The Kitces Report
Advancing knowledge in financial planning... July 2008

A New Approach to Finding the Efficient Frontier

Especially in today's turbulent markets, the importance of proper asset allocation cannot be understated. However, the foundation of determining asset allocation under the tenets of Modern Portfolio Theory – using a mean-variance optimizer – has developed so many criticisms over the years that it is rarely used in practice by financial planners to design actual client portfolios.

This month's newsletter will explore a new way to apply a mean-variance optimizer, by incorporating a Monte Carlo-style approach to modeling inputs to derive an effective result. The underlying research for this month's newsletter, based originally on the work of Dr. Richard Michaud, was written primarily by David M. Blanchett of Unified Trust Company. You may recognize David's name, as he was one of the winners of the Journal of Financial Planning's 2007 Financial Frontiers award. David was kind enough to share some of his latest research work with me, which I have adapted further for your reading. Although the underlying approach that David demonstrates may be difficult for you to apply directly in your practice, I believe nonetheless that the research itself and the conclusions that you can draw from it are valuable.

Resampled MVO

About the Author

Michael E. Kitces, MSFS, MTAX, CFP®, CLU, ChFC, RHU, REBC, CASL, CWPP™, is the Director of Financial Planning for Pinnacle Advisory Group (www.pinnacleadvisory.com), a private wealth management firm located in Columbia, Maryland. In addition, he is an active writer and speaker, and publishes The Kitces Report and his blog "Nerd's Eye View" through his website www.kitces.com.

Introduction – Modern Portfolio Theory, MVO, and Resampling

Portfolio theory is constantly changing. With increasing computing power it is becoming easier and easier to develop more sophisticated strategies to optimize a portfolio. Modern Portfolio Theory (MPT), originally introduced by Harry Markowitz, is based on the concept that an investor should only invest in optimal (i.e., "efficient") portfolios. A portfolio is efficient, according to Markowitz's original research, if "no other portfolio has higher expected return for a given level of risk or less risk for a given level of expected return." [Markowitz, 1952]

Mean Variance Optimization (MVO) is currently the most common methodology for creating portfolios based on MPT. In essence, the process mathematically determines the optimal weightings amongst a list of available asset classes to yield portfolios that have the highest expected return for a given level of risk. This is computed based on the inputs of the returns and standard deviations of each of the available assets, and the correlations amongst those assets.

While MVO is a popular tool for creating efficient portfolios, it can be a victim of its own inefficiencies. MVO often results in impractical allocations, containing exaggerated exposures to a relatively small number of asset classes, which cannot reasonably be implemented for clients. The key problem with traditional MVO is its sensitivity to the underlying assumptions, where slight changes in the inputs can lead to dramatic differences in optimal portfolio allocations. Anyone who has used an MVO software tool before has likely experienced this phenomenon – a modest 1% change in the projected future returns or standard deviation of an asset class or two can cause the "optimal" weightings to change by upwards of 50% sometimes. However, by incorporating the possibility of uncertainty into the optimization process itself, this extreme sensitivity of MVO to its underlying assumptions can be reduced significantly – a process that can yield far more useful and meaningful MVO results.

If it were possible to perfectly predict the future returns of asset classes, building the ideal portfolio would be easy – quite literally, you would just buy whatever will go up the most over your time horizon. With no uncertainty, there would be little need for MVO. However, perfect information about the future financial markets doesn't exist, and in fact not only are actual future returns unknown, the precise *average* return and *range* of returns isn't known perfectly either. In other words, even though a particular asset class have averaged X% over the past 75 years, doesn't necessarily mean it will average X% over the upcoming 30-year time period (e.g., a client's retirement time horizon). It could be Y% more or Z% less, or perhaps have the same average but be more volatile, or simply be correlated differently with the other asset classes in the portfolio. Because of this, it is inevitable that any forward-looking optimization (i.e., forecast) will be impacted by estimation error. In statistics, estimation error is the term for when the assumptions you started with for a forecast (called "ex-ante" assumptions) are different from what actually turns out in the end (called the "ex-post realized values").

Since the future will always be uncertain (i.e., we still don't *know* what future returns, volatility, and correlations will be), there will always be a risk of having estimation error. Thus, one of the underlying goals of an optimization process should be, ideally, to not only maximize the risk/return tradeoff (what MVO does inherently), but to also reduce the potential size of the estimation error. This is like saying "I know my forecasts of the future could be wrong, but at least I want to be certain that if I'm wrong, I'm not wrong by a lot." Particularly when a large estimation error could result in a client's financial ruin!

This month's newsletter will be exploring the effect of using a Monte-Carlo-style statistical technique called resampling, in the process of doing MVO. In essence, resampling allows the user to incorporate a range of possible inputs for each asset class into the optimization process, thereby allowing the output to optimize not only the risk/return trade-off, but also to minimize the impact of estimation error on the results. As you will see, the results of traditional MVO and resampled MVO are similar. However, there are some important differences, most notably that resampled MVO creates portfolios that are arguably more intuitive, better diversified, and more stable than those created using traditional MVO.

Estimation Error & Resampling

In traditional Mean Variance Optimization (MVO) there are three inputs: returns, standard deviations, and correlations. Each of these inputs is implicitly assumed to be 100% certain. In other words, the value is assumed to be only the precise numerical amount inputted into the MVO software/process, and it is not possible to explicitly incorporate any type of confidence interval, forecast range, or any other type of uncertainty into the inputs. The user must select one and only one value for each input, and as discussed earlier, any change in that singular input can result in a significant change in output. In addition, there will be a continual risk that the user's inputting values turn out to be different than what occurs subsequently in reality, which will create estimation error – the risk that actual results differ from what the model forecasted.

Estimation error is inevitable whenever a forecast is made without perfect information. For example, while it is relatively certain that the sun will rise to tomorrow, it is less certain that it will rain next Tuesday, and even less certain that the standard deviation for domestic small cap value stocks will be 17.26% over the next twenty years. In the simplest of terms, since it is impossible to perfectly predict the future, it is impossible (or at least highly, highly improbable!) that any forecast will be exactly correct. Thus, we should anticipate some estimation error as the realized results turn out to be somewhat different than what was expected. (It is important to note for the context of this research that even using average historical results for your MVO inputs still represents an input that yields a forecast of a certain expected result in the future. To the extent that that forecasted expectation differs from what actually occurs in the future, simply using historical averages as assumptions can still produce a potentially significant estimation error.)

Since estimation error is in essence inevitable, it stands to reason that at least one of the goals of an optimization process should be to not only find an optimal risk/return trade-off, but also to minimize the amount of estimation error that could occur if your forecasted return, standard deviation, or correlation, turn out to be something different than the value you anticipated.

One methodology that minimizes the impacts associated with estimation error for portfolio optimization is resampling. Resampling is a process that was originally introduced by Jorion [1992] and has been expanded on a great deal by Michaud [1998], along with many others. At its core, resampling takes a simulation analysis approach (i.e., Monte Carlo) to portfolio optimization, where the optimal portfolio is based on what yields the most favorable results amongst a variety of scenarios. Each of the scenarios (or runs) represents a potential outcome based upon a forecasted range of possible returns, standard deviations, and correlations, for each of the asset class inputs.

With resampling, it is possible to randomly generate thousands of potential scenarios (or more) that can be used to determine the most efficient portfolio, by analyzing and aggregating together the various scenario outcomes. Typically, the efficient frontier for resampled portfolios are determined by averaging the allocations of each of the optimal portfolios amongst all of the tested scenarios. Thus, for example, if scenario 1 found an optimal small cap value weighting of 10%, scenario 2 had an optimal weighting of 20%, and scenario 3 had an optimal weighting of 5%, the overall average optimal weighting would be 11.67%. Of course, in practice this is multiplied by thousands of scenarios, not merely a 3-scenario average.

While resampling can be used with any portfolio optimization approach (e.g., using semi-standard deviation as the definition of risk), this research addresses the relative benefits of resampling when applied to MVO, since MVO is far-and-away the most popular optimization methodology in use today.

Traditional MVO and Resampling

Traditional MVO tends to overweight those asset classes that have large expected returns, small standard deviations, and negative correlations (and underweight asset classes with lower returns, higher standard deviations, and positive correlations). Not surprisingly, MVO tends to allocate mostly (or entirely) to those select few asset classes that are mathematically "most efficient" on a risk-adjusted basis based on the inputs provided. As a result, the process often totally ignores any of the even-slightly-less-efficient asset categories. As a result, slight changes of the input values that lead one asset class to suddenly have a slightly more efficient risk/return trade-off can result in a dramatic change in optimized weightings. The traditional MVO process – where the most efficient asset classes receive maximal weightings, to the near or total exclusion of most other asset classes – has in fact been called by Michaud [2001] (a proponent of resampling) as an "error maximizer", because it effectively makes the most extreme bets possible for a given set of inputs. Of course, if it were possible to perfectly predict the future, extreme allocations would likely be of little concern, and in fact, they really *would* be most effective. However, perfect information does not exist. Therefore, extreme or biased allocations can lead to suboptimal (and often significantly inferior) portfolio performance to the extent that the original underlying assumptions of the optimization do not hold (i.e., if/when future returns, volatility, or correlations turn out to be slightly different than the original optimization assumed).

While traditional MVO is based on one potential outcome – the one "precisely" predicted by the exact inputs utilized – the number of possible outcomes that can be forecasted with resampling is unlimited. By creating multiple (e.g., thousands of) combinations of return, volatility, and correlations to form an array of potential outcomes (numerous efficient frontiers each based on the combination of inputs provided), and then evaluating amongst all of them, it is possible to incorporate the reality of uncertainty into the optimization.

The fundamental differences between traditional MVO and resampled MVO are analogous to comparing the results of a financial plan based entirely on a single fixed-return projection (a "time-value-of-money approach") to one that utilized Monte Carlo analysis. In other words, utilizing a fixed annual return and linearly growing the portfolio precisely at that exact single return, would be the equivalent of the traditional MVO approach, while doing retirement projections using a Monte Carlo approach would be analogous to resampled MVO. The time-value-of-money analysis consists of one "run" or one outcome, without considering any type of variability in the outcome, providing the user information about one middle-of-the-road outcome and little more. Contrast this with a Monte Carlo simulation, where the number of runs is unlimited (or at least very numerous), allowing the client to consider a virtually infinite number of possibilities. Figure 1 demonstrates the difference between a single "run" time-value-of-money calculation and a 50-run Monte Carlo simulation.

Figure 1: Monte Carlo Simulation vs. Time-Value-of-Money Fixed Return Projection

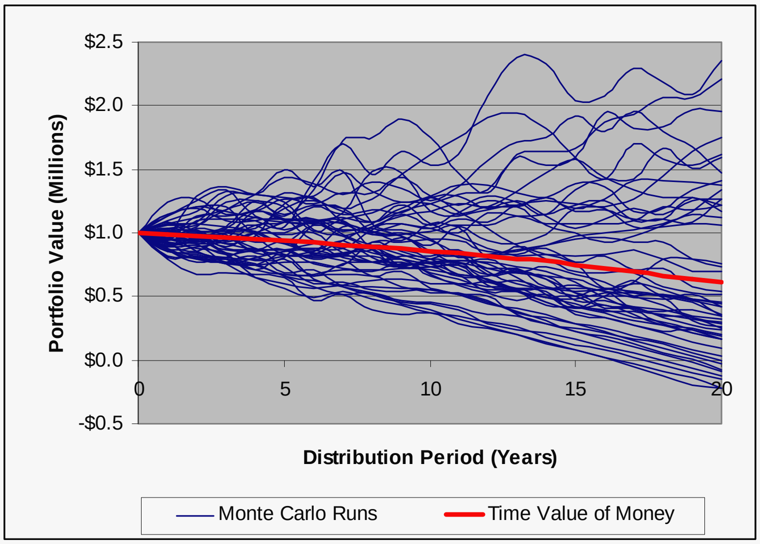

While the time-value-of-money fixed returns line in Figure 1 represents the approximate median result of the Monte Carlo runs, an analysis based entirely on a single projection using one fixed return calculation ignores the possibility of a varying outcome. Monte Carlo tools have become popular because they allow their users to explore the more "colorful" nature of outcomes, as opposed to the simple pass/fail (i.e., black and white) results of a linear retirement projection. While Monte Carlo simulation is becoming more of a standard for financial planning because of its more advanced and textured results and applications, traditional MVO (which is similar in its estimation process to the fixed return approach above) still remains the primary method for portfolio optimization.

Few investment professionals would be willing to state with 100% certainty what the return for an asset class (e.g., domestic small growth) will be in the future (e.g., 9.52%); however, most investment professionals would probably feel more confident if they were allowed to use a range (e.g., 7% - 12%) to forecast the predicted return. The ability to incorporate estimation ranges, and an associated confidence level of results within that range, is a key improvement of resampled MVO over traditional MVO. For example, with traditional MVO, you can only use one return for each asset category (e.g., 8%), but with resampling it is possible to assign a range of values with varying probabilities (e.g., a 20% probability the return will be 7%, a 50% probability the return will be 8%, a 20% probability the return will be 9%, and a 10% probability the return will be 10%, or alternatively to utilize a normal distribution assumption where the return is 8% with a standard deviation of 2%).

MVO Today

The two key reasons that traditional MVO remains popular at all are its aesthetic/visual appeal, and its relative simplicity. A traditional MVO optimizer creates only one efficient frontier that allows the user to easily compare the relative efficiency of two or more portfolios. Traditional MVO is also relatively simple in that it only takes a basic understanding of financial markets and most users can adjust the input parameters to create and/or determine efficient portfolios (i.e., PhD is not required to utilize this approach).

Despite the general acceptance of MVO, some users have dismissed its potential benefits because the results are often so unstable and non-intuitive, where small changes in inputs produce large changes in allocations and the allocations themselves often have "uncomfortably" large exposures to single asset classes. Some users attempt to constrain "unusual" results by placing upper (or lower) limits on the exposure for certain asset classes. However, Frost and Savarino [1988] have noted that while "imposing an upper bound on portfolio weights can improve portfolio performance by reducing estimation bias," the greater the constraints placed on the optimization process, the less pure the results of the optimization. In fact, by imposing overly restrictive constraints on optimization, the financial planner will really do nothing more than force the software to "quantitatively" confirm whatever biases and beliefs the planner already had about what the optimal portfolio should be![1]

A Look at History

One needs only to look back at history to get an idea of the uncertainty associated with "market averages" that are commonly used as inputs for MVO. For example, look at correlation, which is an indication of the strength and direction of a linear relationship between two variables. Traditional MVO requires the user to assume that correlations are fixed for the entire time period, despite the fact that ample research has shown that they are not stable and are constantly changing (e.g., [Coacker 2006 and 2007]). Figure 2 includes the rolling correlations (based on monthly returns) between domestic large blend equity and international equity over the last 80 calendar years (1927-2006) for three different rolling time horizons (5 years, 10 years, and 20 years).

Over the last 80 years the correlation between domestic large blend equity and international equity has varied considerably, although it has clearly increased over time. The maximum ranges in correlations for the 5 year, 10 year, and 20 year rolling periods have been 1.10 (between -.1838 and .9189), 0.82 (between -.1140 and .7083), and 0.62 (between -.0326 and .5838), respectively. The 80-year average correlation is just that, an average. It says nothing of the range of historical returns, the relationship of correlations, or even the fact that correlations have been increasing over time. More advanced relationship techniques, such as copula dependencies (which move beyond linear correlation by capturing the complex inter-relationships between inputs) or GARCH models (Generalized Autoregressive Conditional Heteroskedasticity – a methodology to incorporate time scale parameters into the estimation process) can be used to incorporate these issues into the optimization process, but add a level of complexity that is beyond the scope of this analysis (not to mention the scope of most financial planning practices!). But the bottom line is that to merely assume the historical average correlation between these two asset classes could significantly misstate the future correlations; at a minimum, at least acknowledging and incorporating a range of potential future correlations is more reflective of the past historical relationship.

Beyond just the correlations, the historical returns themselves have also been unstable over time, even over longer periods. Figure 3 below includes information about the rolling annualized historical returns for domestic small value equity over the last 80 years.

Figure 2:

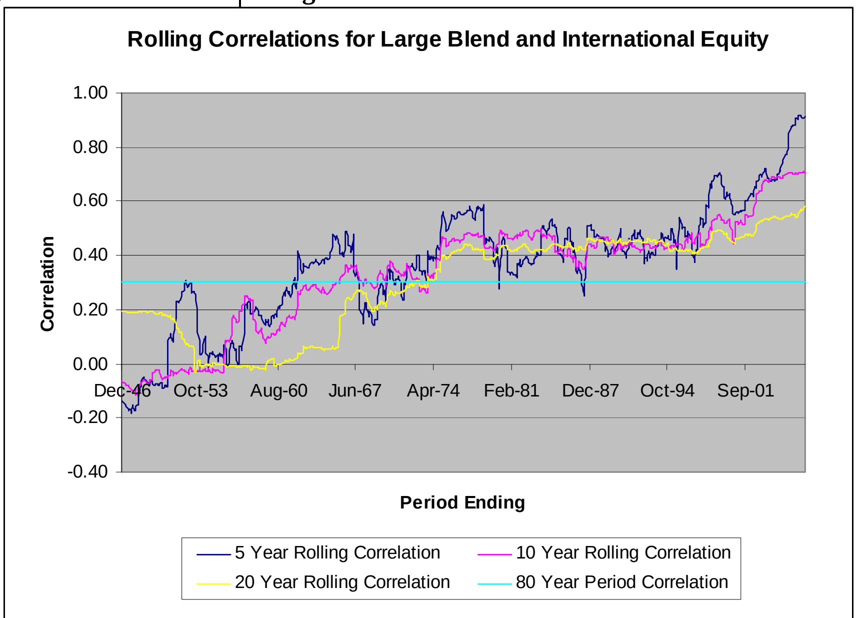

While the 80-year geometric average return for domestic small value equity was 15.57%, there were rolling periods as long as 20 years that still had profoundly different average annualized returns. In fact, 20-year rolling historical returns over the last 80 calendar years have ranged from a low of 6.46% to a

Figure 3:

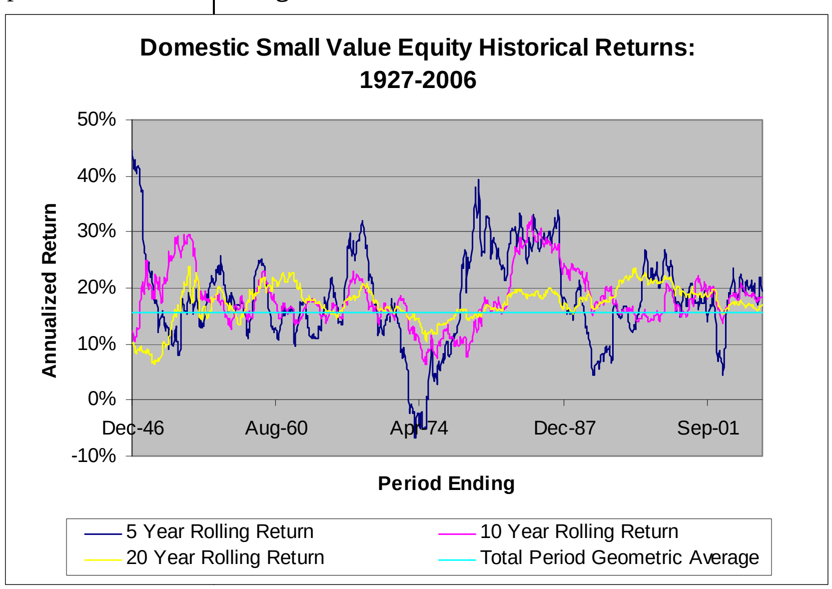

high of 23.93%), producing a standard deviation of 3.26%. Although a 15.57% return may represent a reasonable long-term return for domestic small value, an investor with a shorter time period (e.g., 20 years) is not likely to experience this exact average, but rather some variation of it, and that in turn may affect how desirable exposure to this asset class would be for a particular client portfolio.

All asset classes are likely to experience some varying range of possible returns, standard deviations, and correlations, and different asset classes will have different ranges (e.g., the ranges of return, volatility, and correlation will be different for cash than for domestic large growth). While it is possible to incorporate subjective relative views into portfolio optimization through other techniques (such as the Black-Litterman Model), this is not possible with traditional MVO. Therefore, the ability to incorporate, through resampling, the possibility that a return (or standard deviation, or correlation) may deviate from its long-term average is an important improvement over traditional MVO.

Analysis

There are a variety of resampling strategies of varying complexity that could be used to compare the differences between resampled MVO and traditional MVO. However, for this analysis a resampling tool was built in Microsoft Excel[2] that allowed forecast ranges (i.e., uncertainty) to be incorporated for each of the inputs. The ranges were based on qualitative estimates by the author. The distribution of the ranges for the various inputs were based on normal distributions, although the overall distribution for each asset class was much tighter (i.e., leptokurtic), with each input parameter varying no greater than one absolute value of the range from the expected value. For example, domestic large-cap growth stocks were assumed to be normally distributed around an average annual return of 11.89% with a standard deviation of 19.63%, but the long-term average return for the resampling process could vary no further than from 6.89% to 16.89% (given a range of 5.00%, which is reasonable given that these are supposed to be long-term average returns). The inputs for each run/sample are resubstituted (i.e., adjusted) based on random number generation.

Optimal portfolios were determined for thirteen different overall equity allocations of increasing equity exposure from 20% to 80%, in 5% increments (i.e., an optimal portfolio for a 20% equity allocation, an optimal portfolio for a 25% equity allocation, an optimal portfolio for a 30% equity allocation, etc.). While optimized MVO portfolios are typically found by determining the maximum return for any level of risk (or minimum level of risk for any return) without regard to the equity/fixed allocation, optimal portfolios for this analysis were defined as those with the highest return/risk ratio for a certain equity allocation. For example, the "optimal" portfolio for the 60% equity portfolio would be the various asset class weightings that yielded the highest return/risk ratio, given that the optimized portfolio must end out with a total equity allocation of 60% and a total fixed income allocation of 40%. This approach allows the resampled MVO process to determine an efficient frontier along the entire equity exposure spectrum.

Optimizing based on an equity/fixed allocation, as opposed to a certain return or risk level, was done because client allocation decisions are typically made based on a given equity exposure. For example, a common goal is to have the most efficient 60/40 portfolio. Optimizing based upon equity/fixed income allocation ensures that the portfolio will have the same general allocation regardless of the values for that scenario. The importance of the overall asset allocation decision has been noted by Brinson et al [1986] and Tokat et al [2006]. Optimizing based on a constant equity/fixed allocation also creates portfolios that are easier to compare on an apples-to-apples basis.

The maximum allocation for any asset class for the portfolio optimization was constrained to 25%. This was done to ensure that only allocations that could reasonably be expected to be implemented by a client were considered. While constraints are less important with resampling, the same constraint was applied for both optimization strategies for the sake of comparison (i.e., the constraints were applied similarly for both the resampled and traditional MVO approaches).

It is worth noting that one "run" of the resampled MVO software is effectively the same as a single set of inputs into a traditional MVO analysis. The difference between the two methods is that while traditional MVO was based only on the result of a single set of inputs, the resampled MVO portfolios were predicated on randomly adjusting the input values, resulting in a repeated series of "single" MVO runs with varying inputs, in order to incorporate a reasonable range of uncertainty into the optimization process.

Historical returns over the 28 calendar years (1979-2006) for 13 different asset categories were considered for the analysis. 1979 was selected as the beginning period since it is the inception date for the Russell indexes, which served as the proxies for domestic equity. Alternate proxies were considered, though, during the period for High Yield Bond and Emerging Markets equity to ensure enough data existed for their inclusion in the analysis. Figure 4 shows the asset classes and respective proxies used for the analysis.

Figure 4: Asset Classes and Proxies

	Asset Category	Proxy
1	Cash	3 month Treasury Bill Yield
2	Long-term Govt	LB:Aggr Intmdt Bond
3	Intermediate Bond	LB:Gov Long
4	High Yield Bond	LB:High Yield Intmdt*
5	Domestic Large Growth	Russell 1000 Growth
6	Domestic Large Value	Russell 1000 Value
7	Domestic Small Growth	Russell 2000 Growth
8	Domestic Small Value	Russell 2000 Value
9	EAFE Growth	MSCI:EAFE Growth Index
10	EAFE Value	MSCI:EAFE Value Index
11	International Small	DFA International Small Blend Composite
12	Emerging Markets	MSCI:Emerging Markets**
13	Domestic Real Estate	Lipper:Real Estate

* CAI MF:High Yield Style from 1979-1986
** CAI MF:Emerging Mkts Sty from 1979-1987

The historical returns and standard deviations for each of the asset classes in Figure 4, along with the ranges for the resampling analysis, are included in Appendix I. The historical correlations and correlation ranges used for the analysis are included in Appendix II. Again, note that the resampling ranges are based on qualitative estimates by the author, unlike the returns, standard deviations, and correlations, which are purely historical.

As a starting point for analyzing and comparing the impact of resampled MVO, the efficient frontier generated from the historical values in Appendices I and II using traditional MVO is shown in Figure 5 (right).

Most readers should be familiar with the "bowed" shape of the efficient frontier. The efficient frontier represents those portfolios with the highest return for any given level of risk (where risk is defined as standard deviation). Under traditional MVO it is not possible to have a portfolio that is above (or northwest) the efficient frontier, since the efficient frontier represents the most efficient portfolios possible for any given level of risk or return. Portfolios below the efficient frontier (south and/or east) would be considered "inefficient" since they do not maximize return for a given level of risk.

The equity allocations of the traditional MVO optimized portfolios were heavily skewed towards Domestic Real Estate and Domestic Large Value. This was not surprising given the relative efficiency of the two asset classes during the test period. The average total allocation to Domestic Real Estate and Domestic Large Value was 59.30% of the total allocation to equities. This means that on average, the other six asset classes represented 40% of the equity allocation of the optimal portfolios. Domestic Small Growth and International Growth were not used in any of the traditional MVO portfolios. The fixed income portion of the traditional MVO portfolios was dominated by Cash and Long-term Government Bonds; Intermediate-term Bond and High Yield Bond asset classes are not featured at all in the portfolios with equity allocations greater than 50%.

Figure 5: Efficient Frontier for Traditional MVO

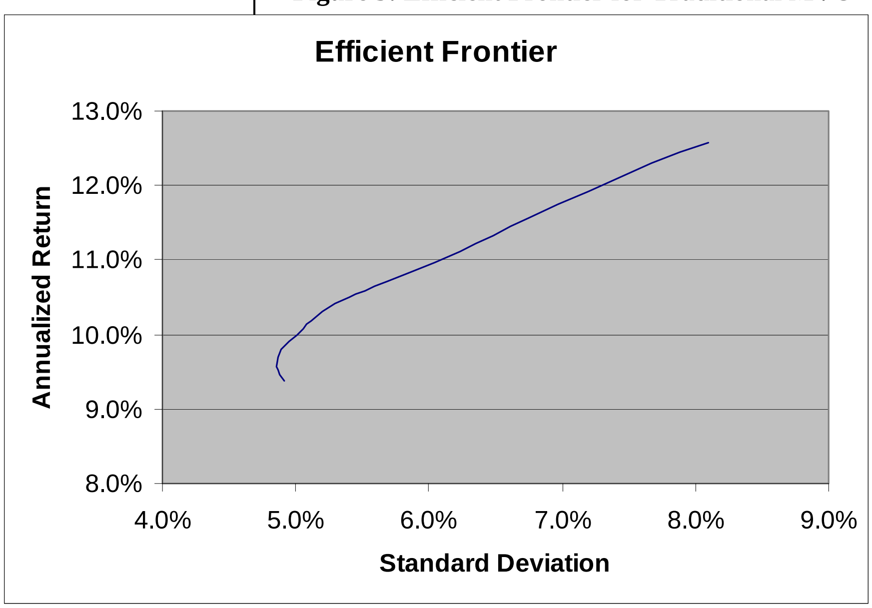

Figures 6 and 7 below contrast the allocations for the optimized traditional MVO portfolios to the optimized resampled MVO portfolios.

Figure 6: Optimal Portfolios defined by Traditional MVO

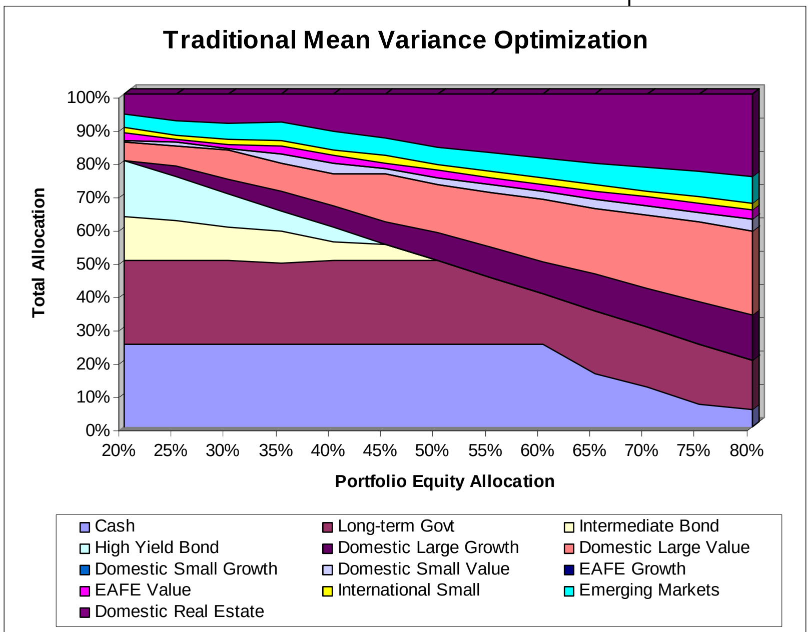

Figure 7: Optimal Portfolios defined by Resampled MVO Based on 2,500 Runs

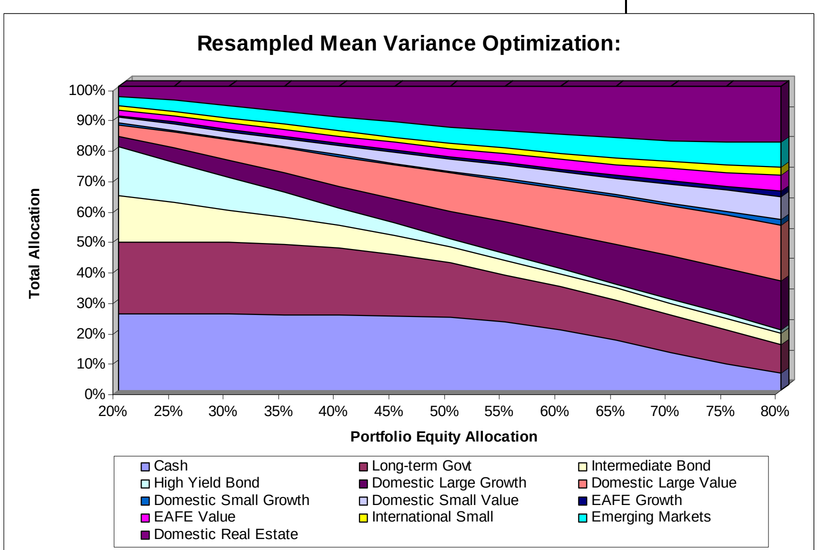

While there are certainly similarities between two methods, there are a number of notable differences. The fact the portfolios of the two methods were similar, though, should not surprise the reader since both methods are based upon the same definition of risk (standard deviation). The allocation differences between the two methods were solely a product of introducing uncertainty in the optimization process (via resampling). Nonetheless, some of the significant and notable differences include:

- While the average equity allocation to Domestic Large Value and Domestic Real Estate was (on average) 59.30% for the traditional MVO portfolios (in Figure 6), the average allocation to these two asset classes for the resampled MVO portfolios was reduced to 47.09%. While these two asset classes are still featured heavily, the resampled MVO portfolios were less sensitive to the "best" asset classes.

- The resampled MVO allocations featured more asset classes than the traditional MVO portfolios. This is a key difference in resampled MVO compared to traditional MVO, in that resampled MVO portfolios tend to be better diversified. For example, the traditional MVO portfolios have no allocation for the 50%+ equity portfolios to Intermediate Bond, High Yield Bond, Domestic Small Growth equity, and International Large Growth equity, while the average resampled MVO allocation to those categories was 4.2%, 1.8%, .9%, and 1.1%, respectively. While these differences may seem minor, the differences between the strategies would likely increase further as more asset classes are considered. In any event, the point still remains that the resampled MVO process yields recommendations of greater diversification than the traditional MVO approach.

- Finally, the relative shifts in exposure to various asset classes from one efficient portfolios risk level to the next (i.e., various equity allocations) was smoother with resampled MVO portfolios than when compared to the traditional MVO portfolios.

Why did the resampled MVO allocations differ from the traditional MVO results? They differ because the resampled MVO portfolios were based on the average of the 2,500 different scenarios considered, while the traditional MVO portfolios were based entirely on the historical data as a single scenario input without accounting for the possibility of any uncertainty. While a user may expect (for example) domestic Small Value to outperform domestic Large Growth, there is a possibility that this may not occur. Also, while a user may expect (for example) domestic Large Value to have a lower standard deviation than domestic Small Growth, again there is a possibility that this may not occur. By introducing the possibility of multiple futures, the best portfolio can be determined given that the realized values may vary from the assumptions used for the optimization.

To give the reader an understanding of how adjusting the inputs can impact the results of MVO, Figure 8 includes the optimized allocations for 100 of the 2,500 runs for the 60/40 resampled MVO test. Note that each simulation has its own set of randomly assigned returns, standard deviations, and correlations, based on the ranges assigned (as shown in Appendices I and II).

Figure 8 – Optimal 60/40 Portfolio Runs using Simulation Optimization

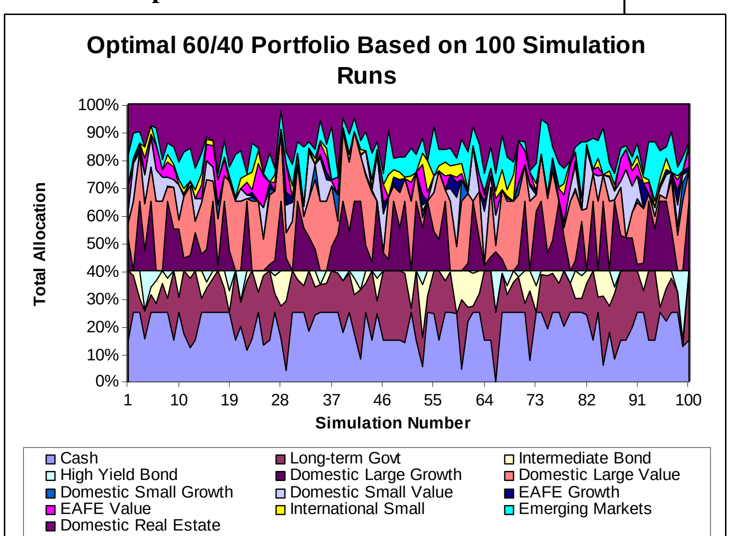

While there are a number of similarities among the optimal portfolios in Figure 8, there are also some notable differences. While Domestic Large Value and Domestic Real Estate where the most featured asset categories, their allocations varied considerably across scenarios. Both had maximum allocations of 25% for at least one of the scenarios; however, the minimum allocation among the 100 scenarios was 0% for Domestic Large Value and 2.08% for Domestic Real Estate. Viewed differently, adjusting the inputs (even slightly) lead to dramatically different optimal portfolio allocations.

Figure 9: Multiple Efficient Frontiers from Resampled MVO

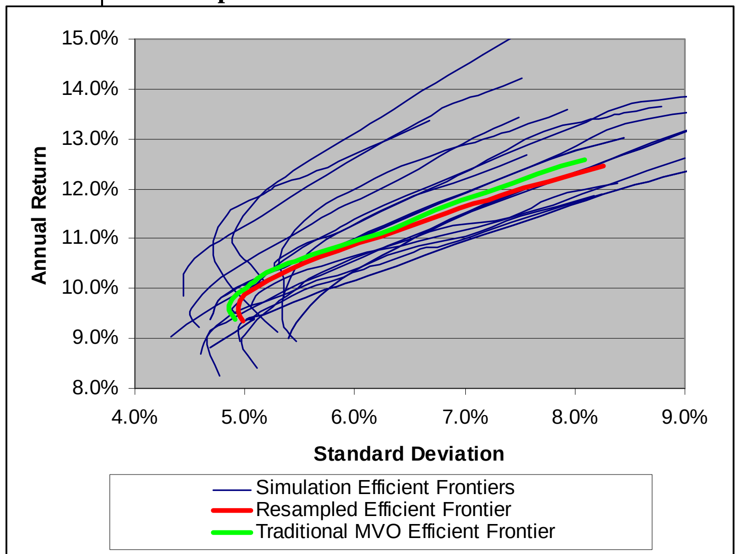

To give the reader an idea of the impact varying MVO assumptions had on the efficient frontier, Figure 9 above compares the efficient frontier for the traditional MVO optimal portfolios to the efficient frontiers created by 20 different simulations. Also shown is the overall resampled MVO efficient portfolio line (established by averaging the exposures for all of the simulations).

As the assumptions change, so does the efficient frontier, although the average resampled MVO efficient frontier is very similar to the traditional MVO frontier. However, the underlying optimal allocations can vary significantly, and the resampled MVO approach yields far more diversified portfolios. By introducing uncertainty into the optimization process, the efficient frontier itself has inherently become uncertain or unstable. Although the resampled MVO efficient frontier is "less efficient" than the traditional MVO efficient frontier (since the resampled MVO frontier line falls slightly south of the traditional MVO frontier line), there were a number of resampled frontiers above and below the traditional MVO efficient frontier based upon the assumptions for that resampling run.

Thus, assuming that each of the 20 different resampled efficient frontiers was equally likely, it would be impossible to suggest that one was "better" than another. After all, the green line is only more efficient than the red line *if* that precise set of assumptions turns out to be the actual future. If it doesn't, though, differences between the predicted returns and actual returns may emerge – and because of the concentrated underlying allocations using the traditional MVO approach, the deviations could be significantly adverse. Therefore, the essence of the resampled MVO approach is to take the average of the 20 efficient frontiers to increase the likelihood that the allocation was the best possible, given the array of potential outcomes.

This is why the resampled MVO approach inherently reduces the risk of significant estimated errors (i.e., reduces the risk that future returns deviated significantly from the expected return).

The Benefits of Resampling

So what do you get for your efforts to determine an optimal portfolio asset allocation using resampled MVO? Research by Michaud [1998] and Fletcher and Hillier [2001] suggest that resampled MVO generally leads to an increase in risk-adjusted performance (defined as the Sharpe Ratio), by literally better managing the "risk" that future returns will be different than anticipated, when compared to traditional MVO. More importantly, Harry Markowitz himself, the father of Modern Portfolio Theory, has even conceded that resampling represents a "better mousetrap" when compared to traditional MVO. Markowitz found that the resampling methodology created by Michuad beat a traditional optimizer in 10 out of 10 cases where the traditional and resampled MVO portfolios were run through a series of tests to see which yielded the better result. The difference in success between the two approaches in Markowitz' study was statistically significant [Chernoff 2003].

Additionally, a test was conducted based upon the research in this paper to determine whether or not any benefit exists when comparing resampled MVO and traditional MVO. The test compared the optimized 60/40 traditional MVO portfolio to the optimized 60/40 portfolio for 1,000 different random scenarios (where the inputs were randomized based upon the values in Appendix I and II). In other words, the return and standard deviation of 1,000 hypothetical return sequences were analyzed using the "optimal" portfolios produced by the traditional and resampled MVO approaches. The resulting difference between the resampled MVO portfolio and traditional MVO portfolio (i.e., improvement due to resampling) is included in Figure 10.

Figure 10: Multiple Efficient Frontiers from Resampled MVO

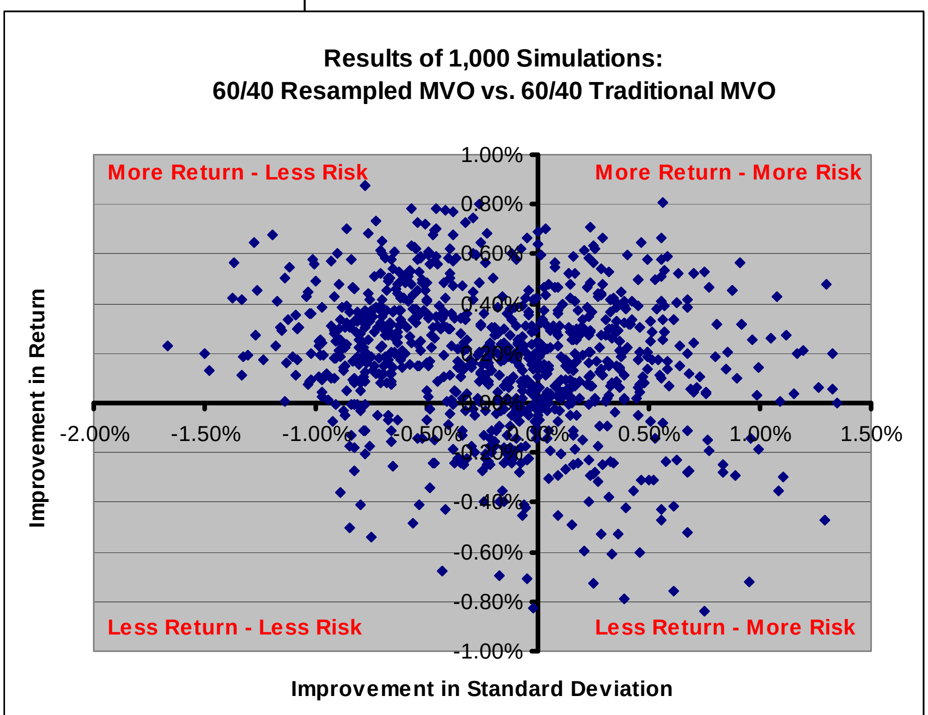

In Figure 10, dots that appear above the horizontal center line had an excess return in favor of the resampled MVO portfolio; dots below the center line favored the traditional MVO portfolio. Likewise, dots to the left of the vertical center line showed less risk (as measured by standard deviation) in the resampled MVO portfolio, while dots to the right had more risk.

Thus, as the reader can see from the concentration of dots, the resampled MVO 60/40 optimized portfolio tended to outperform the traditional MVO 60/40 optimized portfolio on both a risk and return basis for the 1,000 run test – because there are more dots above the line than below the line (higher average returns), and more dots to the left of the line than the right (less risk). In particular, the heavy cluster of results in the northwest quadrant represents a large group of portfolios that had both superior returns *and* less volatility, utilizing the resampled MVO optimized portfolio instead of the one created by traditional MVO.

In fact, the resampled 60/40 portfolio outperformed the traditional 60/40 portfolio on average by 0.182% per year, and had a 0.177% lower average standard deviation. The outperformance differences for both return and standard deviation were statistically different. While 18.2 bps per year in performance may not seem like a material improvement on an absolute basis, it nevertheless represents a "free lunch" to users of resampled MVO (not unlike a similar "free lunch" of improved basis points of performance for Gobind Daryanani's rebalancing methodology; see the January 2008 issue of the Journal of Financial Planning). Also bear in mind, the resampled MVO portfolios were also more diversified, more stable across different equity allocations, and arguably resulted in more "intuitive" portfolios than the traditional MVO portfolios. As a result, future risk and return results for portfolios derived from a resampled MVO process yield results less likely to produce significant "misses" relative to the projected goals (i.e., less potential estimation error), due to the enhanced diversification.

Other Optimization Considerations

When addressing the relative benefits of resampled MVO, it is important to understand the importance of the various inputs from an estimation error perspective. According to research by Chopra and Ziemba [1993], errors in means are about ten times as important as errors in variances, and errors in variances are about twice as important as errors in covariances. Additionally, Best and Grauer [1991] note that "a surprisingly small increase in the mean of just one asset drives half the securities from the portfolio. Yet the portfolio expected return and standard deviation are virtually unchanged." Therefore, the expected return assumptions are the inputs with the greatest importance from a forecasting and estimation error perspective.

A noted criticism of MVO is that standard deviation is an incomplete definition of risk.[3] Markowitz [1959] originally selected variance (which is the square of standard deviation) as the definition of risk for three primary reasons: familiarity, convenience, and cost (because it would have required more expensive tools and computing power to use alternative, "more complex" metrics of risk). Familiarity and convenience are two of the primary reasons for its continued popularity, since cost is no longer a realistic factor. An alternative definition of risk, that was noted in Markowtiz's original research, as well as Nawrocki [1999], Sortino and Satchell [2001], and Kasten and Swisher [2005] is semi-standard deviation, which is also referred to as downside risk. Downside risk focuses on the negative distribution of returns (e.g., those returns below an investor's target, or minimum acceptable, return) based upon the premise investors do not "fear" making too much money (i.e., upside deviation). If asset classes were normally distributed, looking at downside risk would not be necessary (because the standard deviation of the normal distribution would by definition be capturing the volatility of downside risk (since it's exactly the same as the volatility for upside "risk"). However, historical returns have exhibited non-normal characteristics such as kurtosis and skewness – which means that the distribution of returns below 0% (or some other threshold) may be different than the distribution of positive returns – thereby questioning a key assumption of MVO, that the asset classes are normally distributed.

In addition to downside risk, alternate forms of risk measurement such as Value-at-Risk (estimates the maximum dollar loss at a specific probability level) and Expected Shortfall (the average value of total losses beyond a definite probability level) have been introduced as "improvements" to standard deviation. While each definition of risk is going to have its respective pros and cons, it is important to remember that regardless of the definition of risk selected, estimation error will be a factor to some degree. Since the results of any optimization process are only going to be as good as the underlying inputs (i.e., garbage in – garbage out) and since it is impossible to perfectly predict the future, estimation error is a consideration regardless the definition of risk. Therefore, incorporating a methodology to minimize the impact of estimation error, such as resampling, is essential to creating truly optimized portfolios, regardless of the risk parameter used as a benchmark.

So Now What?

Most financial planners probably won't be building their own Excel resampling MVO calculators, which makes it a little difficult to apply some of this research. Nonetheless, there are several important points to take away.

First of all, the optimal portfolios that result from resampled MVO – arguably, a 'superior' optimization process – begin to look much more like the portfolios that financial planners typically build for their clients. So at the very least, this is an affirmation that when we constrain our MVO software, or simply build our allocations in our own 'semi-arbitrary' manner with a level of 'forced' diversification amongst numerous asset classes, we can feel vindicated by a more advanced optimization process that supports the portfolios that have been created. In other words, the sometimes arbitrary decisions that we make as financial planners to include a small allocation to a new asset class that has perceived value as a diversifier is, generally speaking, validated by the results of the resampled MVO process.

Second, these results provide another affirmation for the virtues of using Monte Carlo-style approaches to modeling the uncertainty in the markets to produce a more effective analytical framework for our clients. Just as a Monte Carlo-based approach can yield a more robust analysis of a client's retirement spending plan, so too does it produce a more robust tool for determining optimal asset allocations. Of course, Monte Carlo analysis has its flaws, which can apply both in the MVO and retirement planning context, but as this research shows, it still produces more stable and reasonable results than the overly sensitive alternative linear approach. And no, this doesn't necessarily mean you'll walk your client through every gory detailed step of the resampled MVO approach, but you can educate them more generally on the analytical depth that you use to arrive at the portfolios you design for your clients. And who knows, in the process you might even discover that there's a more efficient asset allocation you should be using for your clients, which in turn can provide a genuine improvement to the risk/return characteristics of their portfolio and their ability to achieve their financial planning goals!

Third, for those of you who are building passive, strategic portfolios, you might consider trying the resampled MVO approach out yourself. Unfortunately, it does not appear that there are very many programs out there to help you get this done. Dr. Michaud himself (the "father" of resampled MVO) has patented the underlying resampled MVO process, and as far as the authors are aware is the only one who makes software commercial available (and reportedly at a cost that makes it impractical for most advisors). However, the underlying principle of a resampled MVO process still has value – at the least, advisors can test the underlying approach by deliberately experimenting themselves with varying inputs to an optimizer and taking an average of the results. Additionally, for the few advisors who are so inclined, a tool can be constructed in Excel to do model resampled MVO, as it really requires only the knowledge in Excel to generate random numbers, determine an optimized portfolio mathematically, and to capture multiple randomized scenarios to aggregate the results.

Caveats

Of course, virtually any theoretical approach to portfolio modeling has its caveats and challenges, and resampled MVO is no exception.

First of all, there are criticisms that the process to actually determine the optimal resampled portfolio – taking the average of the asset class exposures for each of the underlying trail runs – is a somewhat arbitrary solution. Why is taking the average – particularly when the model was already constrained – the right solution? Also, the use of averages in particular can render the result highly susceptible to a small number of significant outlier scenarios – a risk which we already know occurs commonly due to the underlying sensitivity of the entire MVO approach.

In addition, resampled MVO is not the only approach used to resolve some of the underlying problems with the MVO process. Another alternative, although somewhat more mathematically intensive and perhaps even more difficult for financial planners to implement, is the Black-Litterman model (created by Fischer Black, of Black-Scholes model fame, and Robert Litterman of Goldman Sachs) which provides "better" estimates to use for the returns of a mean-variance optimizer.

Furthermore, as discussed above, the resampled MVO process is difficult to apply in practice in a thorough and rigorous manner with clients, due to the lack of software tools available and the complexity relative to the Excel programming skills of most financial planners. (If anyone is aware of a software program that does help financial planners conduct a resampled MVO analysis, please let me know at feedback@kitces.com!).

The final caveat to the resampled MVO approach is that, by its nature, it's also only practical for those who adopt a passive, strategic approach to portfolio design. For those who are more active and tactical, the results of a mean-variance optimization tend to have less value. In fact, even for those who simply use underlying active managers, an MVO approach can be more difficult to apply, as there is always a "risk" that the active manager will change the composition of the portfolio in a manner that makes the underlying original MVO assumptions about expected return and voliatility less relevant. Nonetheless, the resampled MVO approach may at least provide a useful framework to develop a benchmark, or a base optimized portfolio from which the planner can subsequently make tactical adjustments or substitute active managers who are anticipated to provide value in their particular investment space.

Conclusion

Mean Variance Optimization has been a portfolio allocation tool with mixed reviews. It provides an important level of "science" in a profession where allocation decisions have historically been based more on "art." Yet, despite its potential benefits, MVO often results in impractical allocations that are not suitable for clients. Resampling, a Monte Carlo-like optimization approach that allows the user to incorporate uncertainty into the process, reduces the impact of estimation error relative to the traditional MVO approach.

Regardless of the definition of risk used or the optimization methodology, an optimization process is only useful if its inputs are valid. Since it is impossible to perfectly predict the future, any methodology that further minimizes the potential impact of estimation error is worth considering. Resampling is one approach to accomplish this goal; it is a process that tends to lead to more intuitive, better diversified, and more stable portfolios when compared against traditional MVO. Therefore, while resampling is certainly not the last advancement in portfolio theory, it is nevertheless an important one, and notwithstanding some of the implementation challenges, it is something that should be considered in the process of developing optimal portfolio asset allocations.

Endnotes

[1] Along these lines, the user of an optimizer can easily skew the results of the optimization process by adjusting the input forecasts. For example, if the user is fond of Mid-Cap Growth, he or she could change the inputs (even just modestly) by increasing the return, decreasing the standard deviation, and reducing the correlations of Mid-Cap Growth relative to the other asset categories considered (or some variation of three) to ensure it is featured heavily in the resulting "optimal portfolio."

[2] Since the resampling tool was built in Microsoft Excel it was heavily reliant on both the Solver function and Visual Basic for Applications. One problem with using the Solver function in Microsoft Excel, which is based on a general nonlinear programming optimizing algorithm, is that it must be supplied with a "first guess" to the final solution that can impact the results of the optimization process [Nawrocki 2000]. In order to minimize the potential impact of the first guess, the first guess is assumed to be an equal weighted allocation among the investments, based on the target equity/fixed allocation. Therefore, the optimizer would only move away from a diversified, equally weighted portfolio, if a more optimal solution exists. However, since a high level of precision is used for the Solver (.000001), the impact of the first guess is anticipated to be negligible.

[3] One practical reason for the continued popularity of standard deviation when compared to downside risk, is its ability to give the user a "visual" clue as the distribution of returns. For example, assuming a normal distribution, 66% of returns are expected to be within one standard deviation, 95% within two standard deviations. An asset category (or investment) with a downside risk of 5.0% (for example) tells the user little about the overall distribution of expected returns and is only useful as a relative measure (assuming the same MAR, or minimal acceptable return for the calculation).

Works Cited

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." *Review of Financial Studies*, vol. 4, no. 2 (Summer): 315-342.

Brinson, Gary P., L. Randolph Hood, and Gilbert L. Beebower. 1986. "Determinants of Portfolio

Performance." *Financial Analysts Journal*, vol. 42, no. 4 (July/August): 39-44.

Chernoff, Joel. 2003. "Markowitz Says Michaud Has Built a Better Mousetrap." *Pensions & Investments*, vol. 38, no. 26 (December): 3-5.

Chopra V. and W.T. Ziemba. 1993. "The Effect of Errors in Mean and Covariance Estimates on Optimal Portfolio Choice." *Journal of Portfolio Management*, vol. 19, no. 2 (Winter): 6-11.

Coaker, William J. 2007. "Emphasizing Low-Correlated Assets: The Volatility of Correlation." *Journal of Financial Planning*, vol. 20, no. 9 (September): 52-70.

Coaker, William J. 2006. "The Volatility of Correlation: Important Implications for the Asset Allocation Decision." *Journal of Financial Planning*, vol. 19, no. 2 (February): 58-69.

DiBartolomeo, Dan. 1993. "Portfolio Optimization: The Robust Solution." White paper: http://www.northinfo.com/documents/45.pdf.

Fletcher, Jonathan and Joe Hillier. 2001. "An Examination of Resampled Portfolio Efficiency." *Financial Analysts Journal*, vol. 57, no. 5 (September/October): 66-74.

Frost, Peter A., and James E. Savarino. 1988. "For Better Performance: Constrain Portfolio Weights." *Journal of Portfolio Management*, vol. 15, no. 1 (Fall): 29–34.

Jobson J.D., and Bob Korkie. 1981. "Putting. Markowitz Theory to Work." *Journal of Portfolio Management*, vol. 7, no. 4 (Summer): 70-74.

Jorion, Philippe. 1992. "Portfolio Optimization in Practice." *Financial Analysts Journal*, vol. 48, no. 1 (January/February): 68-74.

Kaplan, Paul D. 1991 "Asset Allocation Models Using the Markowitz Approach." White Paper: http://corporate.morningstar.com/ib/documents/MethodologyDocuments/IBBAssociates/MarkowitzApproach.pdf.

Markowitz, Harry. 1952. "Portfolio Selection." *Journal of Finance*, vol. 7, no. 1 (March): 77-91.

Markowitz, H. M. 1959. Portfolio Selection: Efficient Diversification. of Investments. Yale University Press.

Michaud, Richard O. 2001. Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation. Oxford University Press.

Nawrocki, David N. 1999. "A Brief History of Downside Risk Measures," *Journal of Investing*, vol. 8, no. 3 (Fall): 9–25.

Nawrocki, David N. 2000. "Portfolio Optimization, Heuristics and The 'Butterfly Effect.'" *Journal of Financial Planning*, vol. 13, no. 2 (February): 68-75.

Sortino, Frank and Stephen Satchell. 2001. Managing Downside Risk in Financial Markets. Butterworth-Heinemann.

Swisher, Pete K. and Gregory W. Kasten. 2005. "Post-Modern Portfolio Theory." *Journal of Financial Planning*, vol. 18, no. 9 (September): 74-66.

Tokat, Yesim, Nelson Wicas, and Francis M. Kinniry. 2006. "The Asset Allocation Debate: A Review and Reconciliation." *Journal of Financial Planning*, vol. 19, no. 10 (October): 52-61.

Appendix I

	Asset Category	Geometric Return	Range	Standard Deviation	Range
1	Cash	6.30%	2.00%	3.80%	2.00%
2	Long-term Govt	8.72%	3.00%	6.26%	4.00%
3	Intermediate Bond	9.72%	3.00%	12.34%	5.00%
4	High Yield Bond	9.54%	4.00%	11.17%	6.00%
5	Domestic Large Growth	11.89%	5.00%	19.63%	8.00%
6	Domestic Large Value	14.59%	4.00%	13.54%	8.00%
7	Domestic Small Growth	10.32%	5.00%	23.15%	8.00%
8	Domestic Small Value	15.92%	4.00%	17.26%	8.00%
9	EAFE Growth	9.24%	5.00%	21.99%	8.00%
10	EAFE Value	13.47%	4.00%	21.85%	8.00%
11	International Small	13.71%	4.00%	28.96%	10.00%
12	Emerging Markets	13.98%	5.00%	23.08%	10.00%
13	Domestic Real Estate	12.95%	4.00%	14.28%	8.00%

Appendix II

Correlation Coefficients	Cash	Gov	Bond	Hi Yld	Lar Gro	Lar Val	Sm Gro	Sm Val	Intl Gro	Intl Val	Intl Sml	Em Mkt	RE
Cash	1.00												
Long-term Govt	0.21	1.00											
Intermediate Bond	0.06	0.92	1.00										
High Yield Bond	-0.01	0.58	0.50	1.00									
Domestic Large Growth	0.01	0.12	0.13	0.49	1.00								
Domestic Large Value	0.00	0.23	0.23	0.59	0.81	1.00							
Domestic Small Growth	-0.02	0.06	0.06	0.52	0.90	0.81	1.00						
Domestic Small Value	-0.02	0.16	0.14	0.61	0.68	0.88	0.85	1.00					
EAFE Growth	-0.06	0.16	0.18	0.41	0.67	0.66	0.64	0.53	1.00				
EAFE Value	-0.07	0.15	0.15	0.43	0.59	0.68	0.58	0.57	0.91	1.00			
International Small	-0.13	0.01	-0.01	0.50	0.70	0.66	0.79	0.71	0.55	0.56	1.00		
Emerging Markets	-0.04	-0.08	-0.08	0.13	0.21	0.33	0.28	0.32	0.50	0.59	0.35	1.00	
Domestic Real Estate	-0.07	0.42	0.43	0.58	0.41	0.65	0.50	0.70	0.39	0.44	0.45	0.22	1.00

Correlation Range	Cash	Gov	Bond	Hi Yld	Lar Gro	Lar Val	Sm Gro	Sm Val	Intl Gro	Intl Val	Intl Sml	Em Mkt	RE
Cash	n/a												
Long-term Govt	0.30	n/a											
Intermediate Bond	0.30	0.10	n/a										
High Yield Bond	0.30	0.20	0.20	n/a									
Domestic Large Growth	0.30	0.30	0.30	0.20	n/a								
Domestic Large Value	0.30	0.30	0.30	0.20	0.10	n/a							
Domestic Small Growth	0.30	0.30	0.30	0.20	0.10	0.10	n/a						
Domestic Small Value	0.30	0.30	0.30	0.20	0.20	0.10	0.10	n/a					
EAFE Growth	0.30	0.30	0.30	0.20	0.20	0.20	0.20	0.20	n/a				
EAFE Value	0.30	0.30	0.30	0.20	0.20	0.20	0.20	0.20	0.10	n/a			
International Small	0.30	0.30	0.30	0.20	0.20	0.20	0.20	0.20	0.20	0.20	n/a		
Emerging Markets	0.30	0.30	0.30	0.30	0.30	0.20	0.30	0.20	0.20	0.20	0.20	n/a	
Domestic Real Estate	0.30	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.30	n/a