

Developing Robust Asset Allocations¹

Working Paper

First Version: February 17, 2006

Current Version: April 18, 2006

Thomas M. Idzorek, CFA
Director of Research
Ibbotson Associates
225 North Michigan Avenue
Suite 700
Chicago, Illinois 60601-7676
312-616-1620 (Main)
312-616-0404 (Fax)
tidzorek@ibbotson.com

Abstract

Over the last 50 years, Markowitz's mean-variance optimization framework has become the asset allocation model of choice. Unfortunately the model often leads to highly concentrated asset allocations, the primary reason that practitioners haven't fully embraced this Nobel Prize winning idea. Two relatively new techniques that help practitioners develop robust, well-diversified asset allocations are the Black-Litterman model and resampled mean-variance optimization. The first approach focuses on building capital market expectations that behave better within an optimizer while the second approach is an attempt to build a better optimizer. In addition to providing practitioner friendly overviews of the two approaches, this article contributes to the literature by comparing / contrasting empirical examples of both approaches as well as the first empirical example of how the Black-Litterman model and resampled mean-variance optimization can be used together to develop robust asset allocations.

Key Words: Robust asset allocation, mean-variance optimization, Black-Litterman, resampling.

Introduction

In their seminal and extremely influential work, Brinson, Hood, and Beebower [1986] estimates that *over time* 90% of the *variance* in returns of a typical portfolio is explained by the variance of the portfolio's asset allocation. Ibbotson and Kaplan [2000], among others, confirms this important finding supporting the notion that strategic asset allocation (SAA) is the most important decision in the investment process.

Strategic asset allocation is both a process and a result. The strategic asset allocation focuses on how to invest assets to maximize the probability of achieving one's long-term goals at an appropriate level of risk. It is the process of determining the target long-term allocations to the available asset classes. The process results in a set of long-term target allocations to applicable investable asset classes (proxied by market indices). The resulting long-term target asset allocations are often formalized into a strategic policy benchmark (policy benchmark for short) or model asset allocation. Using today's popular alpha-beta vernacular, strategic asset allocation is the beta decision, and as such, investment vehicles like mutual funds and hedge funds are not part of the discussion.²

The most widely used quantitative strategic asset allocation framework is Harry Markowitz's mean-variance optimization, an idea that resulted in a Nobel Prize for Markowitz (see Markowitz [1952, 1959]). Mean-variance optimization is one of the cornerstones of modern portfolio theory and over the last half century has become the dominant asset allocation model. The procedure maximizes expected return for a given level of risk, or equivalently, minimizes risk for a given return. The traditional textbook approach to asset allocation is as follows.

Step 1 – Estimate the returns, risks, and correlations of the relevant asset classes based on capital market conditions.

Step 2 – Use mean-variance optimization to create an efficient frontier.

Step 3 – Select a point on the efficient frontier or select a mix of the risk-free asset and the optimum risky asset allocation based on an estimated risk tolerance level.

By almost all accounts, the maximization of return per unit of risk is a logical and worthwhile objective. However, the Markowitz framework may be too powerful for its own good. Common issues or criticisms of traditional mean-variance optimization include:

1. Mean-variance optimization usually leads to asset allocations in which the majority of the holdings are concentrated in a small number of asset classes that make up the opportunity set, contradicting the common-sense notion of diversification;
2. If the assets exist to help meet a liability, the liability should be considered in the process;
3. Basing one's decision solely on an asset allocation's mean and variance is insufficient to base one's decisions, in a world in which asset class returns are not normally distributed; and,
4. Most investors have multi-period objectives and the mean-variance framework is a single period model.

These potential shortcomings are the likely reasons that practitioners have not fully embraced mean-variance optimization. For a number of practitioners, mean-variance optimization creates the illusion of quantitative sophistication; yet, in practice, asset allocations are developed using judgmental, ad hoc approaches. Recent advances significantly improve the quality of typical mean-variance optimization-based asset allocations that should allow a far wider audience to realize the benefits of the Markowitz paradigm, or at least the intent of the paradigm.

In this article, we focus on the first issue: the lack of diversification that can result from traditional mean-variance optimization. We begin with two examples in which traditional mean-variance optimization

leads to extreme asset allocations. These examples highlight the sensitivity of the output (the asset allocations) to changes in the inputs (the capital market assumptions). We examine the causes and possible solutions to highly concentrated asset allocations. We focus on two relatively new and exciting techniques that can lead to robust forward-looking asset allocations: the Black-Litterman asset allocation model and resampled mean-variance optimization.³ The Black-Litterman model and resampled mean-variance optimization are very different; yet, they both help overcome the problem of highly concentrated asset allocations. Building on the initial example, we apply both of the robust asset allocation approaches individually, prior to demonstrating how the Black-Litterman model and resampled mean-variance optimization can be used together.

Highly Concentrated Undiversified Asset Allocations

Traditional mean-variance optimization or *traditional* MVO, as we will call it, often leads to highly concentrated, undiversified asset allocations. To gain a better feel for this, let's look at a typical example using an opportunity set that includes nine asset classes: US Large Cap Growth, US Large Cap Value, US Small Cap Growth, US Small Cap Value, International Stocks, US Bonds, International Bonds, Commodities, and Cash. When developing an opportunity set, one should select non-overlapping, mutually exclusive asset classes that reflect the investor's investable universe. The development of a robust opportunity set is a critical step in the strategic asset allocation process; yet, all too often, its importance is overlooked. In general, investors should be encouraged to expand their opportunity sets to include the asset classes that make up the hypothetical, all-encompassing market portfolio of the Sharpe-Lintner-Mossin-Treynor Capital Asset Pricing Model (CAPM).⁴

In this article, we focus on two types of graphs, efficient frontier graphs and efficient frontier asset allocation area graphs. Efficient frontiers display returns on the vertical axis and the standard deviation of returns on the horizontal axis. Each point on an efficient frontier represents the risk and return of an *efficient* asset allocation, where an efficient asset allocation is one that *maximizes* return per unit of risk.

Efficient frontier asset allocation area graphs complement the efficient frontier graphs. They display the asset allocations of the efficient frontier across the entire risk spectrum. Like efficient frontier graphs, efficient frontier *area* graphs display risk on the horizontal axis. In the examples presented here, the applicable risk spectrum from the efficient frontier is divided into 101 evenly spaced segments. Position 0 corresponds to the minimum variance asset allocation and Position 100 corresponds to the maximum return asset allocation. The percentage of the allocation to each of the asset classes is shown on the vertical axis. Conceptually, the efficient frontier area graph is similar to a standard asset allocation pie chart that shows the asset allocation that corresponds to a particular spot on the efficient frontier, except the efficient frontier area graph displays all of the asset allocations on the efficient frontier. Each of the vertical cross sections represents a particular asset allocation for that particular risk level. It is helpful to look at the efficient frontier graphs and the efficient frontier asset allocation area graphs together because you can simultaneously see the asset allocations associated with a particular risk-return point on the efficient frontier, and vice versa.

In our first example, we use the most recent 10 years of monthly historical data to construct annual capital market assumptions – historical returns, historical standard deviations, and historical correlations. The historical efficient frontier is displayed in Figure 1 and the asset allocations are displayed in Figure 2.

Figure 1: Traditional MVO Efficient Frontier, Historical Inputs (Jan 1995 to Dec 2004)


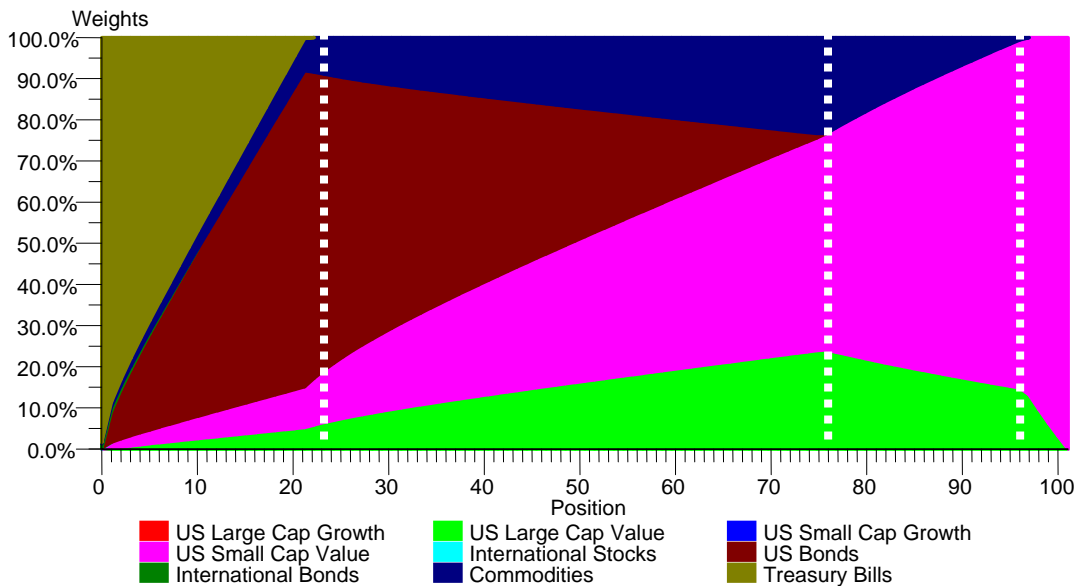


Figure 2: Efficient Frontier Asset Allocation Area Graph (Based on Figure 1 Efficient Frontier)



In Figure 2, Position 0 corresponds to the minimum variance asset allocation and Position 100 corresponds to the maximum return asset allocation. As an alternative to the 0 to 100 position labels, we could have used the appropriate standard deviations. The vertical white dashed lines on the asset allocation area graph represent corner portfolios, the locations where an additional asset class has either entered or exited the asset allocation.⁵ This particular asset allocation area graph is characterized by four sections. In the

first section, five of the nine asset classes in the opportunity set are in the asset allocation. In each of the subsequent sections, an asset class drops out of the asset allocation until 100% is allocated to a single asset class. In each of the sections, the asset allocation is dominated by large allocations to one or two asset classes. Treasury Bills and US Bonds dominate the first section, US Bonds and US Small Cap Value dominate the second section, and US Small Cap Value dominates both the third and fourth sections.

Steep vertical changes in the asset allocation area graph indicate that small changes in standard deviation (the horizontal axis) correspond to large changes in the asset allocations. This is the case in the first section between Position 0 and Position 23, where small changes in the standard deviations correspond to large changes in the allocations between Treasury Bills and US Bonds. Nearly half of the asset classes in the opportunity set, US Large Growth, US Small Growth, International Equity, and International Bonds are completely excluded from the asset allocations.

The nature of *traditional* MVO is such that lower risk asset allocations typically include more of the available asset classes, while the higher risk asset allocations are progressively more concentrated in fewer asset classes. Ultimately, 100% is allocated to the highest return assets class, which is often (but not necessarily) the highest risk asset class.

Next, we repeat the experiment using a different historical 10-year period, January 1985 to December 1994.

Figure 3: Traditional MVO Efficient Frontier, Historical Inputs (Jan. 1985 to Dec. 1994)





Figure 4: Efficient Frontier Asset Allocation Area Graph (Based on Figure 3 Efficient Frontier)



Once again, the vertical white dashed lines on the area graph indicate that an additional asset class has either entered or exited the asset allocation. This particular asset allocation area graph is characterized by seven sections. The highest number of asset classes included in an asset allocation is six. International Stocks and International Bonds, two asset classes that are excluded from Figure 2, play prominent roles in the riskier asset allocations of Figure 4. The US Small Cap Growth asset class is completely excluded from Figure 4.

These two *historical* efficient frontiers tell us the asset allocations that were optimal in the *past*. More specifically, Figures 1 and 2 identify the optimal allocations over the January 1995 to December 2004 time frame and Figures 3 and 4 identify the optimal allocations over the January 1985 to December 1994 time frame. With *perfect* hindsight, or *perfect* foresight for that matter, these highly concentrated asset allocations are desirable; we want to load up on the winners and ignore the losers!

Back on January 1, 1995, when the second of our time periods began, few (if any) investors would have been willing to allocate their assets based on the highly concentrated allocations depicted in Figure 4. Doing so certainly was less than optimal over the ensuing 10-year time frame from January 1995 to December 2004 (see Figure 2). In the absence of a crystal ball, highly concentrated asset allocations are inappropriate and should not be recommended.

There are a handful of import takeaways from comparing the results of the first two optimizations.

- In both cases, the number of asset classes that make up the efficient asset allocations is far fewer than the number of asset classes in the opportunity set.
- Important asset classes are excluded from the asset allocations.
- Different inputs lead to significantly different asset allocations.
- The allocations that were optimal in one period would not have been optimal in the other period.
- In a forward-looking context, few, if any, of these asset allocations could be recommended.

Why are the asset allocations so concentrated?

It is well known that mean-variance optimization is very sensitive to the estimates of returns, standard deviations, and correlations (see Michaud [1989] and Best and Grauer [1991]). Of these three inputs, returns are by far the most important and, unfortunately, the least stable. Chopra and Ziemba [1993] estimated that at a moderate risk tolerance level, mean-variance optimization is 11 times more sensitive to estimation error in returns relative to estimation error in risk (variance) and mean-variance optimization is two times more sensitive to estimation error in risk (variance) relative to estimation error in covariances (which also applies to correlations). Richard Michaud coined the phrase “the Markowitz optimization enigma” to describe the problem of input sensitivity and the highly concentrated asset allocations that result (see Michaud [1989]).

Input sensitivity indicates that the model’s output (the asset allocations) changes significantly due to small changes in the input (the capital market assumptions). *Estimation error* refers to the fact that in a forward-looking context the inputs are forecasts, and as such, are likely to be less than perfect (i.e. they contain errors). Putting these two issues together enables an uninformed practitioner to do more harm than good.

Input sensitivity causes mean-variance optimization to lead to highly concentrated asset allocations. If the inputs contain errors (and in a forward-looking context, they inevitably will), the asset allocations may be concentrated in the wrong asset classes. Mean-variance optimizers are sometimes described as “estimation error maximizers,” referring to the idea that the asset allocations are highly concentrated in the asset classes whose forecasts contain positive errors and the asset allocations exclude the asset classes whose forecasts contain negative errors. The traditional Markowitz algorithm treats the inputs as if they were known with 100% certainty. If one aims to determine the asset allocations that were efficient in the *past*, estimation error is not a factor and highly concentrated asset allocations are *desirable*. In a forward-looking context, in which input forecasts contain estimation error, highly concentrated asset allocations are *undesirable*.

Traditional Solutions – Constraints and Long-term Data

The traditional solutions to input sensitivity and estimation error are to set relatively strict asset class constraints, to use long-term data, or to abandon mean-variance optimization all together. When taken to an extreme, strict constraints represent an abandonment of mean-variance optimization as the asset allocations are essentially assigned by a list of binding constraints. Over short-time periods, the relationship between return and risk varies widely. Over longer time periods, there is typically a more consistent relationship between return and risk. Unfortunately, even very large amounts of historical data may not result in a consistent relationship between risk and return. Additionally, long-term data going back 50 or more years is not available for most of the asset classes that make up a modern opportunity set of asset classes. In the case of our opportunity set, we are limited by our International bond asset class proxy, which begins in January 1985.

State of the Art Solutions – The Black-Litterman Model and Resampling

Two robust asset allocation tools are revolutionizing asset allocation: the Black-Litterman Model and resampled mean-variance optimization (a.k.a. resampling). We introduce both techniques prior to working through some empirical examples.

The Black-Litterman Model


The Black-Litterman model was invented by Fischer Black and Robert Litterman (see Black and Litterman [1989, 1990, and 1992]). The Black-Litterman model enables investors to combine their unique views regarding the performance of various assets with CAPM market equilibrium returns in a manner that results in intuitive, diversified portfolios. More specifically, the Black-Litterman Model uses a Bayesian approach to combine the subjective views of an investor regarding the expected returns of one or more assets with the CAPM market equilibrium expected returns (the prior distribution) to form a new, mixed estimate of expected returns (the posterior distribution).

The CAPM market equilibrium returns are calculated using the reverse optimization technique described in Sharpe [1974], which is an extension of Sharpe's more famous and Nobel Prize-winning idea on the CAPM. Sharpe [1964] elegantly argues that there is a consistent relationship between expected return and market risk. Other names for market risk include systematic risk, benchmark risk, correlated risk, non-diversifiable risk, and beta risk, which is the term we use. Investors demand compensation, or more precisely, the *expectation* of compensation for beta risk. The relationship between expected return and beta risk is the basis for reverse optimization – a procedure that backs out implied expected returns based on a presumed efficient asset allocation. The presumed efficient asset allocation is typically based on market capitalization estimates of the asset classes that form the opportunity set.⁶ Other names for CAPM market equilibrium returns include market returns, implied returns, consensus returns, and reverse optimized returns.

CAPM market equilibrium returns result in return forecasts that are a direct function of beta risk, the type of risk that is *unconditionally* rewarded by the market with a risk premium. Estimation error in the standard deviations and correlations is spread throughout the CAPM market equilibrium returns. The ratio of expected excess return (above the risk-free rate) to beta risk is the same for all of the assets (see Figure 5). This ratio changes when the estimate of the market risk premium changes. The perfect balance between expected return and beta risk results in a set of capital market assumptions that behave well in a traditional mean-variance optimizer. Notice that the excess CAPM market equilibrium returns of the asset classes are proportionate to the risks (as measured by beta) of the asset classes and thus plot on the security market line. The slope of the security market line equals the market risk premium.

It is worth noting that in the context of reverse optimization, changes in the market risk premium change the shape of the efficient frontier, but not the composition of the efficient asset allocations as the ratio of excess return to beta risk is the same for all asset class regardless of the risk premium. The market risk premium determines the slope of the security market line. If a constant representing the risk-free rate is added to each of the excess returns, the excess returns are transformed into total return estimates; yet, the slope of the security market line is unaffected. Changes in the risk-free rate result in parallel shifts up or down in the efficient frontier, but the shape of the efficient frontier and the composition of the efficient asset allocations are unchanged. Even though the risk-free rate and market risk premium (assuming it is positive) do not affect the composition of the efficient asset allocations, the magnitude of the total return estimates will have large implications in any subsequent forecasting / Monte Carlo simulations.


Figure 5: Security Market Line



In Figure 6, the red (lower) efficient frontier is based on a long-term risk-free rate of 4.84% and a market risk premium of 3%. The green (middle) efficient frontier is based on an *unreasonable* long-term risk rate of 25% and a market risk premium of 3%. Changes in the risk-free rate result in parallel shifts in the efficient frontier. The blue (top) efficient frontier is based on a long-term risk rate of 4.84% and an *unreasonable* market risk premium of 25%. Changes in the market risk premium change the shape of the efficient frontier. The key observation is that all three of the efficient frontiers in Figure 6 lead to the same asset allocations!

Starting with CAPM market equilibrium returns enables the Black-Litterman model to produce returns that behave well inside of an optimizer and lead to well-diversified asset allocations. Modern portfolio theory and the CAPM tell us that the asset allocation with the highest expected Sharpe Ratio is the hypothetical market portfolio. The market capitalization weights used in the reverse optimization move us from a hypothetical market portfolio to a working version of the market portfolio that is applicable to the opportunity set of asset classes being optimized. The weights used in the reverse optimization process correspond to a particular definition of the presumed efficient market portfolio and the returns that are implied by those weights. When used in a traditional mean-variance optimizer, the CAPM market equilibrium returns result in an efficient frontier in which the asset allocation with the maximum Sharpe Ratio corresponds to the original market capitalization. This completes an important circle: traditional optimization moves us from returns to asset allocations (weights) and reverse optimization moves us from asset allocations (weights) to returns.⁷

Figure 6: CAPM-Based Efficient Frontiers



This maximum Sharpe Ratio asset allocation also corresponds to the asset allocation in which the ratio of expected excess return to marginal contribution to risk is equal for all of the asset classes.⁸ Moving to either side of the point of tangency (i.e. the maximum Sharpe Ratio asset allocation) on the efficient frontier has two effects: 1) it changes this perfect balance of expected excess return to marginal contribution to total risk relationship, and 2) it decreases the Sharpe Ratio of the asset allocation.

While these starting CAPM market equilibrium returns are sometimes referred to as Black-Litterman returns, they are technically just the starting point for the Black-Litterman model.⁹ For those who do not agree with the CAPM market equilibrium returns and have an opinion or view regarding the return of one or more assets, the CAPM market equilibrium returns can be augmented with this additional information to form a mixed or combined estimate of expected returns. This is done using the Black-Litterman model, of which the mathematical details are beyond the scope of this article (see Idzorek [2005] for a more detailed discussion).

Conceptually, the returns of an individual asset class form a distribution, similar to the traditional bell curve. The center of the distribution represents the expected return of the asset class and the standard deviation of the asset class represents the variation around the expected value. When there is more than one asset, the correlations, a measure which characterizes the degree to which asset returns move together, must also be considered. Collectively, the returns, standard deviations, and correlations form a multi-dimensional distribution of returns, or more accurately, a multivariate distribution of returns. This distribution is the starting point for the Black-Litterman model and serves as an anchor for grounding additional opinions or views regarding expected returns.

Each view also results in a distribution, where the expected return of the view is at the center of the distribution and the uncertainty in the view (around the expected value) is described by the standard deviation of the view. The standard deviation around the expected view value is based on the confidence

one has in the particular view. All else equal, the higher the confidence in the view, the lower the standard deviation in the view and the more aggressively the view will affect the final mixed estimate of expected returns. Views are typically treated independently, thus it is not necessary to take into account the correlation among views as it is assumed to be zero.

In other words, there are two sources of information, each of which is described by its respective distributions: the distribution of CAPM equilibrium returns and the distribution of view returns. The Black-Litterman model combines these two distributions to form a final, mixed estimate of expected returns that is anchored by the CAPM equilibrium returns, but also reflects the view returns. The Black-Litterman model accomplishes this in a manner that still strikes a consistent relationship between expected returns and betas; thus, in most cases, the combined estimate of expected returns still leads to well-diversified asset allocations that reflect the stated views.

The Black-Litterman model is sometimes referred to as a Bayesian asset allocation model because the view returns are shrunk toward the CAPM equilibrium returns.¹⁰ The CAPM equilibrium returns are typically based on market capitalization estimates of the different asset classes that form the opportunity set. As a result, the CAPM equilibrium returns lead to asset allocations that are related to the market capitalization of the asset classes. This helps prevent asset classes with small market capitalizations, such as emerging market equities or small cap stocks, from dominating all but the most aggressive asset allocations.

Given the Black-Litterman model's ability to incorporate other sources of information, in addition to being a powerful strategic asset allocation tool it is a powerful tactical (or dynamic) asset allocation tool. This can be done using market capitalization weights or previously established strategic asset allocation weights / policy benchmark weights in the reverse optimization process.

The Black-Litterman model leads to return estimates that produce to well-diversified asset allocations when used with either *traditional* MVO or *resampled* MVO. *Resampled* MVO is the subject of our next section and a method that results in well-diversified asset allocation, albeit using a very different approach.

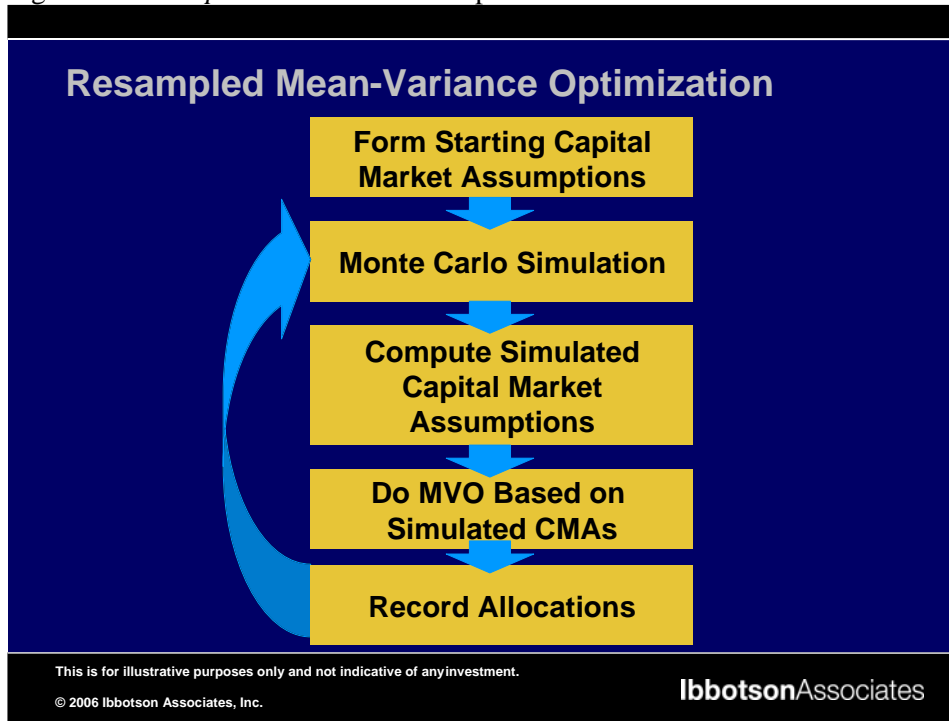
Resampling

Resampled MVO combines *traditional* MVO with Monte Carlo simulation to account for the uncertainty in forward-looking capital market assumptions. The current embodiments of *resampled* MVO grew out of the work of Jobson and Korkie [1980, 1981], Jorion [1992], DiBartolomeo [1993], and Michaud [1998]. The primary theory behind *resampled* MVO is that *traditional* MVO treats the estimated expected returns, standard deviations, and correlations as if they were known with 100% certainty. In a forward-looking context these inputs are not known with certainty. *Resampled* MVO is an optimization technique that accounts for input uncertainty and simultaneously addresses the issues of estimation error, input sensitivity, and diversification.

Conceptually, *resampled* MVO is a large scale sensitivity analysis in which hundreds or perhaps thousands of permutations on the *starting* inputs lead to an equal number of *traditional* MVO frontiers. Based on either the starting MVO capital market assumptions (a parametric approach) or the historical returns of the asset classes that make up the opportunity set (a non-parametric approach), a Monte Carlo simulation produces a *simulated* set of capital market assumptions. The simulated set of capital market assumptions is fed into a *traditional* mean-variance optimizer, resulting in an intermediate frontier called a *simulated* frontier. The resulting asset allocations, or portfolio weights, from these *simulated* frontiers are saved. The process is repeated many times and the asset allocations from each of the simulated

frontiers are averaged (using a variety of methods).¹¹ Ibbotson uses a proprietary “bin approach,” in which asset allocations from the *simulated* frontiers are grouped together based on narrowly defined standard deviation ranges that cover the risk spectrum. The average asset allocations are then linked back to the original inputs to plot the resampled efficient frontier. The resampling process is depicted in Figure 7. *Resampled MVO* is computationally intensive and, depending upon the number of asset classes, can take several minutes to complete.

Figure 7: *Resampled Mean-Variance Optimization*



Resampling employs one of two types of Monte Carlo simulation – historical non-parametric simulation (sometimes referred to as bootstrapping) or parametric simulation. Historical resampling uses historical time series data and does not require distributional assumptions, hence the term “non-parametric.” Parametric resampling requires an assumption about the asset class return distribution or distributions. For a particular opportunity set, the multivariate return distribution is defined by the same three standard mean-variance optimization inputs – returns, standard deviations, and correlations. Parametric Monte Carlo simulation that incorporates correlations, in addition to return and standard deviations, is a multivariate Monte Carlo simulation.¹² The inputs for parametric resampling do not need to come from historical data, although in practice they are often based on historical data. We will focus on parametric resampling.

Parametric resampling will work with any method of estimating returns, standard deviations, and correlations, as long as the correlation matrix is well defined. Nevertheless, it is important to use the best possible forecasts of the capital market assumptions. The standard axiom, “garbage in garbage out,” still applies. Shortly, we will compare *resampled MVO* asset allocations based on historical inputs to *resampled MVO* asset allocations based on the starting Black-Litterman model return, i.e. the CAPM equilibrium returns.

Robust Asset Allocations

Figures 1-4 above demonstrate that historical inputs coupled with *traditional* MVO result in highly concentrated asset allocations. When asking what was optimal in the *past*, these asset allocations provide the answer. If the goal is to set asset allocation policy for the future, highly concentrated asset allocations are untenable.

At this point, we have produced two *traditional* MVO frontiers and corresponding asset allocation area graphs based on two different sets of historical inputs. We will now create three more efficient frontiers and corresponding efficient frontier asset allocation area graphs using various permutations of return estimates and optimization models. All of the ensuing efficient frontiers are based on the same January 1995 to December 2004 historical standard deviations and correlations used to create Figures 1 and 2. The various permutations, including the first two permutations that have already been presented, are described in Table 1.

Table 1 Return Estimate and Optimization Model Permutations

Figure	Return Estimate	Optimization Model
Figures 1 and 2	Historical (1995 to 2004)	Traditional MVO
Figures 3 and 4	Historical (1985 to 1994)	Traditional MVO
Figures 8 and 9	Black-Litterman (CAPM Equilibrium)	Traditional MVO
Figures 10 and 11	Historical (1995 to 2004)	Resampled MVO
Figures 12 and 13	Black-Litterman (CAPM Equilibrium)	Resampled MVO

Traditional MVO with Black-Litterman (CAPM Equilibrium) Returns

Using *traditional* MVO, Figures 8 and 9 are based on the CAPM equilibrium returns that are the starting point for the Black-Litterman model.

Figure 8: Traditional MVO Efficient Frontier, Black-Litterman Inputs

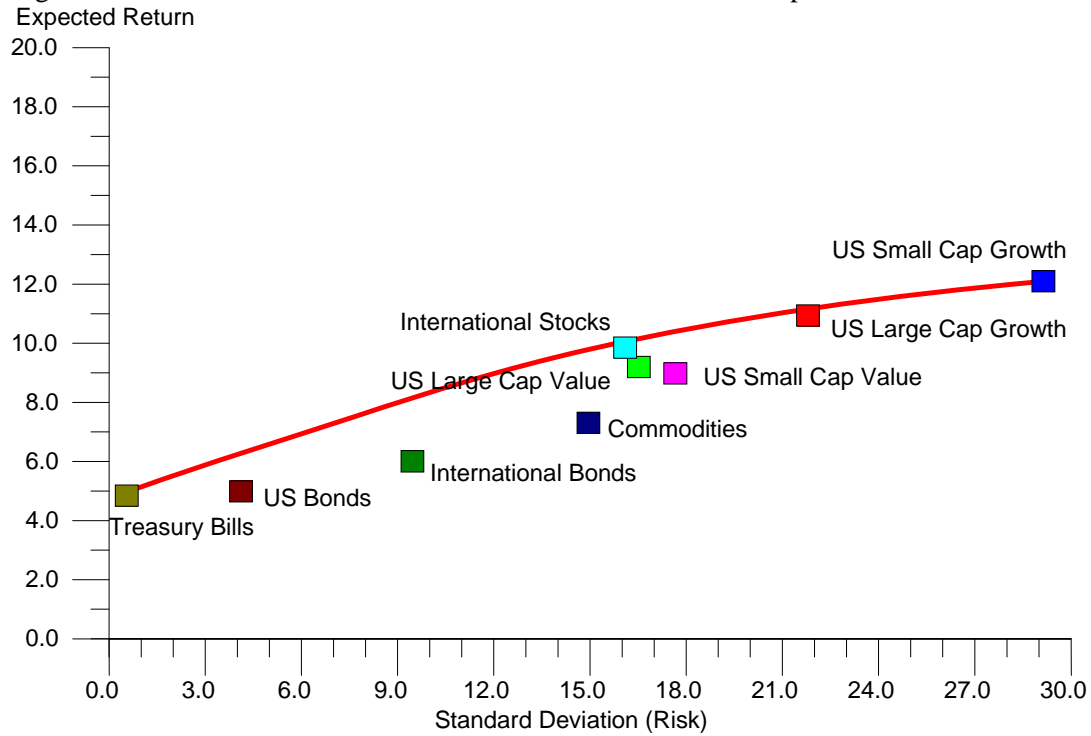
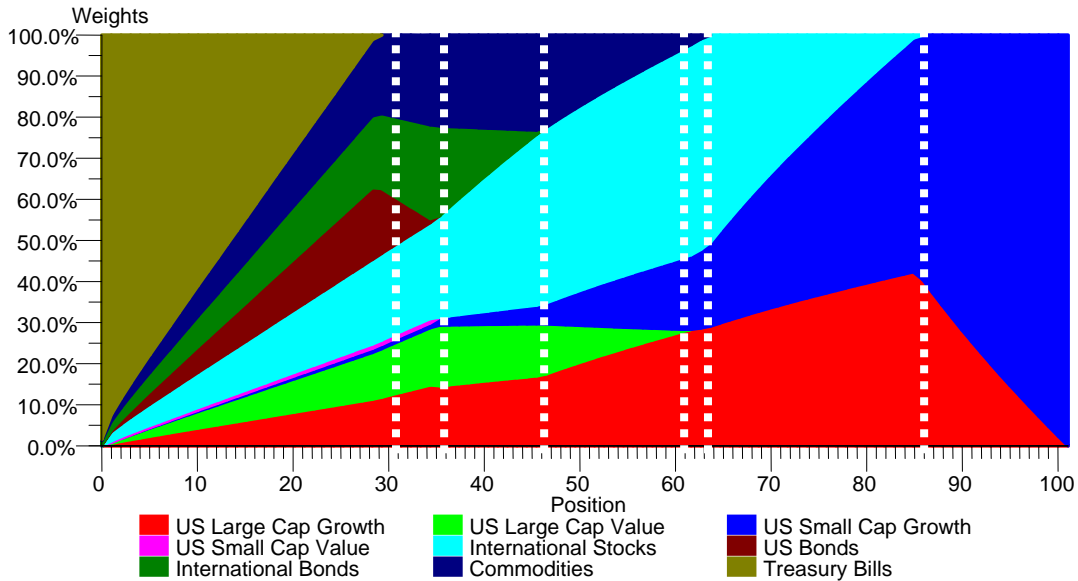


Figure 9: Efficient Frontier Asset Allocation Area Graph (Based on Figure 7 Efficient Frontier)



Relative to the first two historical optimizations, there is a substantial increase in the number of asset classes in the efficient asset allocations. All nine asset classes are included in the first section, eight in the second section, and six in the third section. Once again, the vertical white dashed lines on the area graph indicate the approximate point in which an additional asset class has either entered or exited the asset allocation. In the first section, allocations to Treasury Bills are gradually decreased as the allocations to all of the other asset classes grow proportionately. For the part of the efficient frontier from which asset allocations are typically selected (Position 0 – 50), the asset allocations are well-diversified and intuitive. The ratio of growth asset classes (US Large Cap Growth and US Small Cap Growth) to value asset classes (US Large Cap Value and US Small Cap Value) is almost evenly divided. The ratio of large capitalization asset classes to small capitalization asset classes is grounded in economic reality. The split between the four US equity asset classes and International Stocks is reasonable, as is the split between US Bonds and International Bonds.

Resampled MVO with Historical Returns (January 1995 to December 2004)

Figures 10 and 11 are based on the same historical inputs that were used to create Figures 1 and 2; the only difference is the optimization method. Figures 10 and 11 were created using *resampled* MVO while Figures 1 and 2 were created using *traditional* MVO.

Traditional MVO-based efficient frontiers are always concave (outward bulging). Due to the nature of the averaging procedure used to calculate *resampled* MVO-based efficient frontiers, it is possible to have convex (inward bulging) sections. This is typically most pronounced at the upper end of the risk spectrum, an area of the frontier from which asset allocations are seldom selected. From this perspective, the *resampled* MVO frontier based on historical inputs (Figure 10) is slightly less aesthetically appealing than the corresponding *traditional* MVO frontier based on historical inputs (Figure 1).

Figure 10: Resampled MVO Efficient Frontier, Historical Inputs

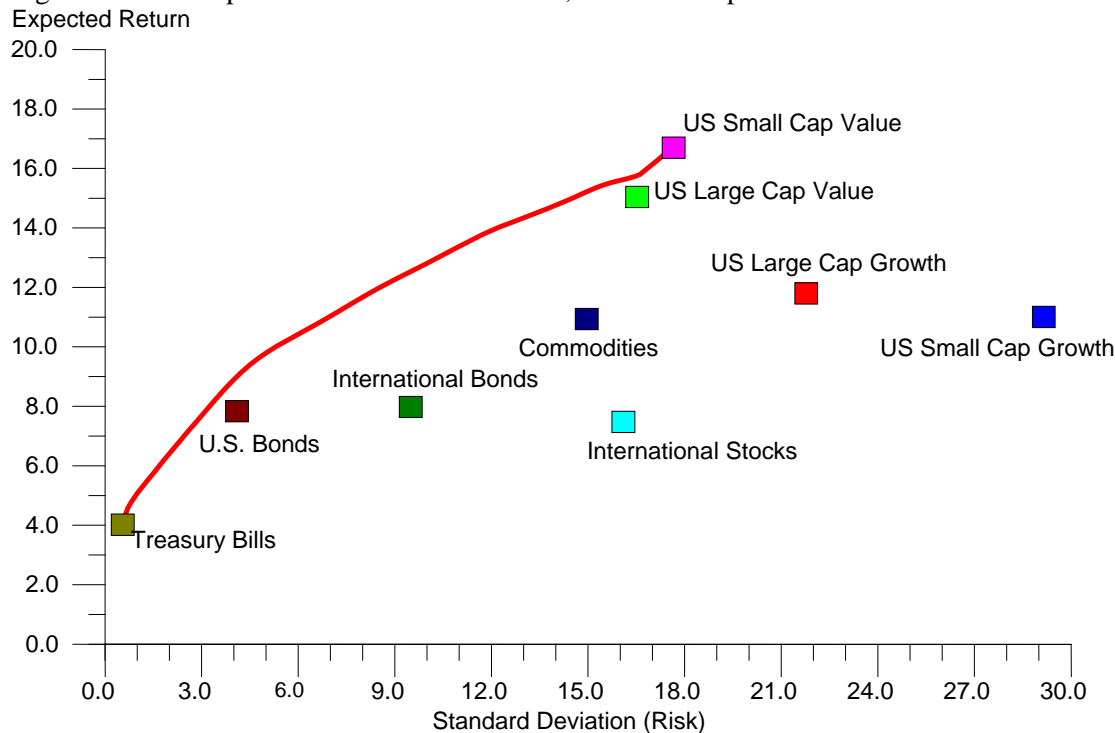



Figure 11: Efficient Frontier Asset Allocation Area Graph (Based on Figure 7 Efficient Frontier)



The *resampled* MVO-based asset allocations of Figure 11 are more appealing than the corresponding *traditional* MVO-based asset allocations of Figure 2. More specifically, the asset allocations evolve smoothly across the efficient frontier asset allocation area graph and are significantly more diversified than the *traditional* MVO-based asset allocation of Figure 2. The evolution of asset allocations is so smooth with small allocations to some asset classes persisting across large sections of the efficient frontier that denoting various segments is less meaningful; thus, it is difficult to use vertical white dashed lines to indicate where an additional asset class has either entered or exited the asset allocation.

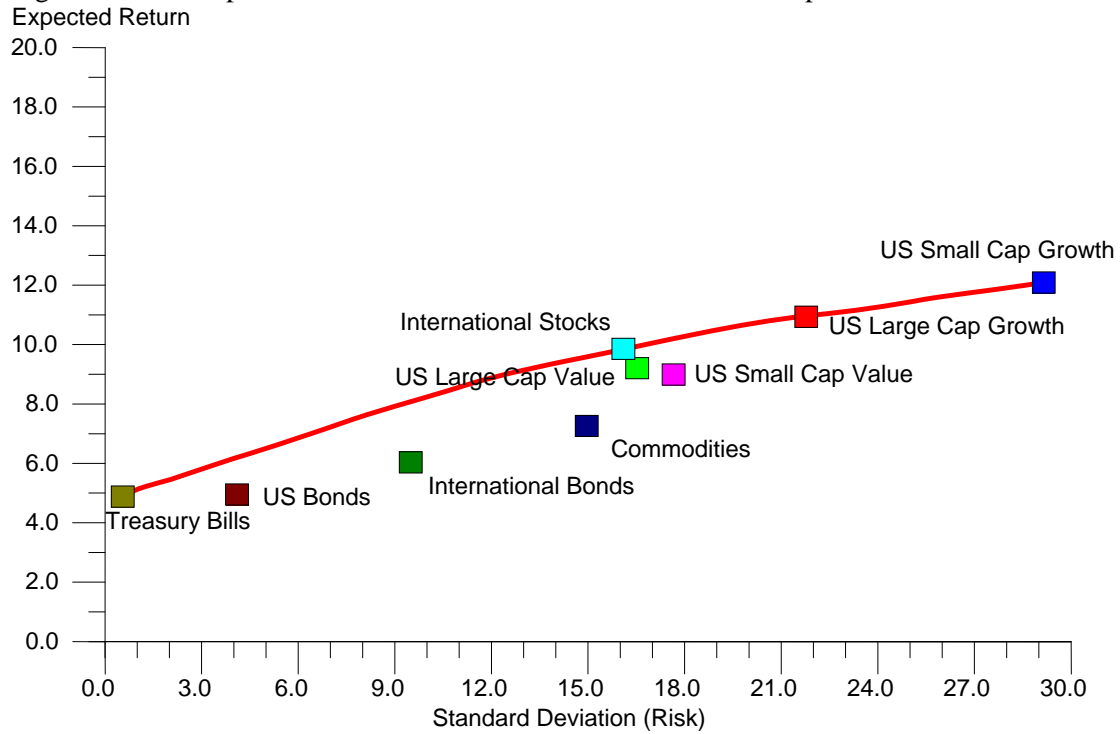
Relative to Figure 2, the core asset allocations are similar; however, the *resampled* MVO-based asset allocations are considerably more diversified with at least eight of the nine asset classes in the majority of the asset allocations. The most noticeable differences are the allocations to International Bonds and US Large Cap Growth stocks, two asset classes that are absent from the *traditional* MVO-based asset allocations. There are several possible shortcomings of the asset allocations in Figure 10. The combined allocation to the two US small cap asset classes is nearly the same size as the allocation to the two US large cap asset classes. Additionally, US Large Cap Growth, US Small Cap Growth, and International Stocks are almost ignored.

While resampling improved the diversification of the asset allocations, the quality of the asset allocations still depends on the quality of the inputs. Optimizations based on historical data have no link to the economic significance (size) of the asset classes in a given opportunity set.

Resampled MVO with Black-Litterman (CAPM Equilibrium) Returns


The final efficient frontier and corresponding asset allocation area graph are based on the Black-Litterman model CAPM equilibrium returns using *resampled* MVO.

Figure 12: Resampled MVO Efficient Frontier, Black-Litterman Inputs



Relative to the previous *resampled* MVO efficient frontier based on historical inputs (Figure 7), the inward bulging sections of the *resampled* MVO efficient frontier based on Black-Litterman CAPM equilibrium returns (Figure 11) are less pronounced. Both the *traditional* MVO and *resampled* MVO efficient frontiers based on the Black-Litterman CAPM equilibrium returns (Figures 8 and 12) are nearly identical in appearance. The differences in the corresponding asset allocation area graphs (Figures 9 and 13) are more pronounced.

Figure 13: Efficient Frontier Asset Allocation Area Graph (Based on Figure 12 Efficient Frontier)



The combination of Black-Litterman CAPM equilibrium returns with *resampled* MVO produces the most diversified asset allocations (Figure 13). The higher risk asset allocations (Positions 50 – 100) are significantly more diversified than those of the previous return and optimization model permutations. We suspect that higher risk asset allocations based on *resampled* MVO are significantly more robust than those from *traditional* MVO.

The lower risk asset allocations (Positions 0 – 50) are similar to those of Figure 9, the asset allocation area graph based on the Black-Litterman CAPM equilibrium returns coupled with *traditional* MVO. Relative to Figure 9, at the lower risk levels, the asset allocations to the two US small cap asset classes are much larger. Small allocations to most of the asset classes persist throughout the higher risk asset allocations.

Comparing Efficient Frontiers

Relatively few measures for comparing efficient frontiers and their corresponding asset allocations exist. We introduce two measures that help us make high-level comparisons.

Using a method analogous to Idzorek and Bertsch [2004] that is applied to rolling asset allocation style graphs, we can measure the variability of the asset allocations across an efficient frontier. The variability of an equally weighted asset allocation is zero. This is not a statement regarding whether or not an equally weighted asset allocation is desirable, it is simply one possible measure of diversification. High average variability indicates extreme asset allocations, asset allocations with high allocations to some assets and low allocation to others. On a typical efficient frontier, 100% is allocated to the asset class with the highest expected return. For a typical efficient frontier, this is the asset allocation with the greatest variability in asset class weights: the weight is either 100% or 0%. All of the other asset classes are less extreme, with the noticeable exception that the minimum variance asset allocation can theoretically contain a 100% allocation.

We call these two new measures the *aggregate* frontier variability and *average* frontier variability. While our intermediate calculations are based on the *variance* of the asset allocations, we express these

measures in the more familiar standard deviation space. Comparisons are more meaningful when the number of asset allocations observed per frontier are the same, the opportunity set is the same, and the constraints are the same, as is the case in all of our examples.

Table 2 contains several measures of diversification, including aggregate frontier variability and average frontier variability, for the four permutations of inputs and optimization models.

Table 2: Efficient Frontier Summary Metrics

	Aggregate Frontier Variability	Average Frontier Variability	Average Percentage of Assets in Asset Allocations	Average Non-Zero Asset Allocation
Historical + <i>Traditional</i> MVO	4.275	20.57%	44.22%	27.13%
Historical + <i>Resampled</i> MVO	3.131	17.61%	97.14%	12.31%
Black-Litterman + <i>Traditional</i> MVO	3.550	18.75%	60.40%	25.03%
Black-Litterman + <i>Resampled</i> MVO	2.987	17.20%	97.91%	12.21%

These measures support what we have already seen and provide us with several metrics for making high-level comparisons. Lower average frontier variability figures indicate that the asset allocations are less volatile across the frontier. Relative to an equally weighted asset allocation, the Historical + *Traditional* MVO asset allocations were the most variable.


Relative to *traditional* MVO, *resampled* MVO decreased the average frontier variability and increased the average percentage of assets in the asset allocations. Relative to historical returns, Black-Litterman CAPM equilibrium returns also decreased the average frontier variability and increased the average percentage of assets in the asset allocations. Using the Black-Litterman CAPM equilibrium returns with *resampled* MVO resulted in the lowest average frontier variability and the highest average percentage of asset in the asset allocations.

Further insights can be gained by examining several of the asset allocations from the area graph in more detail using conventional asset allocation pie charts. We focus on the cross sections with standard deviations of 5%, 7.5%, and 10%, which we might think of as conservative, moderate, and aggressive asset allocations. Notice that the “aggressive” asset allocation comes from around the middle of the efficient frontier. In practice, few investors select asset allocations from the most aggressive third of the efficient frontier, as the marginal increase in return per unit of risk is quite small and the allocations become significantly more concentrated in the high-return assets.

In Figure 14, the conservative 5% standard deviation cross section, the Historical + *Traditional* asset allocation is by far the most concentrated, consisting of four out the nine asset classes in the opportunity set. The lack of Treasury Bills from the Historical + *Traditional* asset allocations is troubling for a conservative portfolio. The other three methods of determining asset allocations resulted in diversified asset allocations in which all nine of the asset classes from the opportunity set are included in the asset allocations. Relative to the two Black-Litterman-based asset allocations, the Historical + *Resampling* asset allocation overweights US Bonds relative to Treasury Bills. When studied closely, there are noticeable differences between Black-Litterman + *Traditional* and Black-Litterman + *Resampling*: the

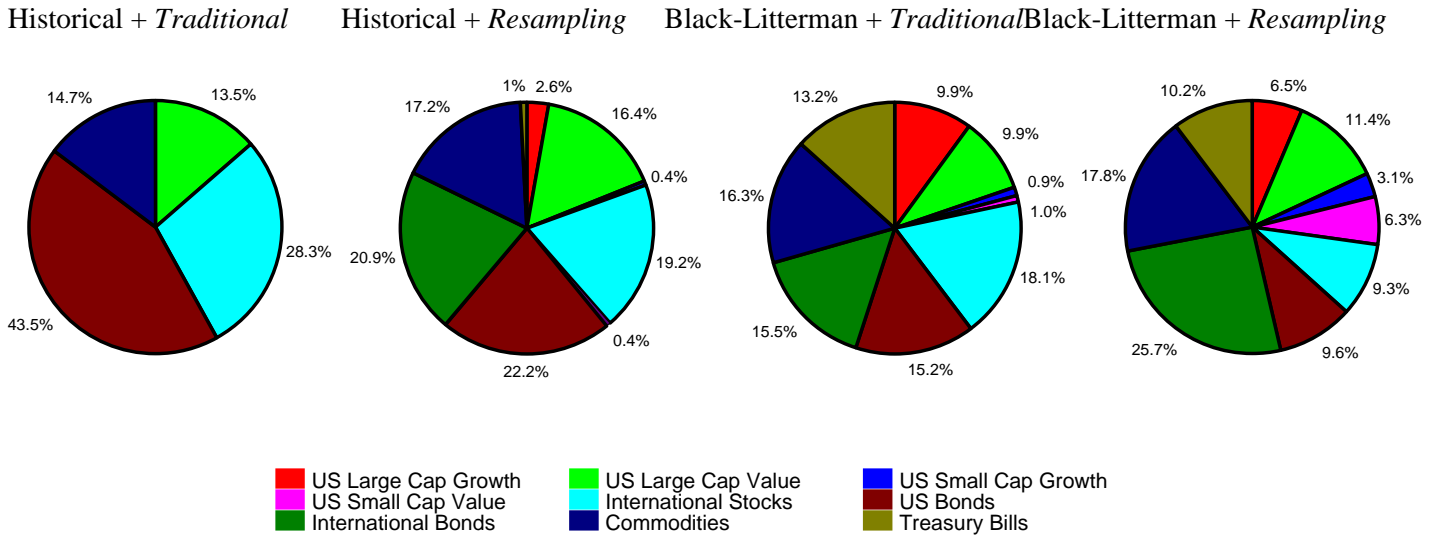
allocation to Treasury Bills decreased by nearly 18% and the allocation to US Bonds and International Bonds both increased by 10%.

Figure 14: Conservative Asset Allocation (Standard Deviation = 5%)




In Figure 15, the moderate 7.5% standard deviation cross section, the Historical + *Traditional* asset allocation remains the most concentrated, consisting of four of the nine asset classes in the opportunity set. The other three methods of determining asset allocations resulted in diversified asset allocations in which all nine of the asset classes from the opportunity set are included in the asset allocations. The reasonableness of the Historical + *Resampling* asset allocation is questionable. It includes a substantial asset allocation to US Small Value, a large value vs. growth bias, and virtually no international stocks. The differences between the two Black-Litterman-based asset allocations are subtle. The Black-Litterman + *Resampling* example has a value vs. growth bias, larger positions in the two US small cap asset classes, and it overweights International Bonds relative to US Bonds.

Figure 15: Moderate Asset Allocation (Standard Deviation = 7.5%)



In Figure 16, the aggressive 10% standard deviation cross section, the Historical + *Traditional* asset allocation remains the most concentrated, consisting of four out the nine asset classes in the opportunity set. The *resampled* MVO-based asset allocations use all nine of the asset classes in the opportunity set. Treasury Bills are no longer included in the Black-Litterman + *Traditional* asset allocation. The differences between the two Black-Litterman-based asset allocations are more pronounced than they were in previous two examples. The Black-Litterman + *Resampling* continues to have a value vs. growth bias, larger positions in the two US small cap asset classes, and over weights International Bonds relative to US Bonds.

Figure 16: Aggressive Asset Allocation (Standard Deviation = 10%)



The Black-Litterman model and resampled MVO are two very different techniques. Relative to traditional MVO with pure historical inputs, both approaches help practitioners overcome the most common criticism of *traditional* MVO: highly concentrated asset allocations. For the typical practitioner,

using the Black-Litterman model with *resampled* MVO may be the most powerful combination for consistently developing a forward-looking efficient frontier and corresponding asset allocations that are not only usable, but robust.

Relative to *traditional* MVO, *resampled* MVO always decreased the size of the largest asset allocation and increased the size of the smallest asset allocation. The value-over-growth bias that seemed to result from using *resampled* MVO deserves further study.

Our conversation on the Black-Litterman model would not be complete without a brief introduction explaining how the Black-Litterman model can be used to combine CAPM equilibrium returns with an investor's view to create a combined return estimate. We will illustrate this ability using a hypothetical view.

Our fictional view is that a market capitalization weighted portfolio of the two US *value* asset classes (US Large Cap Value and US Small Cap Value) will outperform a market capitalization weighted portfolio of the two US *growth* asset classes (US Large Cap Growth and US Small Cap growth) by 2% on an annual basis. Depending upon the context, this may be a long-term strategic view or a short-term tactical view. Whether or not this is a bullish or bearish view on US value stocks relative to US growth stocks depends on the CAPM equilibrium returns. Based on the CAPM equilibrium returns used in our example, the weighted average returns on the two value asset classes and the two growth asset classes are 9.16% and 10.99%, respectively, a difference of 183 basis points in favor of growth over value. Thus, the view that value will outperform growth stocks by 2% is an extreme view stating that the return differential is closer to 383 basis points.

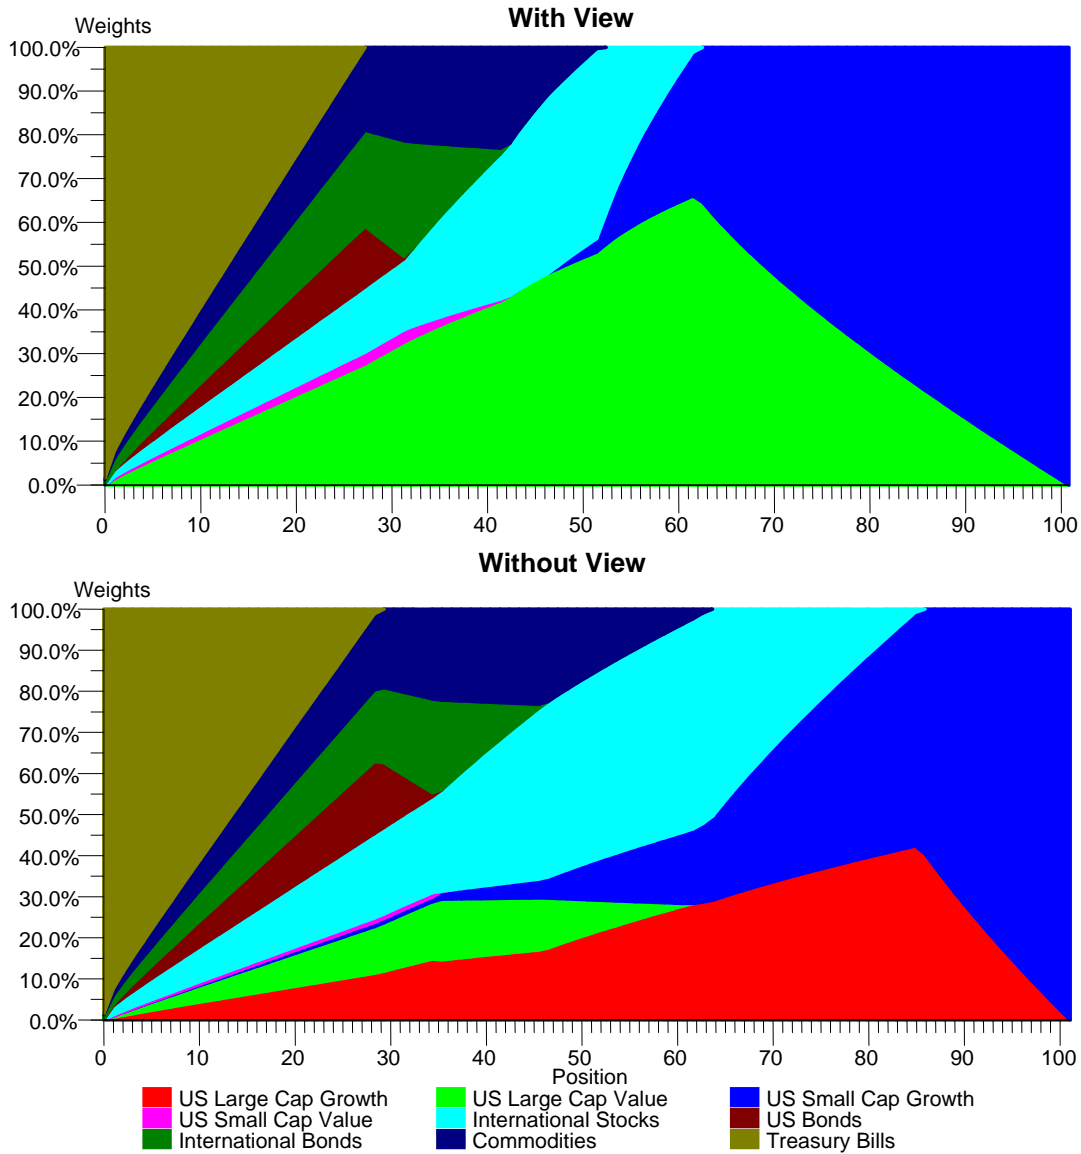
Traditionally, investors have used ad hoc approaches to incorporate such a view into their return forecasts.¹³ For example, an investor might *increase* the return of the two value asset classes and *decrease* down the return of the two growth asset classes. Unfortunately, given the sensitivity of traditional mean-variance optimization to changes in returns, this is likely to have profound and unintuitive implications on the resulting asset allocations. The Black-Litterman model provides us with an elegant framework for blending this type of view with the CAPM equilibrium returns in such a way that the resulting asset allocations remain intuitive.

In the absence of views, the Black-Litterman model leads back to the market portfolio. The asset allocation on the efficient frontier to the *left* of the market portfolio can be interpreted as more *conservative* versions of the market portfolio. Likewise, the asset allocation on the efficient frontier to the *right* of the market portfolio can be interpreted as more *aggressive* versions of the market portfolio. In the presence of views, the asset allocations can be interpreted as variations on the market portfolio that reflect the views.

Figure 17 compares the allocations *with* and *without* the view based on *traditional* MVO. The top half of Figure 17 presents the asset allocations of a *traditional* MVO based on Black-Litterman returns that incorporate the hypothetical view that value stocks will outperform growth stocks by 2%. The bottom half of Figure 17 presents the asset allocations of a *traditional* MVO based on the Black-Litterman CAPM equilibrium returns without the view.

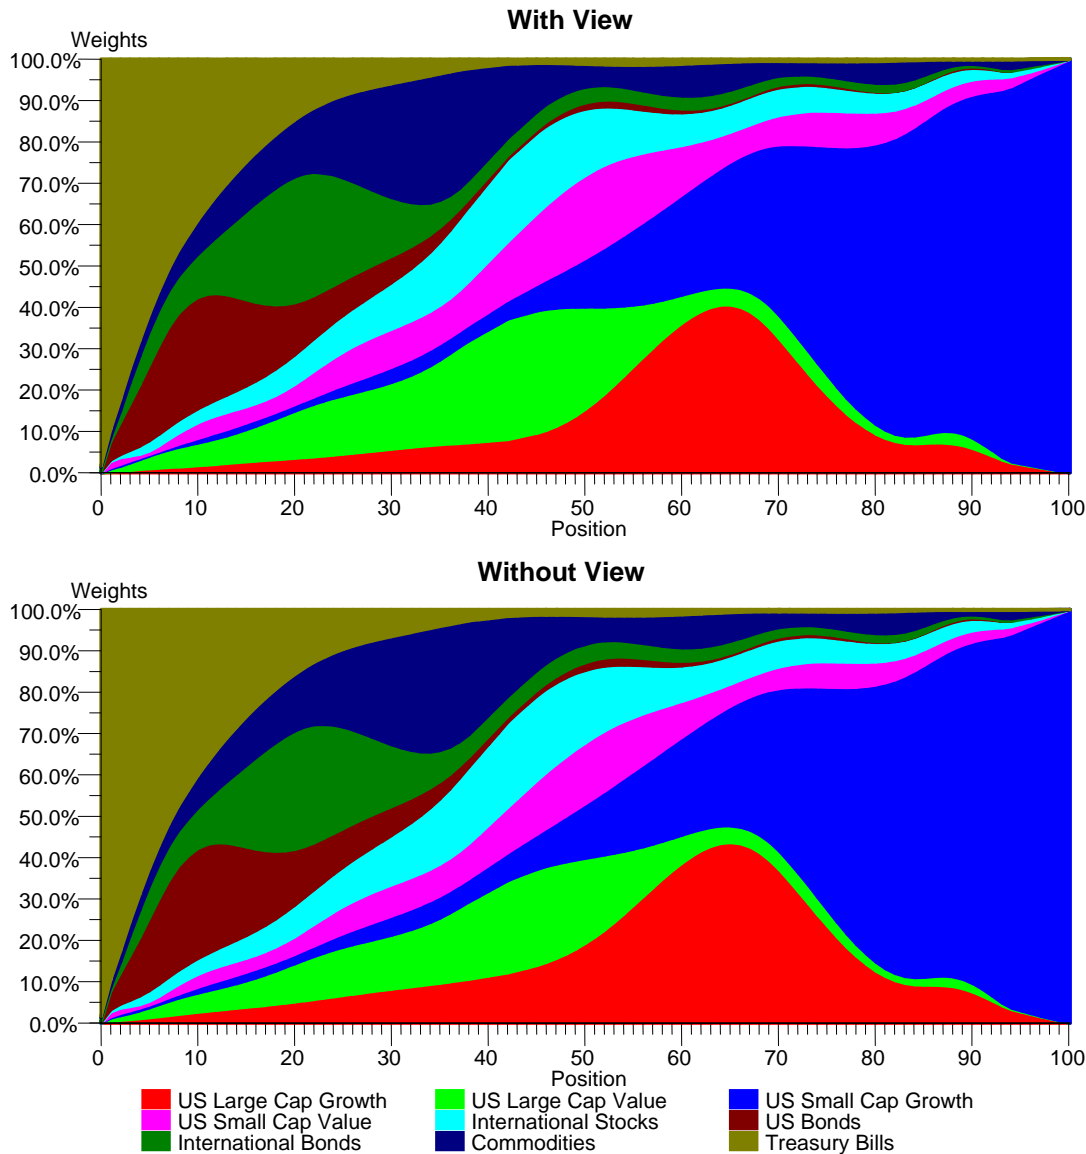
As one would expect with this extremely bullish view on value stocks, the asset allocations of the top panel significantly overweight value stocks relative to the lower panel. Intuitively, this overweight position in value stocks was primarily funded by decreasing the asset allocations to growth stocks. The most important observation is that incorporating the view into the combined return estimate did not wreak the usual havoc associated with the traditional ad hoc approaches to incorporating views.

Figure 17: Asset Allocations with and without view based on *traditional* MVO



Next, we repeat the process using *resampled* MVO. The top half of Figure 18 presents the asset allocations of a *resampled* MVO based on Black-Litterman returns that incorporate the hypothetical view that value stocks will outperform growth stocks by 2%. The bottom half of Figure 18 presents the asset allocations of the *resampled* MVO based on the Black-Litterman CAPM equilibrium returns without a view.

Figure 18: Asset Allocations with and without view based on *resampled* MVO



Notice that two value asset classes, US Large Cap Value and US Small Cap Value are slightly larger in the top panel than they are in the bottom panel. Similarly, the two growth asset classes, US Large Cap Growth and US Small Cap Growth are slightly smaller in the top panel than they are in the bottom panel. A particularly nice feature of the Black-Litterman model is that asset allocations of the five asset classes that were not part of the view remain essentially unchanged.

Conclusion

Mean-variance optimization is widely regarded as the premier asset allocation tool. *Traditional* MVO treats the inputs (the capital market assumptions) as if they were known with certainty. Coupling this with the related problems of input sensitivity, input estimation error, and the common practice of using historical return to estimate the inputs, will typically lead to unusable asset allocations in which the

efficient asset allocations are highly concentrated in a small sub-set of the assets in the opportunity set. We examined two robust asset allocation procedures, the Black-Litterman asset allocation model and *resampled* MVO, which help practitioners overcome the weaknesses of *traditional* MVO and result in robust asset allocations. The two approaches are very different. The Black-Litterman model is a model for developing a set of expected returns that behaves well within an optimizer. The Black-Litterman model starts with CAPM expected returns and provides an elegant framework for augmenting the CAPM expected returns with proprietary research. It can be used for either strategic asset allocation or tactical asset allocation. *Resampled* MVO is an attempt to create a better optimizer. It combines Monte Carlo Simulation with *traditional* MVO. Because they are so different, they can be used together. Individually, both approaches led to more diversified asset allocations relative to *traditional* MVO asset allocation based on historical data. Finally, combining the two approaches is a powerful combination for developing robust forward-looking asset allocations.

References

- Athayde, Gustavo M. de, and Renato G. Flôres, Jr. (2003). "Incorporating skewness and kurtosis in portfolio optimization: a multidimensional efficient set." In Stephen Satchell and Alan Scowcroft, eds., *Advances in Portfolio Construction and Implementation*. Oxford: Butterworth-Heinemann.
- Athayde, Gustavo M. de, and Renato G. Flôres, Jr. (2004). "Finding a maximum skewness portfolio—a general solution to three moments portfolio choice." *Journal of Economic Dynamics and Control*, vol. 28, issue 7, 1335–1352.
- Best, Michael J., and Robert R. Grauer. (1991). "On the Sensitivity of Mean-Variance-Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results." *The Review of Financial Studies*, January, 315-342.
- Black, Fischer, and Robert Litterman. (1992). "Global Portfolio Optimization." *Financial Analysts Journal*, September/October, 28-43.
- Brinson, Gary P., L. Randolph Hood, and Gilbert L. Beebower. (1986). "Determinants of Portfolio Performance." *Financial Analysts Journal*, July/August, 39-44.
- Ceria, Sebastian, and Robert Stubbs [2005]. "Incorporating Estimation Error into Portfolio Selection: Robust Efficient Frontiers." Axioma, Working Paper.
- Chopra, Vijay. K., and William T. Ziemba. (1993). "The effect of errors in means, variances, and covariances on optimal portfolio choice." *Journal of Portfolio Management*, Winter, 6-11.
- DiBartolomeo, Dan. (1993). "Portfolio Optimization: The Robust Solution." Prudential Securities Quantitative Conference. Available online at http://208.15.47.66/Papers/NorthfieldResearchDocs/19931221_optimization_robust.pdf.
- Herold, Ulf, and Raimond Maurer. (2003). "Bayesian Asset Allocation and U. S. Domestic Bias." *Financial Analysts Journal*, November/December, 54-65.
- Jobson, David J., and Bob Korkie. (1980). "Estimation for Markowitz Efficient Portfolios." *Journal of the American Statistical Association*, Vol. 75, September, 544-554.
- Jobson, David J., and Bob Korkie. (1981). "Putting Markowitz Theory to Work." *Journal of Portfolio Management*, Summer, 70-74.

- Jorion, Phillippe. (1992). "Portfolio Optimization in Practice." *Financial Analysts Journal*, January-February, 68-74.
- Idzorek, Thomas M. (2005). "Monte Carlo Simulation in EnCorr," *Ibbotson Research Paper*, Ibbotson Associates, Fall 2005.
- Ibbotson, Roger G., and Paul D. Kaplan. (2000). "Does Asset Allocation Policy Explain 40, 90, or 100 Percent of Performance?" *Financial Analysts Journal*, January/February, 26-33.
- Idzorek, Thomas M., and Fred Bertsch. (2004). "The Style Drift Score." *Journal of Portfolio Management*, Fall, 76-83.
- Idzorek, Thomas M. (2005). "A Step-By-Step Guide to the Black-Litterman Model." Working Paper, April 26.
- Idzorek, Thomas M. (2006). *Modern Asset Allocation: Implementing Theory*, Working manuscript.
- Lintner, John. (1965). "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets." *Review of Economics and Statistics*, February, 13-37.
- Markowitz, Harry M. (1952). "Portfolio Selection." *The Journal of Finance*, March, 77-91.
- Markowitz, Harry M. (1959). *Portfolio Selection: Efficient Diversification of Investments*. New York: Wiley, 1959.
- Michaud, Richard O. (1998). *Efficient Asset Management*. Boston: Harvard Business School Press.
- Mossin, Jan. (1966). "Equilibrium in a Capital Asset Market." *Econometrica*, October, 768-783.
- Perret-Gentil, Cedric and Maria-Pia Victoria-Freser. (2004). "Robust Mean-Variance Portfolio Selection." *National Centre of Competence in Research: Financial Valuation and Risk Management: Working Paper No. 173*, July, 1-43. Available online at www.nccr-finrisk.unizh.ch/media/pdf/wp/WP173_10.pdf.
- Scherer, Bernd. (2005). "How Different is Robust Optimization Really?" Working Paper.
- Sharpe, William F. (1964). "Capital Asset Prices: A Theory of Market Equilibrium." *Journal of Finance*, September, 425-442.
- Sharpe, William F. (1974). "Imputing Expected Security Returns from Portfolio Composition." *Journal of Financial and Quantitative Analysis*, June, 463-472.
- Treynor, Jack L. (1961). "Market Value, Time, and Risk." Unpublished Manuscript dated August 8, 1961.
- Treynor, Jack L. (1962). "Toward a Theory of Market Value of Risky Assets." Unpublished Manuscript dated Fall 1962.

End Notes

¹ This article borrows from sections of my working manuscript, "Modern Asset Allocation: Implementing Theory."

² All too often investors mistakenly think hedge funds are an asset class. If an investment can theoretically add alpha it is a product and not an asset class. Mutual fund managers and hedge fund managers are both managers; they simply play by different sets of rules.

³ By no means are these the only "robust" techniques; rather these are two sophisticated asset allocation enhancements that are enjoying relatively large scale adoption by practitioners. Fabozzi, Forcardi, and Kolm [2005] refer to these types of approaches, which mitigate the affects of estimation error, as robust *frameworks*. A number of authors have proposed higher moment and downside risk portfolio selection models as well as other robust techniques. See for example, Athayde and Flôres [2003, 2004], Ceria and Stubbs [2005], Perret-Gentil and Victoria-Freser [2004], and Scherer [2005].

⁴ As the number of asset classes in the opportunity set increases, so does the likelihood of multicollinearity (i.e. linear combinations of asset classes that are highly correlated with other asset classes or combinations of asset classes). Multicollinearity results in unstable asset allocations because the same efficient risk-return trade-off can be obtained with different combinations of the asset classes.

⁵ Corner portfolios also occur when the condition of a constraint becomes binding, or alternatively, is no longer binding. Our examples do not include constraints; nevertheless, our dashed lines are only approximations of corner portfolios. If multiple asset classes enter or exit the efficient asset allocation at nearly the same point, we use a single line to indicate the change (i.e. multiple corner portfolios) rather than drawing dashed lines on top of one another. The asset allocations between two corner portfolios are linear combinations of the adjacent corner portfolios.

⁶ At any point in time, the current market capitalizations of the various asset classes represent market clearing figures and, while the market may not be perfectly efficient, it is certainly reasonably efficient.

⁷ When unconstrained, these two (forward and reverse) optimization procedures do not require a numerical optimization routine and can be solved with (matrix) algebra, in which we have an equation with one unknown variable (vector).

⁸ The marginal contribution to total risk of an asset class is obtained by dividing the covariance of the asset class (with the total asset allocation) by total risk (standard deviation). The only difference between this and a beta calculation is the use of the standard deviation rather than the variance in the numerator.

⁹ While reverse optimization is attributed to Sharpe [1974], the CAPM was developed separately in Sharpe [1964], Lintner [1965], Mossin [1966], and Treynor [1961, 1962].

¹⁰ Bayesian approaches attempt to create better estimates of expected return by shrinking the expected returns toward a reference set of returns. For excellent descriptions of the common Bayesian asset allocation techniques, see Herold and Maurer [2003].

¹¹ The different methods for grouping together asset allocations that are then averaged are typically what differentiate the various resampling approaches. Michaud [1998] links the asset allocations of the simulated frontiers using an index system, where the minimum variance asset allocation is assigned an index value of one and the maximum return asset allocation is assigned the maximum index value. Jorion [1992] groups asset allocations together that are in statistically similar regions.

¹² For additional information on Monte Carlo simulation, see Ibbotson Associates [2005].

¹³ Excellent examples of these ad hoc approaches are provided in He and Litterman [1999].