QS Investors

PortfolioChoice®: A New Approach to Portfolio Optimization

QS Investors Research Group

A half-century ago, a revolution got underway in the world of finance with the birth of modern portfolio theory. The ground-breaking work of economist Harry Markowitz provided a rigorous mathematical framework for modeling the risk-return characteristics of portfolios. However, in today's world, investor preferences are often not as simple as choosing a direct trade-off between expected return and volatility.

PortfolioChoice was designed to add new depth to the discipline of portfolio modeling and optimization. It is intended to define a portfolio's efficiency according to how well it meets the investor's objective, rather than simply its volatility and return characteristics. It also addresses a range of issues that fall outside the scope of traditional modeling and optimization techniques—issues such as incomplete manager data and non-normal returns, which are both common phenomenons in today's markets.

PortfolioChoice represents an important evolutionary step in practical finance that has been made possible by the recent advances in finance and finance related fields, such as statistics, econometrics and operations research. Ultimately, we believe that these advances and insights will help investors find better solutions to their investment challenges—and so better help them reach their financial goals.

Table of Contents

Introduction

Where it all started

A half-century ago, a revolution got underway in the world of finance with the birth of modern portfolio theory (MPT).

The ground-breaking work of economist Harry Markowitz provided a rigorous mathematical framework for modeling the risk-return characteristics of portfolios. This framework, known as mean-variance analysis, contained two innovative ideas. First, it characterized the investor's objective as a desire for high expected returns with low risk, where risk is measured by portfolio volatility. Second, it identified the investor's best, or optimal, choice of portfolio given this tolerance for risk. The model is summarized in Markowitz's famous efficient frontier, which shows the maximum expected return an investor can attain given the tolerable level of risk. The efficient frontier is now the standard expression of the risk-return trade-off inherent in investment decisions.

From this beginning, the fields of theoretical and practical finance blossomed. Building on Markowitz's work, the discipline of finance has been evolving rapidly, developing ever-more sophisticated investment theories and models, such as the Capital Asset Pricing Model (CAPM) and the Black-Scholes option-pricing formula, which have led in turn to revolutionary new investment concepts, such as benchmarking, and new investment products such as index funds and index options.

The continuing evolution of finance

We begin with a review of modern portfolio theory (MPT) to illustrate both the strengths and weaknesses of conventional quantitative approaches to portfolio decisions. These strengths and weaknesses spring from the same source: Markowitz chose a very special mathematic model of investor preferences, which had the advantage of a simple analysis but limited its flexibility and precision.

Given the limitations of computing power 50 years ago, this simplified approach made sense. But in the real world, investor preferences are often not as simple as a direct trade-off between expected return and volatility. In addition, the increasing sophistication of financial markets and the growing proliferation of asset managers and investment strategies have all contributed to the complexity of portfolio modeling and optimization—especially for investors with large, diversified portfolios.

The next generation of portfolio selection tools

PortfolioChoice was designed to address these issues and add new depth to the discipline of portfolio modeling and optimization.

In broad scope, it is a tool that is designed to quantify investor preferences and beliefs, and then identify the portfolio that best addresses those preferences and beliefs. One of the most significant advances allowed by PortfolioChoice is to define a portfolio's efficiency according to how well it meets the investor's objective, rather than simply its volatility and return characteristics.

PortfolioChoice seeks to quantify both the probability that any given portfolio would actually achieve the investors' objective, as well as what the expected shortfall might be. This approach allows the investor to see which specific elements of the portfolio may contribute to meeting the objective and which may help minimize the potential shortfall.

Thus, PortfolioChoice offers a dynamic approach to portfolio selection, taking into account how each investor resolves the tradeoff between seeking maximum probability of outperformance, versus minimizing the expected shortfall. This very flexible approach adds an important new dimension to the portfolio selection process—and can yield very different results.

Finally, the rigorous quantitative framework offered by PortfolioChoice is based on modern statistical and optimization techniques. These techniques can help investors address a range of issues that fall outside the scope of traditional modeling and optimization techniques—issues such as incomplete manager data and non-normal returns, which are both common phenomenon in today's markets. As a result, PortfolioChoice is designed to provide consistent and reasoned results, even in situations where conventional quantitative analysis practice may be of limited use.

PortfolioChoice represents an important evolutionary step in practical finance that has been made possible by the recent advances in finance and finance-related fields, such as statistics, econometrics and operations research.

Ultimately, we believe that these advances and insights will help investors find better solutions to their investment challenges—and so better help them reach their financial goals.

Section 1 | The Potential Benefits of a New Approach

PortfolioChoice re-defines portfolio optimality in new, more complex terms, with a goal of delivering more appropriate investment solutions. In doing so, however, it raises a number of issues that are difficult to address mathematically, and indeed, such computational complexities were beyond reach as recently as a decade ago. It's not that no one was asking these kinds of complex questions—they were—but statistical techniques and optimization tools available at that time were capable of handling only simplified approaches to issues that PortfolioChoice now addresses with relative ease:

Investor objectives and preferences

The conventional approach to designing an optimal portfolio is encapsulated in the question "How much volatility am I willing to tolerate in order to achieve a high expected return?" This is the heart of mean-variance analysis.

In reality, however, we know that different investors have different goals, and many investors would prefer to express these goals in relation to a target return—for example, to maximize the consistency or probability of meeting or exceeding a target, or to minimize the expected shortfall below the target. This leads us to formulate the portfolio decision according to new questions. For example: "How much expected shortfall am I willing to tolerate in order to achieve a high probability of reaching my target?"

The mean-variance approach is not capable of answering this more complex question, and based on PortfolioChoice analysis, we often find that the optimal portfolio in the mean-variance sense is not as effective in addressing this objective.

Views and beliefs

A second challenge posed by traditional mean-variance analysis is that its conclusions are based completely on historical data or implied current returns. Historical data, however, can be an unreliable predictor of future performance. Therefore, we believe that forward-looking forecasts are essential to the portfolio modeling and optimization process.

The use of views to complement historical data is necessary whenever the available data does not fully reflect such views. Examples might include an investor who takes a contrarian view on the possibility of short-term market volatility and interest rate movements, or the possibility that returns for a specific asset class may be higher or lower than its historical mean.

By adding investor views to the portfolio modeling and selection process, however, we are introducing an element of subjective judgement. And the ultimate success of the portfolio's performance will partly depend on the quality and accuracy of the views that were used in the selection process.

A key, novel feature of PortfolioChoice is its ability not only to blend historical data with investor views, but also to take account of our level of confidence in those views. This approach allows us to add an important element of judgement while also placing constraints on how much influence those views have on the portfolio selection process.

We believe that, overall, the ability to blend historical data with investor views is especially helpful in dealing with situations where strict quantitative analysis might not capture the entire story.

Statistical challenges

A final challenge to current standard quantitative practice is that real-world markets often raise complex statistical issues that are beyond the abilities of mean-variance to solve in a comprehensive way—issues such as incomplete data, non-normal returns, and uncertainty about estimates of statistical parameters (ie, "estimation risk"). Given the proliferation of asset managers and investment products, and the need to incorporate the investor's views into the investment decision, these complex statistical challenges are becoming increasingly common.

We believe that the technology and analytical techniques employed in PortfolioChoice deliver substantial advances in all of these areas. Its *Bayesian statistical analysis** (see below) is able to handle a wide range of issues that arise in real-world situations. Rather than side-stepping issues such as incomplete data and expressing the investor's true preferences and views, PortfolioChoice offers a powerful and flexible framework for addressing these issues explicitly—which enables us to navigate around such difficulties with confidence, and we believe, to design more insightful and appropriate investment solutions.

Bayesian statistical methods and quantitative investment decision-making

PortfolioChoice is based on Bayesian statistical methods, and is an approach that most people find more intuitive than traditional statistical techniques.

The key distinguishing feature of the Bayesian approach, as compared to the standard or "classical" approach, is that all uncertainty is expressed in terms of probabilities. Phrases such as "significance level," "confidence level," and "P-value," which the reader may recall from a college statistics course, arise from the requirement in classical statistics that we avoid the word "probability" in describing uncertainty about statistical parameters or "unknown constants."

A Bayesian statistician and a classical statistician may agree that the probability of heads in the toss of a coin is one-half. However, if the coin has already been tossed but the outcome has not yet been revealed, then the Bayesian will still say that the probability of heads is one-half, having no new information, while the classical statistician may insist that the outcome of the toss is now determined and so the concept of probability no longer applies. This example gives the flavor of the difference rather than the whole truth, but it serves to highlight how the classical approach imposes restrictions on the use of probabilities.

These restrictions are problematic in investment decision-making. For example, we may have a forecast that the mean excess return on stocks (that is, the equity risk premium) is 6%, but have some uncertainty about this. We may wish to hedge against the possibility that the forecast of the equity risk premium is incorrect. In the Bayesian framework, we can represent such uncertainty about parameters explicitly; this is not possible in the classical statistics framework. The uncertainty we have about parameters in a quantitative investment model is called estimation risk.

In the quantitative finance world, the classical approach to statistics is simply inadequate. Parameter values are uncertain, and we must explicitly account for that uncertainty in making decisions because it represents risk to the investor. The Bayesian approach enables us to do this.

*Throughout the paper, a term presented in italics indicates that it is defined in the glossary on page 22.

Section 2 | An Overview of the PortfolioChoice Framework

The generally accepted framework for quantitative investment decisions is based on two characteristics of investors: their *preferences* and their *beliefs*.[i] These are formalizations of natural human motivations: how a person feels about risk and reward (his or her preferences), and what he or she believes are the prospects for the investments under consideration (beliefs).

Like any software tool, PortfolioChoice takes inputs and processes them into outputs. There are two stages to any PortfolioChoice analysis, with the output from Stage 1 used as the input into Stage 2.

STAGE 1

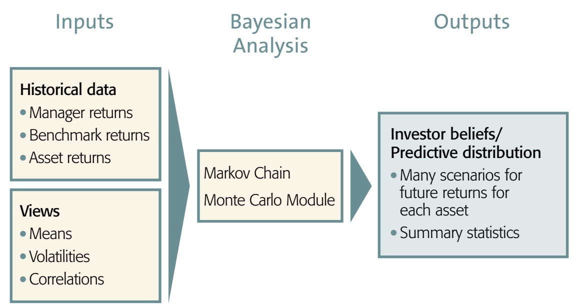

Stage 1 | Statistical analysis

In Stage One, illustrated in the diagram at right, PortfolioChoice quantifies the investor's beliefs by combining two inputs:

- Historical returns on the asset classes, managers and benchmarks under consideration
- *Views*—information or opinions about the potential future behavior of these assets

Using a form of Bayesian statistical analysis, called the Markov Chain Monte Carlo method (described on page 9), the investor's views are seamlessly combined with the historical returns. The output of this statistical analysis is a mathematical representation of the investor's overall beliefs, and is known as the *predictive distribution*.

The predictive distribution is simply a collection of possible scenarios for future returns—the predictive returns. It can be thought of as a compromise between the possible returns that could reasonably be expected based on the historical data and the performance expectations expressed in the investor's views. All of this information, from history and views, is synthesized in the predictive distribution. In a typical PortfolioChoice analysis, at least 20,000 predictive return scenarios are produced.

Stage 2 | Stochastic optimization

Once PortfolioChoice has produced the predictive returns, this information is combined with the investor's preferences in Stage 2, to identify the optimal portfolio—see diagram.

STAGE 2

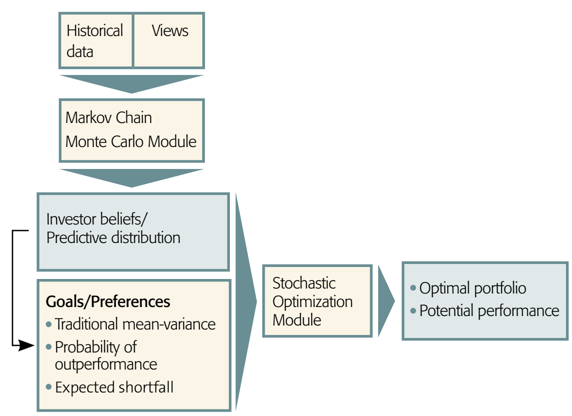

As discussed in Section 1, PortfolioChoice has the ability to accommodate many non-standard preferences. For example, an investor's goal may be to maximize the probability of beating the S&P 500 return by 1%. A different investor may want to minimize the expected shortfall below a fixed 7% return target. These goals, as well as many others, are expressed in PortfolioChoice in terms of the *outperformance probability* and/or *expected shortfall* relative to the specified target. To many investors, these goals are more intuitive and less narrow than the traditional Markowitz *mean-variance preferences*.

At this point we are ready to choose the optimal portfolio. PortfolioChoice uses a process called *Stochastic Optimization* to identify the portfolio that best addresses the investor's goals, based on the predictive distribution. (This process is also described on page 9.)

For example, if the investor's goal is to beat the S&P 500 benchmark by 100 basis points, PortfolioChoice's stochastic optimization module would identify the allocation that offers the largest probability of achieving this objective, based on the performance information contained in the predictive distribution.

We could also work with a wide range of other preferences or views. For example, the investor might want to see a range of options that describe the trade-off between the probability of outperforming the target and the expected shortfall below that target. In a typical client assignment, we will re-run the entire PortfolioChoice process multiple times, identifying optimal portfolios that correspond to various scenarios—e.g., preferences that range from conservative to aggressive, or views of market behavior that range from pessimistic to optimistic.

This detailed analysis enables us to understand the role played by each input in our decision-making process and, we believe, to arrive at a more informed and better designed investment solution.

One particular analysis of this type serves as the "flagship" of PortfolioChoice—the *probabilistic efficient frontier* (PEF). This is a chart much like the Markowitz efficient frontier, but exhibiting the trade-off between the outperformance probability and expected shortfall, rather than the trade-off between portfolio expected return and volatility.

How we crunch the numbers

Now that we have covered the general overview of how PortfolioChoice analysis works, and the benefits it offers in terms of making investment decisions, we can review the mathematics behind the process. This section is directed to readers with expertise in quantitative methods who want to know more about the algorithmic techniques that enable PortfolioChoice to deal with statistical and optimization models of great generality.

Markov Chain Monte Carlo algorithms

The rapid advances in Bayesian statistics of the past 30 years have been powered by parallel advances in computational techniques. First among these is the *Markov Chain Monte Carlo method*.

The *Monte Carlo method* was among the first computationally intensive algorithms to be implemented on computers. It is a method for computing probabilities or expected values using a random number generator. For example, to compute the expected shortfall below zero of a normal random variable, the Monte Carlo method provides an easy approach. We can simply generate a large number of normal random variables in a spreadsheet, compute the shortfall for each, and take the average. We could also use calculus or numerical integration, but the Monte Carlo approach requires no specialized mathematical knowledge. For many important problems in finance, engineering and physics, the Monte Carlo method also turns out to be the most effective computational approach.

The *Markov Chain Monte Carlo method* is a special Monte Carlo technique in which the random quantities generated are not independent. This dependence in the successive random quantities is the key variable to which the term “Markov Chain” refers. The advantage gained by introducing dependence is that it greatly increases the range of problems that can be solved by the Monte Carlo method. Without it, the Portfolio-Choice statistical analysis would be computationally infeasible.

Stochastic optimization

The Markowitz portfolio selection problem without constraints can be solved very easily. Even with constraints, the optimization problem is of a standard type and can be solved by efficient *quadratic programming* algorithms. Because of the generality of the PortfolioChoice model, there is no corresponding direct algorithm for dealing with the portfolio optimization problems that it generates.

PortfolioChoice again turns to a Monte Carlo technique to overcome this hurdle, this time with an optimization framework called Stochastic Optimization or Stochastic Programming. This framework was first developed in the 1950s and 1960s by George Dantzig, creator of the simplex algorithm for linear programming, to address multi-period portfolio selection problems that arose from Markowitz’s work.

To outline the PortfolioChoice stochastic optimization process, the portfolio selection problem is first formulated as a mixed integer linear program (MILP) based on Monte Carlo samples from the predictive distribution. The MILP is solved using a state-of-the-art, third-party software package. The process is then repeated to improve accuracy. Interestingly, solving minimum expected shortfall problems turns out to be relatively straightforward, because the resulting MILP formulations are simple linear programs. Problems involving probability of outperformance are much more difficult to solve because their formulations require integer variables.

Section 3 | Investor Preferences and Objectives

The benefits of PortfolioChoice's approach to objective setting go beyond simple intuitiveness—we have found that PortfolioChoice can potentially help investors arrive at better investment solutions.

As it turns out, the idea of stating an investor's preferences in terms of the probability of attaining a target is not new. A mere four months after Markowitz's 1952 paper appeared (see Markowitz 1991), Arthur Roy (Roy 1952, discussed in Bernstein 1992) of Cambridge University published an alternative approach to portfolio selection which treated preferences in just this way. But, because of the mathematical intractability of such preferences, it was Markowitz's work, and not Roy's, that became the foundation of modern portfolio theory.

Thanks to the computer revolution, we now have the ability to handle the mathematical challenges of Roy's preferences. Whereas Markowitz's preferences are characterized by a single *risk-aversion* parameter, PortfolioChoice allows a range of preferences related to attaining a return target. The extremes of this range are:

- To maximize the probability of reaching or outperforming a target, and
- To minimize the expected shortfall below the target

These extreme objectives are analogous to Markowitz's objectives of maximizing expected return and minimizing volatility. And just as in Markowitz's model, intermediate objectives are possible, expressed in PortfolioChoice as a progressive trade-off between maximizing outperformance probability while limiting expected shortfall. Graphically, this trade-off is shown in a similar kind of graph, which we call the "probabilistic efficient frontier" or PEF—see Exhibit 1.

To create this PEF chart, we took a hypothetical portfolio return target of 7.5% and calculated a range of portfolio options based on moderate views about potential market performance. At the right-hand end point, we see the portfolio with maximum probability of outperforming the 7.5% target, together with the expected shortfall associated with that portfolio. Then, as we move to the left, we place increasing emphasis on minimizing the expected shortfall (shortfall-aversion). When we reach the left-hand end point, we can see the minimum expected shortfall and the associated probability of outperformance.

How does this compare to a mean-variance efficient frontier? Well, in Exhibit 2 we see that the mean-variance efficient frontier suggests that there is only one "optimal" portfolio for every return target. But the PortfolioChoice efficient frontier identifies a wide range of choices based on how the investor chooses to resolve the tradeoff between avoiding a shortfall and seeking outperformance.

EXHIBIT 1*

Probabilistic Efficient Frontier
(7.5% return target)

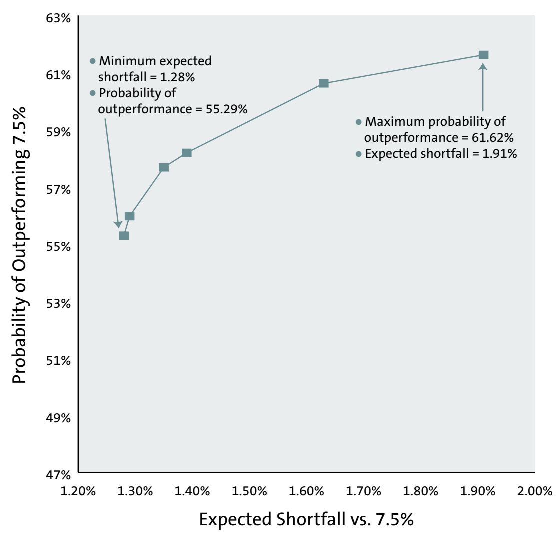

*This is a hypothetical example being used for illustrative purposes only. The scenarios noted are not exhaustive to the various outcomes to the markets. The portfolios in these examples were constructed with four asset classes: domestic equities, international equities, domestic bonds and international bonds. The indexes used to represent asset performance were, respectively, Russell 1000, MSCI EAFE Net, Barclays Capital Aggregate Bond Index, JP Morgan Nondollar Bond Index. Please see notes and disclosures on page 21 for more information.

The PortfolioChoice framework for stating the objective leads to a much different result. When we plot the mean-variance "optimal" portfolio for a 7.5% return against the frontier described by our PEF analysis, we see that the mean-variance portfolio falls well below the line. Not surprisingly, the mean-variance optimal portfolio is inefficient from the point of view of our probability objective.

EXHIBIT 2*

Efficient Frontier Comparison

Probabilistic Efficient Frontier
(75% return target)

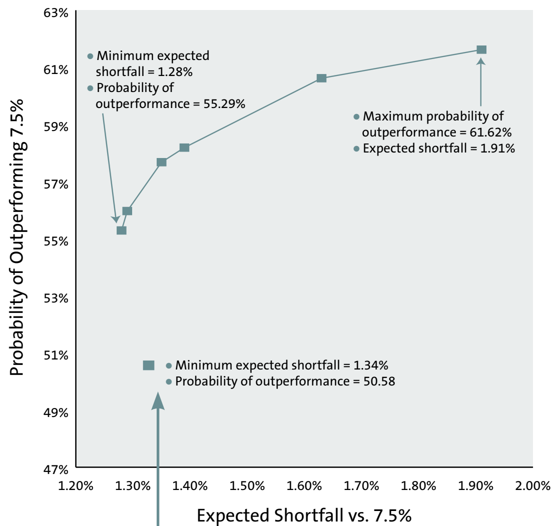

Mean-variance only identifies one "optimal" portfolio for our 7.5% return target. We can analyze this portfolio with PortfolioChoice to identify its outperformance probability and expected shortfall. When we plot these characteristics against the probabilistic efficient frontier, we find that the mean-variance portfolio is not efficient at either maximizing outperformance probability or minimizing shortfall.

Mean-Variance Efficient Frontier

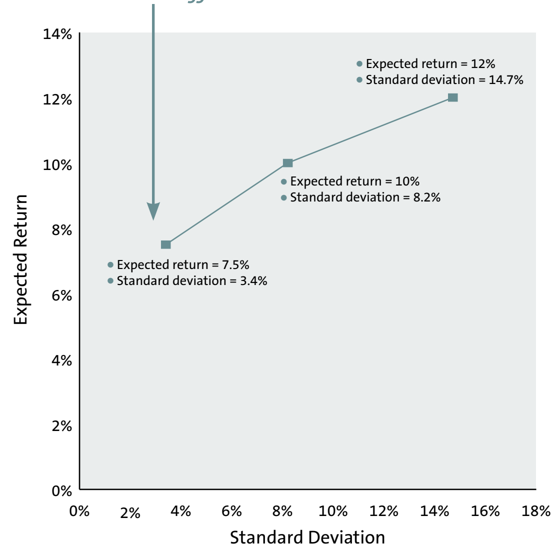

*This is a hypothetical example being used for illustrative purposes only. The scenarios noted are not exhaustive to the various outcomes to the markets. The portfolios in these examples were constructed with four asset classes: domestic equities, international equities, domestic bonds and international bonds. The indexes used to represent asset performance were, respectively, Russell 1000, MSCI EAFE Net, Barclays Capital Aggregate Bond Index, JP Morgan Nondollar Bond Index. Please see notes and disclosures on page 21 for more information.

Section 4 | The Investor’s Views and Beliefs

PortfolioChoice also makes important advances in our ability to express views and beliefs.

In broad terms, an investor’s beliefs are comprised of perceptions and expectations about future returns. These can be based on a wide variety of source material, such as: media reports, professional investment advice, personal experience, and anecdotes of friends and acquaintances. In the world of professional money management, we customarily develop our beliefs within a rigorous framework, drawing on such sources as historical and current financial and economic data.

Yet, one of the limitations of many traditional portfolio selection approaches is that they rely *solely* on historical data as quantitative inputs. As we know from the standard marketing disclaimer: “Past performance is no guarantee of future results,” how a manager has performed over the past three years may not be a good indication of the performance we can expect over the next year.

We believe that the portfolio selection process is greatly enhanced by tempering historical data with views—both quantitatively generated return forecasts as well as investor opinions. At QS Investors, we use return forecasts based on quantitative models of equilibrium in financial markets and in the larger economy. We can also incorporate investor opinions, as well as our level of confidence in those opinions. For example, consider a client who may have a strong opinion about a particular asset class—e.g., the likelihood of short-term interest rate movements and the resulting impact on fixed income securities. PortfolioChoice enables us to capture that important information in the portfolio selection process.

Consider the following example: the investor expects a return on US stocks over the next year to be some fixed value, say 6%. To express a somewhat weaker belief, that the return will be *close* to 6% rather than *equal* to 6%, we might declare that while the actual value is unknown, we believe that it has a normal distribution with a mean of 6% and standard deviation of .5%. Calling on some facts about the normal distribution, this means roughly that we are 95% sure that the expected return on US stocks will be between 5% and 7% over a one-year horizon. This is an example of how we can impose a view on mean return, but we can also impose views on volatilities and correlations.

This highly flexible framework enables us to bring information from any source to bear on the portfolio selection decision. The Bayesian statistical analysis then combines the diverse inputs to produce a unified output, the predictive distribution, which is customized to the specific market views of the investor.

Section 5 | Statistical Strengths of PortfolioChoice

PortfolioChoice has been designed to address a range of statistical challenges that are beyond the capabilities of standard approaches—challenges such as:

Non-normal return distributions

Suppose that the mean return on a particular investment over the next year is 8% and its volatility, or standard deviation, is 17%.

This information is incomplete as a description of the randomness of returns. It is not sufficient to tell us the probability of a 20% drawdown or the probability of meeting a target return of 10%, because the mean and volatility provide only a sketchy profile of randomness or uncertainty. Traditionally, this shortcoming of the mean and volatility as a description of the randomness of returns is met by further assuming that the distribution of returns is *normal*, or, in other words, that returns follow a *bell curve* (see figure 1 below).

FIGURE 1

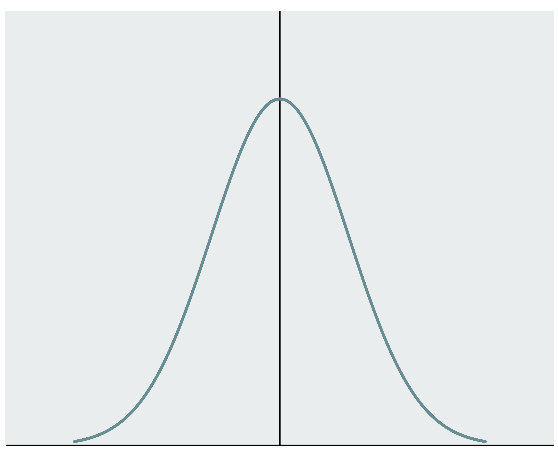

This normality assumption is generally treated as a "neutral" condition, expressing the idea that there is "nothing unusual" about the return distribution. Under the normality assumption, the probability of a 20% drawdown in the example here is precisely 4.98%, as may be verified by spreadsheet, calculator or normal statistical tables. The normal model of returns, like Markowitz's mean-variance analysis, has become popular because of simplifying mathematical properties.

The Long Term Capital Management (LTCM) crisis of late 1998 (Lowenstein 2000) was said to have been brought about by a 7 standard deviation event.[ii] The probability of such an event is approximately one in a trillion under the normal distribution. But with a suitable non-normal distribution, this probability could be as high as 1%.[iii] So either fate singled out LTCM for unfair treatment or LTCM's returns were non-normal.

The problematic feature of non-normal return distributions in investment decisions is that they may produce large negative returns more frequently than a normal distribution of the same mean and volatility. And in today's investment markets, investors may encounter a wide range of securities with potentially non-normal returns. For this reason, we believe that non-normality should be dealt with explicitly in the portfolio selection decision.

Here are some common situations in which non-normal distributions arise:

- Certain types of hedge funds may carry a somewhat larger likelihood of a substantial drawdown than is possible under a normal distribution (see Fung and Hsieh 1999). This characteristic manifests in the histogram of the returns as a skew to the left.

FIGURE 2

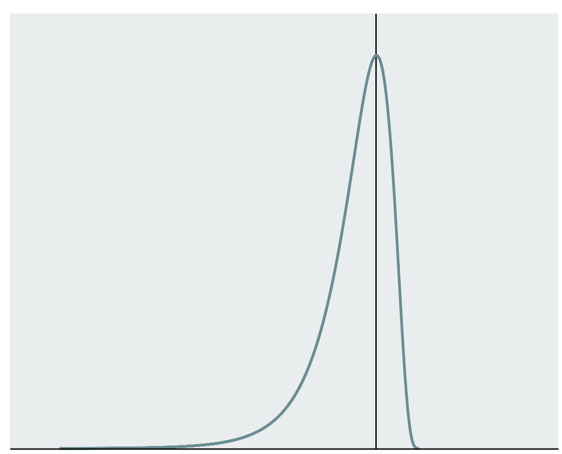

- The payoff of a call option at the expiration date is clearly non-normal: it is typically like a bell curve with the portion to the left of the strike price removed—the histogram *is skewed to the right*.

FIGURE 3

- Closer to home, equity returns over a long horizon tend to follow a distribution called the *log normal distribution*[iv]—which again gives a histogram that is skewed to the right.

FIGURE 4

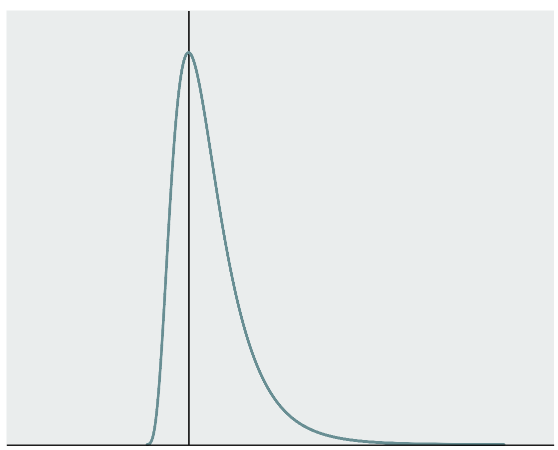

Non-normal distributions are challenging to deal with, both in the statistical analysis of historical return data and also in solving the portfolio selection problem. One of the key features of PortfolioChoice is its ability to handle beliefs that returns follow non-normal distributions. In fact, in PortfolioChoice beliefs are virtually always non-normal to some extent.

Accounting for estimation risk

When we invest quantitatively, we must devise a model of future returns, often relying on statistical parameters such as means and volatilities. Typically, these are known only approximately, and this uncertainty poses a risk to investors that we call *estimation risk* (Klein and Bawa 1976, Jobson and Korkie 1980). The term refers to the fact that the unknown parameters may have to be estimated. In most standard approaches to quantitative analysis, estimation risk is simply ignored and estimates of parameters are treated as if they were the exact values.

The PortfolioChoice framework explicitly takes account of estimation risk, in that the predictive distribution of returns incorporates the investor's uncertainty about parameters. As a result, PortfolioChoice generally attributes higher risk to asset returns than the more standard quantitative approaches to portfolio selection, consequently it recognizes that assets for which we have only a short history are more risky than those for which we have a long history simply because we don't know them as well. This is one of the key benefits of the Bayesian statistical analysis in PortfolioChoice.

Dealing with differing start dates and missing values

It often happens that we do not have the same length of return history on all of our managers, assets or benchmarks. We may have one manager who has run the same active product for 10 years, while another has been in business for only two years. Performing a reasonable quantitative analysis in situations like this is technically quite challenging.

One common heuristic approach is simply to truncate all data to match the shortest available manager history—all managers will then have histories of the same length. This means throwing away valuable information on managers with longer track records. It may also mean throwing away valuable information about different regimes of the business cycle.

Another approach is to extend a manager's short history by substituting a suitable proxy for missing returns; for example, the proxy may be the manager's benchmark or the returns of a manager who follows a similar strategy. But using a proxy may underestimate the manager's active risk or otherwise distort active returns. This heuristic also leads us to make decisions as if we know the manager well through a long track record, when in fact we do not, because it ignores the estimation risk implicit in a short track record.

These problems are resolved in Portfolio-Choice using Bayesian statistical analysis—specifically, 'data imputation,' which is explained in the box below.

'Filling in missing values' through data imputation

Computing the means and covariance matrix of historical returns is a basic skill in quantitative analysis. But when there are missing values, or when the return series have different start dates, this becomes a challenging problem.

For general incomplete data, there is no explicit formula available, and the task requires heavy-duty computational power. In fact, it is so burdensome that it is traditionally handled by heuristics like truncating the data, or extending a manager's returns using a suitable benchmark. While these heuristic approaches may be appropriate on occasion, a decision as important as asset allocation requires a much more rigorous methodology.

Data imputation is a general approach to solving the problem of incomplete data (see Tanner 1996). The key idea is that the missing values are 'filled in' through a Monte Carlo simulation. The random data generated by the simulation follows a precise distribution that reflects not only the relationships between the missing values and the available data, but also their relationships with any parameters in the underlying statistical model. Once the missing values are filled in, the mean and covariance matrix can be computed in the usual way. However, this calculation itself provides us with new information about the unknown parameters. We then replace the filled-in missing values using a fresh Monte Carlo simulation that exploits our new information. The process continues iteratively, and ultimately leads to the means and covariance matrix. (This iterative process is in fact an example of the Markov Chain Monte Carlo method.)

Oversensitivity to views

Here is another well-documented pitfall in portfolio selection—heuristic approaches can overreact to views and provide inappropriate allocation recommendations.

For example, consider the decision to add small-cap stocks to a large-cap portfolio. Suppose we feel that small stocks will outperform large stocks by 1% over the next year.[v] We wish to add small caps to our portfolio to capture their perceived better prospects, in a risk-adjusted way.

The difficulty we run into is that, because of the high correlation between small and large cap stocks, a naive analysis will lead to a surprisingly large weight in small caps in response to the view. The optimizer perceives a near arbitrage. When using PortfolioChoice, the impact of the views is tempered by correlation: the greater the correlation between two assets, the more a view on the spread between them is reigned in. In this way, Portfolio-Choice balances history and views according to the strength of the information they contain. Example C in Section 6 demonstrates this feature, which Portfolio-Choice shares with the Black-Litterman method (Black and Litterman 1991; Bevan and Winkelmann 1998).

Because PortfolioChoice is based on a rigorous methodology it is not stymied by pitfalls such as these difficulties in handling views, and no special intervention is needed to handle them.

Section 6 | Examples*

This section provides three examples of how PortfolioChoice deals with the key issues discussed in this paper.

Example A | The importance of making full use of the data

This example illustrates the importance of the statistical approaches used to handle data.

Imagine it is 1999, after four years of strong performance for growth stocks, and we are comparing two managers, a growth manager with a short history and a value manager with a long history. How do we compare them fairly? (In this example we use the Russell 1000 Value and Growth indices to represent our value and growth managers.)

- Do we truncate the longer history, which would throw away valuable information?
- Do we extend the short history? If so, what do we use as a reasonable proxy?

FIGURE 5

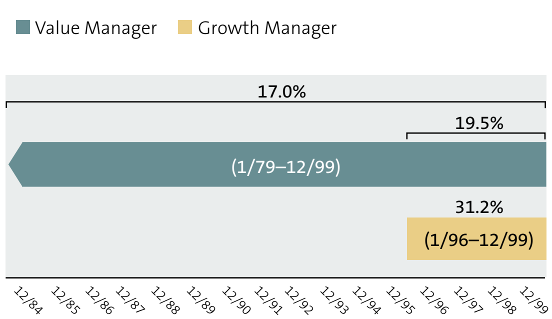

If we truncate the longer history to the same length as the shorter history, the value manager's 19.5% track record during that period appears significantly weaker than the growth manager's 31.2%. And it would lead a traditional optimizer to grossly overweight the growth manager—in this case giving the growth index a weight of 100%.

FIGURE 6

Mean-variance portfolio using truncated history

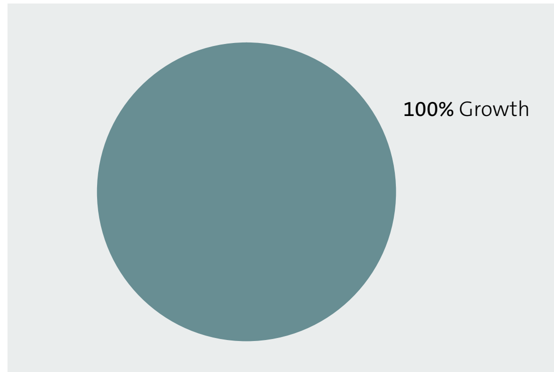

FIGURE 7

PortfolioChoice portfolio using longer history

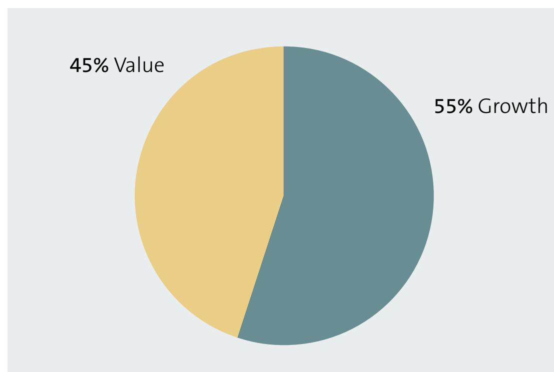

PortfolioChoice, however, estimates the missing history by data imputation. Estimates are constructed using the growth manager's actual correlation to the benchmark and the value manager. Based on this extrapolation, PortfolioChoice estimates that the growth manager's track record over a much longer history would likely have been closer to 19%—not nearly as high as the 31% turned in during its short history at the height of the growth market. Thus, PortfolioChoice gives us a much more realistic estimate of the relative value of these two strategies and hence, a much more realistic allocation.

*Please see page 21 for important notes and disclosures.

The bar chart below shows that in the ensuing two years, the PortfolioChoice allocation to these two indices would likely have outperformed the growth index alone. While this example is extreme (it is unlikely that an investor would make an entire large cap allocation to the growth category), an inappropriately high allocation to growth in a situation such as this could result in a drag on overall portfolio performance.

FIGURE 8

Historical excess return 2000–2002 (annualized)

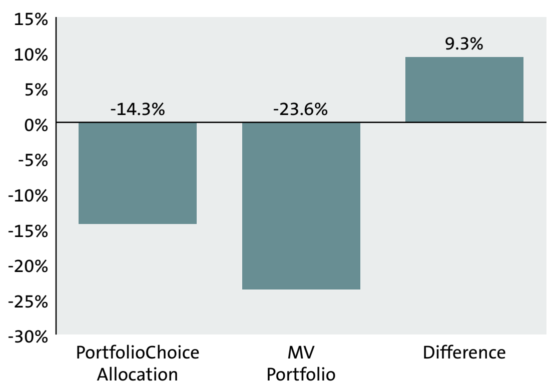

Example B | 'Views' are important, too

To meet the challenges of portfolio selection, it is often not enough to make full use of the available data, your views about future returns can be equally important—an issue we focus on in this example.

First, to simplify, we will work with a two-asset universe consisting of the S&P 500 index, with almost 13 years of historical data, and an active equity fund (that is highly correlated to the S&P 500), with available history of only 19 months.

FIGURE 9

	No. of Data Points (monthly returns)	Average Annual Return During This Period
S&P 500 Index	154	10.6%
Active Equity Fund	19	-16.7%

If we look at the returns of the two assets, it would appear that the fund has vastly underperformed the index. In reality, over the 19 months it has been in existence, the fund has outperformed the index on an annualized basis. We can use the data imputation process that was illustrated in Example A to create a predictive distribution for this manager and compare it to that of the index. Because there is a high correlation between the manager and the index, and we are predicting a positive return for the index, it follows that the manager's return will also be positive. In addition, because the manager has historically outperformed the index, a return forecast based purely on historical data will show the manager continuing to outperform the benchmark on an annualized basis—as in the bar chart below.

FIGURE 10

Historic returns vs. predictive returns

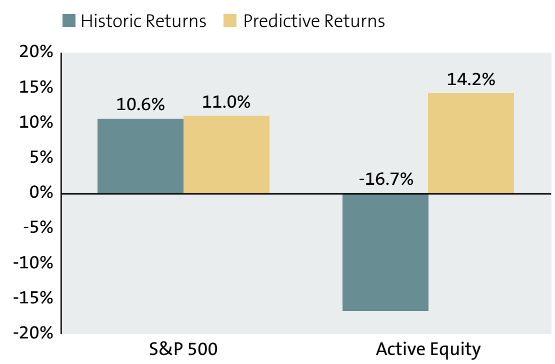

But suppose that, on a one-year view, the investor believes that these assets will underperform their historical means, with the S&P 500 index returning only 7.5% on an annualized basis, and the active equity fund returning 9%. And let us also say that we have a low level of confidence in these views.

FIGURE 10

	No. of Data Points (monthly returns)	Average Annual Return During This Period	View on Potential Return	Confidence in View
S&P 500 Index	154	10.6%	7.5%	Low
Active Equity Fund	19	-16.7%	9.0%	Low

In the following chart, the investor's views for index and manager returns are represented by the grey bars. PortfolioChoice blends the view with the historical performance, but since we have specified the investor's low confidence level in these views, the historical performance gets a heavier weight, and so the predicted return for each asset (green bar) is still pretty close to the original forecast (yellow bar), which was based purely on historical performance (i.e., no view on potential performance).

FIGURE 12

Annual returns: Prior vs. historic vs. predictive

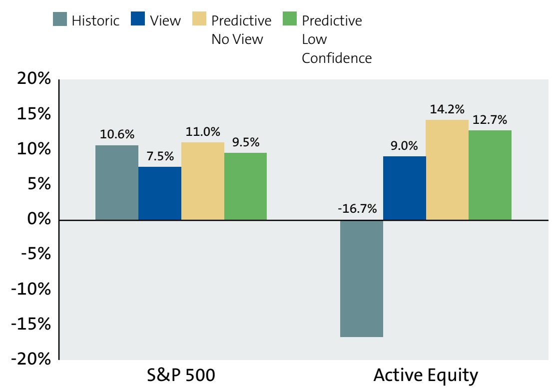

But now, let us increase our confidence in our investor's view for reduced returns. We see that increasing our confidence in the view from low, to medium, to high (the three bars at the right of each group) has a dramatic impact on the predicted return, which decreases away from the forecast based purely on history (yellow bar) until it almost equals the forecast based purely on the investor's view (dark blue bar).

FIGURE 13

Annual returns: Prior vs. historic vs. predictive

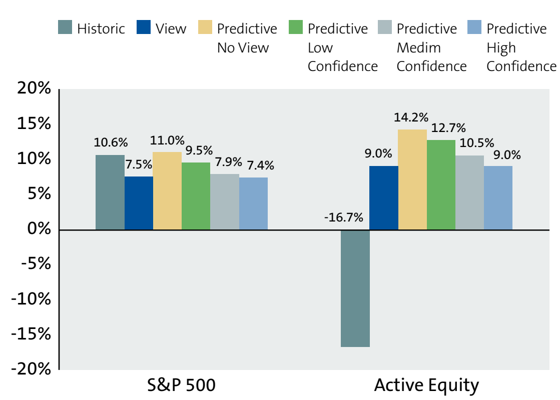

So it is not only important to get the data right. If one has a view about future returns, it is critical to be able to incorporate that into the portfolio selection process as well, as it could have a substantial impact on portfolio allocations—to asset classes and/or individual managers within asset classes. The success of the portfolio selection process, of course, will depend in part on the quality and accuracy of the views, and so we believe that the ability to impose constraints on them (by stating a confidence level) is also an important part of the process.

Example C | Mean-variance over-sensitivity to inputs

This last example shows that even where underlying issues with data and views have been addressed, mean-variance analysis can still deliver an inappropriate result, based on its over-sensitivity to changes in such inputs.

Consider an investor who wants to allocate money to small cap stocks—for the sake of simplicity the allocation is being made away from large cap. The historical mean returns for large cap and small cap stocks are highly correlated, and over the last 15 years are:

- S&P 500 = 11.4%
- Russell 2000 = 10.4%

The investor, however, is making the small cap allocation based on the view that small cap stocks will outperform their historical mean, returning 12.4% annually for the next few years. We incorporate the view into both PortfolioChoice and mean-variance and we get allocations that look like this, with mean-variance allocating a significantly higher portion of the portfolio to small cap:

FIGURE 14

	PortfolioChoice	Mean-Variance
Russell 2000 View	12.4%	12.4%
S&P 500 Allocation	89.2%	71.9%
Russell 2000 Allocation	10.8%	28.1%

The reason that mean-variance suggests a higher allocation to small cap is that it takes a naive approach to analyzing the view. Remember that large cap and small cap are highly correlated, so if you think one will outperform its historical mean, then the other should also outperform. Mean-variance, however, naively assumes that the view only applies to small cap: that small cap will outperform, while the S&P will perform at its historical mean, which runs contrary to the historical correlation between these assets.

PortfolioChoice, by incorporating correlation into the analysis, raises the forecast for large cap in line with the bullish view on small cap and makes a smaller allocation to small cap as a result.

And what if we are even more optimistic about small cap returns, increasing our estimate another 1%... to 13.4%? Traditional mean-variance techniques are not only naive in the way they ignore correlations, they are also overly sensitive to changes in views.

In the table below we see that, within the mean-variance portfolio, a 13.4% return forecast results in a huge jump in the small cap allocation, while the large cap allocation plummets. By contrast, we believe Portfolio-Choice provides a more realistic allocation, even with an optimistic view:

FIGURE 15

	PortfolioChoice		Mean-Variance	
R2000 View	12.4%	13.4%	12.4%	13.4%
S&P 500 Allocation	89.2%	83.4%	71.9%	56.8%
R2000 Allocation	10.8%	16.6%	28.1%	43.2%

From these examples, we see that it is important to:

- Make full use of the data
- Be able to incorporate views about future returns
- Analyze the data correctly

PortfolioChoice makes significant advances in all these areas, and as a result, we believe it can deliver better solutions that more effectively address each investor's unique objectives.

Section 7 | PortfolioChoice—A New Approach to Portfolio Optimization

PortfolioChoice takes advantage of many recent advances, in such areas as computing power, statistical analysis and practical finance, to enhance the discipline of portfolio selection.

Investors can now make portfolio decisions according to a new set of criteria, based on the likelihood that the portfolio will actually achieve the objective, and in the event of a shortfall, how much it might be. In addition, PortfolioChoice makes important advances in our ability to address complex statistical challenges that are becoming much more common in today's sophisticated, volatile markets. Finally, PortfolioChoice allows us to take into account a wide range of beliefs and opinions, including both the investor's outlook for market performance and level of confidence in the return forecasts on which the allocation is based.

This approach represents a powerful new resource in the investment decision-making process.

Rather than simply trying to maximize the relationship of risk and return, each investor can now understand what portfolio characteristics may enhance the probability of meeting the investment objective and then manage, explicitly, the underlying dynamics of the portfolio selection process.

For these reasons, we believe that PortfolioChoice delivers important advances in understanding, technique and potential outcome, all to the benefit of today's investor.

Section 8 | Important Notes and Disclosures

Information, case studies, and examples cited in this document are for illustrative purposes only and are being provided solely as illustrations of the PortfolioChoice selection process and approach. Any quoted PortfolioChoice performance is an estimate that is relevant only to the specified time period and return assumptions, compared with portfolios based on historical extrapolation. The data does not reflect actual performance of any single PortfolioChoice portfolio, nor was a contemporaneous investment model run. Simulated performance results have inherent limitations. Unlike an actual performance record, simulated results do not represent actual trading and are subject to the fact that they are designed with the benefit of hindsight. Also, since the trades have not actually been made, the results may not reflect the impact that certain material economic and market factors might have had on an investment adviser's actual decision-making. Therefore, performance numbers used in illustrations are not necessarily indicative of the results you would obtain as a client of QS Investors, and no representation is being made that these or similar results are guaranteed. Results are generally based on security selection, client investment restrictions (if any) market economic conditions and other factors which would all influence portfolio returns.

Further, investment in international markets can be affected by a host of factors, including political or social considerations, diplomatic relations, limitations or removal of funds or assets or imposition of (or change in) exchange control or tax regulations in such markets. Additionally, investments denominated in an alternative currency will be subject to changes in exchange rates that may have an adverse effect on the value, price or income of the investment. The value of investments and income arising therefrom can fall as well as rise, and no assurance can be given that the investment objectives will be met or that an investor will receive a return of all or part of his or her investment.

[i] Here is how Campbell and Viceira (2002) express the key ideas in describing the challenge of quantitative portfolio selection. "... individuals and institutions must first think systematically about their preferences and the constraints they face. They must form beliefs about the future... Then they must combine preferences, constraints, and beliefs to form optimal portfolios ..."

[ii] Jorion (2000) describes the LTCM crisis as "an 8.3 standard deviation event," while Fung and Hsieh give it only 6 standard deviations. Alan B. Krueger, quoted in "Exposing the Fraying Edges in the Fabric of the Economy" (*New York Times*, Dec. 18, 2000), raises the bid to 9 standard deviations. The number depends on the data you use to characterize the crisis. We chose 7 standard deviations to make our subsequent calculations more easily remembered.

[iii] This value comes from a result known as Chebyshev's inequality, which tells us that the probability that a random quantity falls within 7 standard deviations of its mean is at least 1-1/49 or 2.04%. In the case of a symmetric distribution, the largest possible probability for a 7 standard deviation drawdown is 2.04%/2, or about 1%.

[iv] Returns follow a log-normal distribution whenever geometric returns follow a normal distribution.

[v] We take this to mean that our forecast of the expected return on the spread between small and large cap stocks over the next year is 1%.

Section 9 | Glossary

Beliefs
The investor's perception of the possible future returns on assets. Beliefs are represented mathematically by the predictive distribution in PortfolioChoice.

Bayesian statistics
A standard statistical methodology in which all uncertainty is expressed in terms of probabilities.

Distribution
The possible values of a random quantity, or quantities, and their probabilities.

Efficient frontier
The collection of portfolios that are optimal in Markowitz's sense. These are the portfolios that maximize the mean return for a given volatility.

Estimation risk
The risk to the investor due to uncertainty about the true values of statistical parameters, such as expected returns or volatilities.

Monte Carlo simulation
Computational algorithms based on computer-generated random numbers.

Markov chain
A sequence of random quantities whose inter-dependence is defined as follows: a given quantity in the sequence may be dependent on its predecessor, it has no further dependence on any earlier members of the sequence. [Also see sidebar on page 9.]

Markov Chain Monte Carlo method
A Monte Carlo method using random quantities that follow a Markov Chain, rather than the independent random quantities usually used in the Monte Carlo method.

Mean-variance analysis
Portfolio selection using Markowitz preferences.

Normal distribution
The standard distribution of returns, which implies, for example, that the return is within 1 standard deviation of the mean with probability .68 and within 2 standard deviations with probability .95. These relationships between probabilities, means, and standard deviations do not hold for other distributions. The normal distribution has many attractive mathematical properties and also describes many empirical phenomena as well, including returns on many financial assets.

Non-normal distribution
Distributions that are not normal. Skewed and heavy-tailed distributions are typical examples.

Parameter
A quantity describing a distribution, such as an expected return or volatility. The value is typically unknown.

Probabilistic efficient frontier
The portfolios that maximize the probability of outperforming the investor's target for a given expected shortfall.

Predictive distribution
The distribution of future returns given the historical returns and our views. It is the mathematical representation of the investor's beliefs.

Preferences
A mathematical representation of the investor's goals.

Prior distribution
A mathematical representation of the investor's views.

Stochastic optimization
The process of maximizing a probability or expected value, usually using the Monte Carlo method.

Views
The investor's special knowledge or opinions about future returns. These comprise one of the two ingredients in the formation of beliefs in PortfolioChoice, the other being historical returns.

Section 10 | References

Bernstein, Peter (1992) *Capital Ideas: The Improbable Origins of Modern Wall Street*. The Free Press, New York.

Bevan, A., and Winkelmann, K. (1998) Using the Black-Litterman Global Asset Allocation model: Three years of practical experience. *Goldman-Sachs Fixed Income Research*. June 1998.

Birge, J.R., and Lepingle, F. (1997) Introduction to Stochastic Programming. Springer-Verlag. New York.

Black, F., and Litterman, R. (1991) Asset Allocation: combining investor views with market equilibrium. *Journal of Fixed Income* 1(2), 7-18 1991.

Campbell, J.Y., and Viciera, L.M. (2002) *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press.

Fung, William, and David Hsieh (1999) "A Primer on Hedge Funds," with William Fung, *Journal of Empirical Finance*, 6 (1999), 309-331.

Jorion, P. (2000) Risk Management Lessons from Long-Term Capital Management, *European Financial Management* 6, 277–300.

Jobson, J.D., and Korkie, R. Estimation of Markowitz efficient portfolios. *Journal of the American Statistical Association* 75, 544–554. 1980.

Klein, R.W., and Bawa, V.S. (1976) The effect of estimation risk on portfolio optimal choice. *Journal of Financial Economics* 3, 215-231. 1976.

Lowenstein, R. (2000) *When Genius Failed*. Random House.

Markowitz, H.M. (1991) *Portfolio Selection*. Second Edition. Blackwell. Malden, Mass.

Roy, A.D. (1952) Safety First in Holding Assets. *Econometrica* 20, 431–439.

Stambaugh, R.F. (1997) Analyzing investments whose histories differ in length. *Journal of Financial Economics* 45, 285–331. 1997.

Stigler, Stephen M. (1990) *The History of Statistics: The Measurement of Uncertainty before 1900*. Harvard University Press.

Tanner, M.A. (1996) Tools for Statistical Inference. Springer. New York.

J.P. Morgan Non-Dollar Bond Index is an unmanaged, total return, trade-weighted index of over 360 government and high-grade bonds in 12 developed countries.

Barclays Capital Aggregate Bond Index is an unmanaged index representing domestic taxable investment grade bonds, with index components for government and corporate securities, mortgage passthrough securities, and asset-backed securities with average maturities of one year or more.

MSCI EAFE (Net) Index is an unmanaged, market capitalization weighted index composed of companies representative of the market structure of developed market countries in Europe, Australasia and the Far East, calculated with dividends net of foreign taxes reinvested.

Russell 1000 Index is an unmanaged index that measures the performance of the 1,000 largest companies in the Russell 3000 Index, which measures the performance of the 3,000 largest US companies based on total market capitalization. The Russell 1000 Index represents approximately 92% of the total market capitalization of the Russell 3000 Index.

Russell 2000 Index measures the performance of the small-cap segment of the U.S. equity universe. The Russell 2000 Index is a subset of the Russell 3000 Index representing approximately 8% of the total market capitalization of that index. It includes approximately 2,000 of the smallest securities based on a combination of their market cap and current index membership.

S&P 500 Index includes 500 leading companies in leading industries of the U.S. economy, capturing 75% coverage of U.S. equities.

Although the information contained in this document has been obtained from sources believed to be reliable, we do not guarantee its accuracy, completeness or fairness, and it should not be relied upon as such. Opinions and estimates, including forecasts of conditions, reflect our judgement as of the date of this document and are subject to change without notice. Such opinions and estimates, including forecasts of conditions, involve a number of assumptions that may not prove valid.

Performance tables and charts herein do not reflect the deduction of investment management fees. In the event that such investment management fees and other fees were deducted, the performance of an account would be lower. For example, if an account appreciated by 10% a year for five years, the total annualized return for five years prior to deducting fees at the end of the five-year period would be 10%. If total account fees were .10% for each of the five years, the total annualized return of the account for five years at the end of the five-year period would be 9.89%.

Past performance is not indicative of future results.

This document may not be reproduced, circulated or used for any purpose without our written consent. The manner of circulation and distribution of this presentation may be restricted by law or regulation in certain countries, including the United States. Persons into whose possession this presentation may come are required to inform themselves of, and to observe such restrictions. An investment is not a deposit and is not insured by any government agency or by QS Investors.

This material is intended for informational purposes only and it is not intended that it be relied on to make any investment decision. It was prepared without regard to the specific objectives, financial situation or needs of any particular person who may receive it. It does not constitute investment advice or a recommendation or an offer or solicitation and is not the basis for any contract to purchase or sell any security or other instrument, or QS Investors, LLC to enter into or arrange any type of transaction as a consequence of any information contained herein. QS Investors, LLC does not give any warranty as to the accuracy, reliability or completeness of information which is contained in this document. Except insofar as liability under any statute cannot be excluded, no member of QS Investors, LLC, the Issuer or any officer, employee or associate of them accepts any liability (whether arising in contract, in tort or negligence or otherwise) for any error or omission in this document or for any resulting loss or damage whether direct, indirect, consequential or otherwise suffered by the recipient of this document or any other person.

The views expressed in this document constitute QS Investors' judgment at the time of issue and are subject to change. The value of shares/units and their derived income may fall as well as rise. Past performance or any prediction or forecast is not indicative of future results. This document is only for professional investors. Investments are subject to risks, including possible loss of principal amount invested.

QSI-00037 (9/10)