brinton eaton
beyond wealth management

What We Should Be Demanding from Our Asset Allocation Software

By Jerry Miccolis, CFA®, CFP®, FCAS, MAAA,
& Marina Goodman, CFP®

Executive Summary

The deep losses suffered during the market turmoil of 2008-09 exposed operational vulnerabilities across the entire financial services industry, but one area that has thus far not received due attention is the inadequate technology available to financial planners for asset allocation.

Asset allocation and rebalancing are a financial planner's best tools when seeking to protect client wealth. However, despite several decades of research demonstrating the best way to optimize asset allocation and manage risk, the software tools available to do this are simplistic and severely lacking.

Commercially available asset allocation programs – while capable of producing professional-looking, client-friendly reports – are not nearly as good as they could be at capturing market realities during the portfolio construction process. With no other choice to run their operations, some firms design their own spreadsheet-based systems, which can be inefficient; others go without tools to optimize their asset allocations.

This paper is a call to action to those who service and provide technology to financial planners, to encourage the development of tools that can better optimize asset allocation, rebalancing and risk management. It explores the following areas that can be developed or improved upon:

- Optimizing the asset allocation over multiple time periods rather than over a single time period
- Incorporating the effects of rules-based rebalancing into the optimization
- Allowing for investment returns to have non-normal distributions
- Utilizing risk metrics that accurately reflect clients' fears instead of simply the dispersion of returns
- Allowing for volatility clustering – where one volatile time period is more likely to be followed by another

- Treating correlations as dynamic functions that depend on the returns and volatility of assets, instead of single static numbers

It is the authors' hope that software vendors and service providers will recognize the market opportunity: To offer financial planners high-quality, research-based tools to manage risk, protect their clients' wealth and grow assets under management. In today's marketplace, where investors are jittery and skeptical of financial service providers, better technology can provide firms with the operational edge they need to improve performance and better service their clients.

Note: Concepts from this White Paper were featured in a two-part series in Michael Kitces' financial planning industry newsletter, The Kitces Report, *published in December 2009 and January 2010.*

"Solutions should be simple, but not too simple." *-- Albert Einstein*

The events of late 2008 showed us what happens when investors rely on simplistic models and fail to recognize their limitations. The asset allocation programs that are commercially available, even to financial planners and sophisticated investors, while very good at producing impressive, professional-looking reports, are not nearly as good as they can be at capturing reality.

We (the authors) are doubly frustrated. First, because we had to develop our own (inefficient, spreadsheet-based) system years ago to overcome some of these limitations, and have so far been unsuccessful in motivating the industry's professional software vendors to come to the rescue. And second, because these limitations — in the industry's software, not in the theory — have led to a cacophony of complaints that the theory failed us (for example, the oft-repeated, headline-grabbing "Is Asset Allocation Dead?"). The complainers want to throw the baby out with the bathwater.

This, then, is our hopeful call to action.

Considering that asset allocation is by far the most important aspect of managing investments, you would think that asset allocation software would reflect the latest and greatest risk management research available. Unfortunately, commercially available software is several decades of research behind. One reason is that this research may not be well known to financial advisors generally, and therefore there is little demand for improvements to be made.

In this paper, we will review some of the simplistic assumptions that are common in asset allocation software, and the ways that these assumptions can be improved.

Investors' Time Horizons are Not Single-Period

Using multiple time periods is arguably the single most important improvement that asset allocation software developers can make. Without it, none of our other suggested improvements will matter all that much.

When you invest, the return you receive at the end of your investment horizon is really the compounded result of a series of shorter-period returns. What effect does this compounding have on your returns? That depends on two things: the expected return in each period and the volatility of this return.

For example, say you had two investments where the expected return for each investment in each period is 0%, and you are investing for two periods. The first investment actually returned 0% in each period whereas the second returned 10% in one period and -10% in the next. In both cases, the average return was just as expected — 0% for each investment. However, after the two periods, you would have just as much money as you started with in the first investment, but only 99% of your original amount in the second.

This difference between the average, or expected, return (calculated as the arithmetic mean) and how much money you actually have in your pocket in the end (represented by the compound return and calculated as the geometric mean) is called risk drag.

How important is risk drag? Consider the arithmetic and geometric means of the annual total returns for the ten S&P 500 industry sectors from 1990 – 2008. Information technology has by far the highest arithmetic mean of all the sectors. However, the risk drag is so great that its geometric return is lower than many of the other sectors.

1990-2008	Geometric Mean Return	Arithmetic Mean Return	Standard Deviation	Risk Drag
Consumer Discretionary	6.20%	8.63%	22.9%	2.43%
Consumer Staples	9.91%	11.04%	16.3%	1.13%
Energy	11.00%	12.61%	18.3%	1.62%
Financials	7.05%	10.92%	27.6%	3.87%
Health Care	9.94%	12.51%	25.0%	2.57%
Industrials	7.63%	9.64%	19.9%	2.00%
Information Technology	7.78%	13.34%	35.2%	5.56%
Materials	5.56%	7.44%	18.7%	1.88%
Telecom Services	3.95%	7.18%	26.1%	3.23%
Utilities	6.75%	9.27%	23.1%	2.52%

Source: Standard & Poor's, compiled by Morningstar Encorr

We used a popular asset allocation program to determine the "optimal" set of portfolios along the efficient frontier using the above ten sectors. It heavily favored information technology, assigning it the single highest allocation for portfolios in the moderate to aggressive risk ranges of the frontier.

We then looked at how these portfolios actually performed over the same historical period used in the optimization. This is a heavily biased method of analysis, of

course, because it gives such an unfair advantage to the optimized portfolios (since we are using the same time period that the optimization itself was based on). Even under these most unfairly favorable circumstances, the historic performance of the "optimal" portfolios was nowhere close to the efficient frontier.

The reason is that when we calculated the historic performance, we compounded the actual annual returns, thereby introducing the risk drag that investors actually experience. The program completely ignored risk drag, rendering its optimization totally unrealistic, misleading, and woefully inferior.

This example highlights this simple, but little-understood, fact: Making the best one-year decision, repeated over multiple years, may lead you to a very wrong multi-year decision. And let's face it, most every investor invests over multiple years.

This is a serious problem, and one from which virtually all commercially available asset allocation optimizers suffer.

Portfolios Are Rebalanced

An integral aspect of asset allocation is rebalancing. Rebalancing serves two purposes. One is that it allows portfolios to stay on target. Another is that it can enhance the portfolio's return and decrease its volatility.

The key is to find a combination of assets where some zig while others zag, and to rebalance among them.

In the example below, we have constructed two fictitious assets with perfect negative correlation, each of which is very volatile. Diversifying between them would produce a stable portfolio with a return that is the average of the two assets. However, if you also rebalance then you can get an even higher return than either of the assets! At its theoretical best, rebalancing eliminates risk drag. In reality, it can substantially reduce it.

	Asset A		Asset B			
Year	Annual Return	Account Balance	Annual Return	Account Balance	50/50 Mix	Rebalanced 50/50 Mix
0		$100		$100	$100	$100
1	20%	$120	0%	$100	$110	$110
2	-10%	$108	30%	$130	$119	$121
3	30%	$140	-10%	$117	$129	$133
4	-10%	$126	30%	$152	$139	$146
5	30%	$164	-10%	$137	$151	$161
6	-10%	$148	30%	$178	$163	$177
7	30%	$192	-10%	$160	$176	$195
8	-10%	$173	30%	$208	$191	$214
9	30%	$225	-10%	$187	$206	$236
10	0%	$225	20%	$225	$225	$259

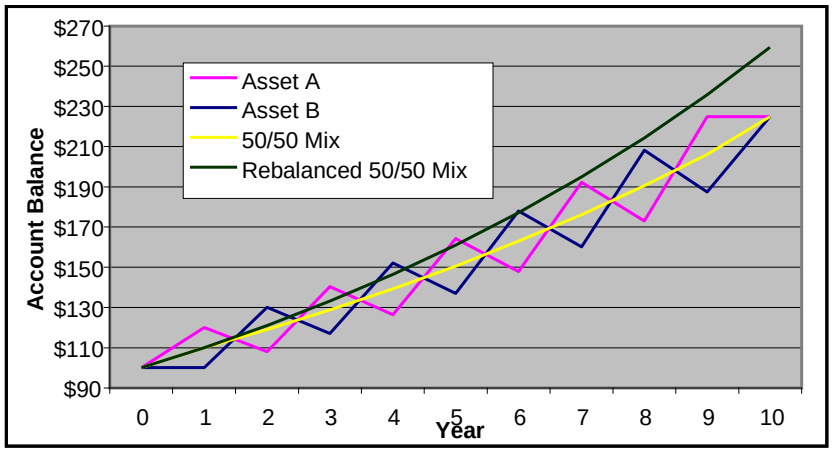

Figuring out which assets work well together and in what proportions they should be held is often counter-intuitive and is impossible for a human being to figure out unassisted. This is an ideal problem for a program to solve.

Here's the good news. Once we can get software vendors to solve the first problem — i.e., multi-period optimization — getting them to incorporate automatic, rules-based rebalancing algorithms should be cinch.

Why? Because solving the first problem requires a big leap. It requires abandoning closed-form, instantaneous, equation-solving approaches to the optimization problem (which single-period optimization makes computationally tractable) and instead adopting a simulation-based, enlightened trial-and-error approach.

Once this leap is made, the rebalancing problem is a no-brainer. We know, because we've done it.

Returns are Not Normal

Many asset allocation programs still assume that investment returns are normally distributed. A number of studies have been done to test if investment returns follow a normal distribution. The universal answer is: They do not. Investment returns have many more extreme values than a normal distribution would allow.

For example, consider the monthly total returns of the S&P 500 from January 1926 to August 2008 (we are purposely excluding returns from September 2008 and later). If these returns were normally distributed then you would expect monthly returns to be below -12% during only 10 of the 992 months, instead this happened 17 times. You would expect returns below -13% during 5 months, not the actual 14. A monthly return below -16% should be so rare that it should have happened only once. It happened 9 times (prior to September 2008!).

The beauty of the normal distribution is that it is easy to deal with algebraically. This should have ceased being a benefit with the advent of the computer age. The downside of using normal distributions is that you greatly under-estimate how often truly disastrous returns occur. An asset allocation optimizer should not assume that returns are normal. It should instead use probability distributions that assign higher probabilities to extreme events.

In our own spreadsheet models, we have used probability distributions ranging from Weibull to logistic to Pareto to extreme value. A promising family of distributions for asset class returns appears to be the so-called stable Paretian distributions.

Risk is Not the Same Thing as Volatility

The ubiquitous measure of risk is standard deviation. However, standard deviation simply measures how diffused the returns are around the mean. Once you equate standard deviation with risk, then you are led to such useless statements as: an investment that can return anywhere from 5% to 25% is riskier than an investment that will reliably lose 3% every year. Standard deviation has been a beloved metric for the same reason as has the normal distribution, it is algebraically easy to deal with. Ease of use is no longer an issue — we should use a risk metric that actually makes investment sense.

There are a number of risk measures that are superior to standard deviation. One general type of risk metric measures the risk of underperformance. For example, you can calculate the probability that the investor might lose money. Another may be

that the investor might earn less than inflation, thereby losing purchasing power. A related measure of risk might be the probability that the investor underperforms a specific benchmark (e.g., a 50/50 weighting between a stock and bond index). This benchmark can be adjusted for investors at different points along the efficient frontier.

Another way to view risk is to quantify the magnitude of "really big" losses. The most comprehensive way to do this is by calculating the Conditional Value at Risk (CVaR) of the portfolio. The intuition behind CVaR is straightforward. Consider the worst returns for a given portfolio, say returns that you expect to happen at most 5% of the time. The expected value of the loss from all these returns is the CVaR at 5%. This can be calculated for any probability of loss. Using CVaR describes investment performance when the performance is poor, and helps distinguish between investments that will cause a loss you can deal with versus investments that will cause you to lose your shirt (and your client's business).

An Aside

The first four problems above (i.e., multi-period optimization, rebalancing, non-normality, and realistic risk metrics) are all easily solvable. We did so long ago with our spreadsheet-based models. But these models are slow and inefficient, and could use some professional programming expertise to be practical for use across the industry. Moreover, we believe the following two problems are beyond our current computational capabilities.

Again, for the benefit of the financial planning profession at large, we beseech the industry's software service providers to take up the cause, and we would be pleased to render whatever guidance and assistance they require.

Correlations are Not Linear

"Correlations" in asset allocation software may not mean what you think they do. There are several methods of calculating correlation and software typically uses the "Pearson Rho" method. This is the standard method taught in statistics class. It assumes that whatever two random variables are being correlated have a linear relationship with each other. We know that asset class returns do not have a linear relationship with one another.

Another nasty feature of the Pearson Rho correlation calculation is that it is not robust — a single piece of data can have a very large effect on the calculation. For

example, the five-year Pearson Rho correlation[1] between Dow Jones US REIT and S&P 600 Small Cap index for the period ending in September 2008 was 58%. If you move the period one month forward so that it ends in October 2008 then it jumps to 74%.

One way to solve both of the above problems is to use the "Spearman Rank" correlation. You no longer have to implicitly assume that the investment returns are linearly related, and the statistic will not go haywire due to a few months of extreme data. In the above example, the Spearman Rank correlations increased from 56% to 60%.

The methodology is fairly straightforward. First you calculate the rank of each return. For example, the Dow Jones US REIT index returned 1.48% in October 2003, which was the 34th highest return during the period October 2003 – September 2008. The 1.48% return is replaced with the rank, 34. Rank can be easily calculated using Excel. Once you calculate the rank for all returns, take the Pearson Rho correlation of the ranks. This is something that would be computationally simple for an asset allocation program to incorporate.

Source: Morningstar Encorr			**Ranks of Returns**	
	Dow Jones US REIT	S&P 600 Small Cap	Dow Jones US REIT	S&P 600 Small Cap
Oct-03	1.48%	8.67%	34	1
Nov-03	4.41%	3.78%	20	15
Dec-03	3.00%	1.77%	27	28

An even better, more realistic, way to treat correlations is through copula functions. Instead of correlation between two securities being reduced to a single number, it is now expressed as a function. This allows for the correlation to change depending on the return of one of the assets.

This opens up a whole area of research into strategic and tactical adjustments to be incorporated into asset allocation models. For example, this would allow the asset allocation to model the tendency of correlations to increase to 1 when returns are very negative (sound familiar?).

It would also allow for correlations to behave one way when returns for one of the asset classes is positive and a different way when it is negative. Hopefully, copula functions and how they can be used by investment managers will be topics of much future research in our industry — they are already quite popular in the actuarial community.

[1] The correlation was taken between the monthly total returns of the indexes over a 60 month period.

Volatility Comes in Clusters

Asset allocation programs assume that volatility in one time period is independent of volatility in the following period. However, history shows that periods of high volatility are not uniformly distributed over time. One month of high volatility is likely to be followed by another month of high volatility.

In the graph below, the S&P 500 monthly total returns are graphed from January 1926 to June 2009. The horizontal blue bars are the average monthly return +/- 1 ½ standard deviations. The vertical yellow bands highlight when a monthly return was more than 1 ½ standard deviations away from the mean (i.e., outside the blue bars). Notice how the yellow bands are not uniformly distributed throughout the time period, but rather they tend to cluster. Volatility in one month shows that there is a higher chance of the next month being volatile as well.

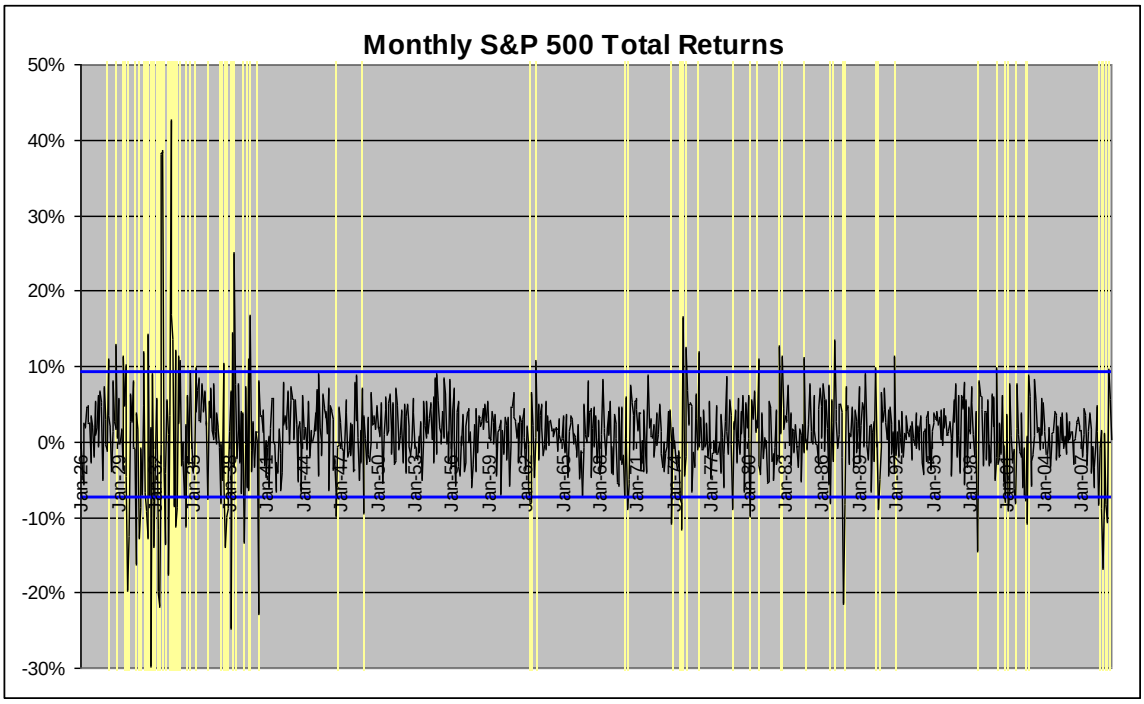

Source: Standard & Poor's, compiled by Morningstar Encorr, analyzed by Brinton Eaton

GARCH (Generalized Auto-Regressive Conditional Heteroskedasticity) models are a way to generate returns where the volatility in one period is related to the volatility in previous periods. These models are well known in operations research circles and widely used in applied statistics. Why not use them in financial modeling as well?

Conclusion

Using multiple time periods, and incorporating rules-based rebalancing, would change the very nature of the optimization process and guide the asset allocation program to look for the combination of assets that would produce the highest returns for the lowest risk in the real world. It would greatly help reduce wrong decisions.

Using non-normal, "fat-tailed" probability distributions for investment returns would more realistically model the chances of large losses.

Measuring risk with respect to a benchmark as well as in terms of the magnitude of a "worst case scenario" loss would allow the asset allocation program to think about risk the way a real investor does and to focus on its most pertinent features.

Incorporating a more sophisticated understanding of correlations and volatility would help guide investment managers to identify periods where correlations and volatility are likely to increase and provide a means of developing and analyzing informed tactical adjustments to strategic asset allocations.

Incorporating these features in asset allocation software would greatly improve the quality of the decisions we make on behalf of our clients.

As we mentioned, we implemented the improvements mentioned early in this article in our spreadsheet-based models some time ago. But our models are slow and clunky — we are not trained programmers. Some of the later-mentioned improvements are simply beyond our programming capabilities at the moment. We hope that professional software developers take up the challenge to make all the improvements we suggest. The financial planning community, and the investing public at large, will forever be better for it.

About the Authors

Jerry A. Miccolis, CFA®, CFP®, FCAS, MAAA, is a co-owner of and Senior Financial Advisor with Brinton Eaton, a full-service, fee-only firm providing integrated financial panning, tax advisory, and investment management services to high-net-worth individuals and non-profit institutions. Marina Goodman, CFP®, is an Analyst with Brinton Eaton. Mr. Miccolis is co-author of, and Ms. Goodman is a contributor to, Asset Allocation For Dummies® (Wiley, 2009).