Computing Return Estimation Error Matrices for Robust Optimization." Axioma, Inc., April 2005.

Wallace, S. and W. Ziemba, eds. *Applications of Stochastic Programming.* Philadelphia, PA: Cambridge University Press, 2005.