COMPLEMENTARY GENIUS

Building portfolios with “outside-the-box” mutual funds

BY GARY A. MILLER, CFA
AND SCOTT MACKILLOP

Frontier Asset Management, LLC
201 North Connor Street, Suite 250
Sheridan, WY 82801
307-673-5675

March 2009

Introduction

Harry Markowitz's seminal work on the efficient market hypothesis and modern portfolio theory, for which he won the Nobel Prize in 1990, has transformed asset management. Legions of financial advisors now construct "efficient" portfolios using the tenets of modern portfolio theory. Markowitz's work was profound in many ways. However, efforts to implement his ideas have presented practitioners with challenges and led to investment practices that are flawed and undermine the potential of those ideas.

In response to these challenges, we have developed an investment process that we call Complementary Genius. In this article we discuss the beliefs underlying our process, how it works and why we think it solves some of the most common problems associated with current portfolio management practices.

The Standard Approach

Before we discuss the specifics of Complementary Genius, let's review the standard approach to managing asset allocation-based investment portfolios. We recognize that financial advisors employ numerous variations of this approach, but the description below captures many of today's most common portfolio management practices.

The standard approach typically involves the following four steps:

Establish Asset Allocation Strategy. The first step is to determine an appropriate asset allocation mix that is "efficient" and falls within the investor's risk tolerance parameters. This entails identifying an appropriate set of asset classes that will be used in developing the asset allocation strategy. An index is selected to represent each asset class. Assumptions are made about the future expected returns, volatilities and correlations of the asset classes. Then a computer determines the optimal mix of asset classes for the target risk level. Often the mix remains fixed throughout time.

Manager Selection. The next step is manager selection. The goal is to find a superior manager to represent each asset class. Managers are screened based on their performance relative to the index associated with the asset class they will represent in the portfolio. Performance is also examined to be sure it closely tracks that index over time. Managers who exhibit excessive style drift may be eliminated. There are other considerations, too, but performance relative to a single benchmark and style purity are the most important traits. Index funds are now popular as a way to represent an asset class in a portfolio. Although they do not beat the index, they track it nearly perfectly.

Ongoing Monitoring. Once a manager is selected to represent each asset class, the process of ongoing monitoring begins. There are two primary ways a manager can get in trouble. The first is poor performance relative to the assigned index. The other is style drift. If a manager's performance falters relative to its assigned index, or the manager wanders too far from its assigned "style box," it may be replaced by another manager whose performance better meets the applicable criteria.

Rebalancing. The portfolio is also monitored to ensure the integrity of the asset allocation strategy. The strategy includes a determination about what percentage of the portfolio each asset class should represent and establishes tolerance ranges. Rebalancing is initiated if an asset class falls outside of the established percentage ranges. In practical terms, this means that each manager's representation in the portfolio is monitored. If a manager's representation falls within the established range, no rebalancing is necessary. If the manager falls outside of that range, rebalancing is initiated. Some rebalancing strategies are done periodically (e.g., quarterly, annually, etc.). In these strategies, rebalancing is initiated periodically regardless of the extent to which the manager's representation differs from the established asset class targets.

Problems with the Standard Approach

There are two types of problems that arise from the standard approach to portfolio management. The first is problems associated with asset allocation methodology. The second is problems associated with manager selection and monitoring.

Rearview Mirror Asset Allocation. Let's look first at the problems with the asset allocation process. Many advisors use historic (backward-looking) asset class returns, volatilities and correlations to estimate future returns, volatilities and correlations. They study long time periods and develop historical averages, which they use as the inputs for their optimization process. These inputs reflect no analysis of the future at all. They rarely, if ever, change regardless of prevailing market conditions. Yet they are used to develop strategies for a financial future that is constantly changing.

There is No "Right" Answer. The optimization process is extremely sensitive to variations in data input. A change of only 0.5% per year in the expected return of a single asset class can change the allocation for that asset class in a portfolio by over 25%! Yet advisors often use different inputs in constructing portfolios, causing the asset class composition of efficient portfolios created by two advisors to differ significantly.

There are a number of reasons for this. First, advisors study different time periods in deriving their historic asset class return, volatility and correlation averages. Thus, they come up with different inputs to use in the optimization process.

In addition, asset allocation strategies are typically developed using indices to represent the different asset classes. The firms that create these indices define each asset class in a slightly different way and include different securities in them. This results in performance differences between indices that purportedly represent the same asset class. So advisors using different indices can derive differing optimization inputs.

It's Not Magic. An asset allocation strategy is not magic. It is a mathematical solution created by a computer based on a number of assumptions provided by the advisor. Unfortunately, reality almost never lines up with the assumptions that were made in developing the strategy. Thus, there will always be alternative strategies that would have performed as well or better than the ones produced by optimizers.

Compromising in the Name of Comfort. Left unconstrained, optimizers tend to create asset allocation solutions that do not include all of the asset classes an advisor may wish to use in a portfolio. Instead, they create solutions that favor the asset classes that look best based on the specific inputs that are used. Most advisors, being aware of this problem, constrain their optimizers so that they force more broadly diversified solutions that "feel about right." They know that their clients will be more comfortable with portfolios that emphasize asset classes they are most familiar with. These portfolios may serve a useful purpose, but they are not optimal solutions.

For all of these reasons, slavish adherence to the precise asset allocation recommendations produced by our computers makes little sense. These solutions may be a good starting place for portfolio construction, but they are not a silver bullet. Yet many of today's common portfolio management practices are directed at rigidly maintaining a particular asset allocation to the exclusion of more important issues.

The Problem with Single Index Comparisons. Now let's turn to some of the problems associated with manager selection and monitoring. First, it is very difficult to find managers that consistently outperform a single index. There is a good reason for this. Many of the best managers pursue opportunity wherever they find it. They do so without regard to the artificial asset class categorizations that have been imposed on them by consultants and data providers. Therefore, their performance may differ markedly from a single index in the short-term.

The same problem arises when we assess managers based on the degree to which they drift from the style box we want them to represent. When managers drift, the problem is often not that these managers have lost their discipline or forgotten their mandate. They are simply seeking opportunity and ignoring the measuring sticks others use to gauge their performance. As Craig Callahan and C. Thomas Howard have demonstrated, style boxes help in describing the characteristics of a portfolio—that is, what types of securities a portfolio holds—but tell us little about a manager's investment approach. (*Outside the Box*, Callahan and Howard, 2005.)

Also keep in mind that there is no evidence that managers who invest within a single style box are better managers than those who do not. In fact, research seems to suggest the opposite. A study done by Professor Russ Wermers of the University of Maryland found that the managers with the most style drift added significantly more value than managers with little style drift. (*A Matter of Style: The Causes and Consequences of Style Drift in Institutional Portfolios*, Wermers 2002.) This study seems to indicate that we should value, rather than punish, independent thinking.

Changing Classifications. When a manager seeks opportunities outside of its assigned asset class or style, the fund is often re-categorized—even if the manager has not changed its investment approach. For example, in the last ten years, Bill Miller's Legg Mason Value Trust has been classified as Large Growth, Large Blend and Large

Value. This confounds the efforts of those using the standard approach to maintain the integrity of their asset allocation strategies. But why worry so much about the labels? Shouldn't we worry more about figuring out how to determine if the managers are adding value, and, if they are adding value, how to use them in a client's portfolio?

Choosing in Isolation. Another difficulty with the standard approach is that managers are usually chosen in isolation, with no regard for how they will fit together in a portfolio. As legendary football coach Knute Rockne once said when asked why he was so successful, "I play my best eleven, not my eleven best." Greater thought should be given to how the team will work together once it is assembled.

Here is a simple example. Bond Fund of America and PIMCO Total Return are both in Morningstar's Intermediate-Term Bond category. According to Morningstar, they both reside in the High Quality/Intermediate Duration style box. They even have somewhat similar risk-adjusted returns. For the five years ending December 2007 Bond Fund of America returned 5.8% annualized, beating PIMCO Total Return's 5.3%. Plus, it did so with less volatility. So viewed individually, Bond Fund of America might look a little better. However, Bond Fund of America consistently holds lower quality bonds than PIMCO, even though they are both in the same Morningstar style box. When stocks are doing poorly, the lower quality bonds do poorly, too. And that's exactly what happened in 2008 as stocks tumbled. Bond Fund of America lost 12.3%, while PIMCO returned a positive 4.8%. So even though Bond Fund of America might look better on its own, PIMCO Total Return is a better choice in a portfolio with a large equity allocation.

A Better Approach

To solve these problems, we must think outside the box. Complementary Genius is our approach to portfolio management that deals with some of the problems we've just described. Complementary Genius is founded on a number of core beliefs.

Asset allocation is an important tool, but is not an end in itself. Modern portfolio theory helps us use assumptions about the future to create asset allocation strategies that we believe will perform well down the road. The better we get at developing estimates about the future, the closer to "optimal" our asset allocation strategies will be. However, as long as there is a gap between our estimates and what actually transpires—and there always will be—our portfolios are not necessarily hurt, and may even be helped, by deviations from the asset allocation strategy.

As in any field of endeavor, some managers are more skilled than others. Our efforts should be directed at finding "genius" managers who possess demonstrable skill, rather than at finding those that best represent an asset class or demonstrate the highest degree of style purity. We want managers who have the talent and the flexibility to consistently add value through changing market environments. Managers who show no style drift lock us in to the shortcomings of our asset allocation strategies. Managers who are more flexible can help us overcome those shortcomings by providing us with a valuable second opinion.

Manager selection should be about creating a team that performs well together. The performance of a portfolio represents the collective returns of all of the managers. Every manager has strengths and weaknesses. So the goal should be finding a group of managers whose unique approaches to investing complement one another. Manager selection should not be about selecting isolated representatives of particular asset classes. It should be an exercise in team-building.

Finding Genius

Quantitative Analysis. Let's look at how the Complementary Genius process works in practice. The first step in the process is manager selection. In trying to identify "genius" managers we use a quantitative process called returns-based style analysis. It helps us to assess whether a manager has added value historically through its unique approach to investing. Style analysis was developed concurrently by Nobel Laureate Dr. William Sharpe and the co-author of this article, Gary A. Miller, CFA.

A fundamental precept of style analysis is that a single index cannot be robust enough to explain the returns of an actively managed portfolio. In order to produce a more accurate, customized benchmark, we compare the returns of each fund against the returns of a group of asset class indices. The result is a set of style factors that better represents how the fund "acted" in the past than a comparison to a single index.

Style factors are really just a collection of indices combined proportionately in a way that best describes the manager's performance pattern over time—a sort of historical average showing how a manager invested in different environments. Because they reflect more clearly how the manager actually invested, style factors provide a more reliable basis for assessing manager skill than performance relative to a single index.

Style analysis is a powerful approach to manager selection for a number of reasons. First, all managers can be considered, not just those that fit neatly into specific style categories. Many great managers simply do not conform to single style boxes. So style analysis greatly expands the universe of funds we can consider in building portfolios.

Further, each manager is compared to its own unique benchmark, regardless of how it invests. A manager that outperforms a single, arbitrary index may simply be lucky. A manager that consistently outperforms a robust best-fit benchmark consisting of multiple indices is more likely to be skillful.

Style analysis can, in some ways, tell us more about the types of securities a manager owns than an analysis of its holdings would. For example, funds that buy stocks with below-average debt or above-average cash on their books have style factors that show that they hold some cash in their portfolios, even if they don't. You can't see it in the holdings, but it's there. It's in the stocks! Funds that own stocks with high dividend yields have fixed income factors. No bonds are evident in the holdings, but the fund "acts" like they are there. And so on. Dynamic style analysis, an advanced offshoot of

returns-based style analysis developed by Mr. Miller, can even tell us how a fund might react in the future to different market environments.

Style analysis is not perfect. The standard style indices are based on market capitalization and growth-value characteristics, so they do not reflect other characteristics that investment managers might use. However, style analysis still appears to be at least as good, or better, than other options for an initial assessment.

Consistency. Once we understand how a fund has acted in the past, we must have good reason to believe that its manager will use the same style of investing in the future. When we refer to consistency of investment style, we are not talking about style boxes, but rather a methodology of investing. We will only use a manager in a portfolio if we can count on it continuing its approach to investing—whatever that might be—in the future. We pay as much attention to consistency of methodology as we do to identifying skill itself.

Qualitative Screens. We also screen managers based on a number of qualitative factors that are a very important part of the Complementary Genius process. We use these factors to identify managers that are independent, experienced, dedicated and flexible specialists with unique investment approaches that can add value to our portfolios. Combining these subjective screening criteria with more quantitative approaches underscores the fact that this process is truly part art and part science.

Creating Complementary Combinations

Once we have identified a stable of "genius" investment managers, we use our Complementary Genius process to combine them in portfolios so their unique investment styles complement one another. We believe that *how* you combine managers in a portfolio is as important as *which* managers you choose to include in the portfolio. In our process, combining managers is really a two-step operation.

Establishing Targets. The first step is to establish the risk and return targets for the portfolio. To do this, we create a "Long-Term Allocation" for each portfolio. The Long-Term Allocation is a combination of asset classes, each represented by an appropriate index that we believe represents the most "efficient" mix over the long-term for the chosen level of risk. Creating the Long-Term Allocation looks very much like the process of establishing an asset allocation strategy under the standard approach to portfolio construction. However, we don't stop there.

At any point in time, our estimates of the future returns, volatilities and correlations for each asset class may vary from those we use in developing our Long-Term Allocations. The assumptions that we use in creating our Long-Term Allocations reflect many decades of history and do not take the current market environment into account. Yet, we know that, year in and year out, the markets will not consistently deliver the investment experience that one might expect from looking solely at the long-term

historical averages. So we develop what we call “Strategic Allocations” that we believe are the best allocations going forward from today, given current conditions.

The estimates that we use to create our Strategic Allocations are derived using economically sound methodologies, not backward-looking historical averages. They are based on an understanding of history, but also reflect current conditions. For example, our return expectations for US stocks change as the normalized earnings yield of that asset class changes. Thus, our current return expectation for US stocks is affected by the current recession since stock prices have fallen more than normalized earnings have. Using the standard approach, the current recession would have little impact on the outlook for stocks, due to its minimal impact on the long-term historical averages.

We have developed distinct models for each asset class, not just US stocks, based on factors that actually drive the returns and volatilities of those asset classes. Our models are based on the actual events that cause change in asset class performance, rather than on factors that simply reflect the effects of those changes.

We derive our expectations about correlations based on periods of extreme market movements, when correlations tend to be higher. The standard method can surprise investors during market downturns because it tends to underestimate the extent to which asset classes are correlated at such periods. Many investors weren’t prepared to have all of the asset classes in their portfolio fall at the same time like they did in 2008.

Because our models for estimating future expected returns, volatilities and correlations utilize variables that change over time, our Strategic Allocations do, too. It seems self-evident that a single asset allocation strategy cannot be expected to perform well in all circumstances. Yet it is common practice to leave an asset allocation strategy unchanged in the face of dramatically different market and economic conditions. We allow our Strategic Allocations to reflect factors that could impact future performance.

The Strategy is Only the Starting Place. Another difference between the standard approach and Complementary Genius is the purpose behind the creation of our Strategic Allocations. Under the standard approach, the asset allocation strategy becomes a rigid framework that drives portfolio management activity. Managers are assessed based on performance relative to a single index. They are punished for style drift. Deviations from the strategy are eliminated through rebalancing. Adhering to the strategy becomes an end in itself.

Our Strategic Allocation is more like a straw man that we try to beat when we combine our “genius” managers in portfolios. We look at the return patterns of the Strategic Allocation in different market environments over long time periods. Then we seek a combination of managers that we believe will outperform the current Strategic Allocation in all environments, particularly in poor markets. This combination of managers usually has asset class characteristic that are close to the Strategic Allocation, but that is not our goal. Variations from the Strategic Allocation are acceptable as long as the

managers produce good results. We think that clients are more interested in good returns than they are in maintaining a particular asset allocation strategy.

There are usually between 100 and 125 “genius” managers on our approved list, so there are always many possible combinations to consider. We use an optimization process that tests thousands of possible combinations. In effect, we treat each manager as though it were a separate and distinct asset class with its own return, volatility and correlation patterns. This gives us far more possibilities for generating return and reducing risk than if we were to use only traditional asset classes.

To get on our approved list, a manager must have demonstrated an ability to consistently add value relative to its unique style benchmark. To be used in a portfolio, however, a manager’s unique investment approach must improve the overall portfolio. Being a good manager is not enough. We are looking for good managers that will make the team stronger.

Our process doesn’t stop once we find a combination of managers that we believe can beat the current Strategic Allocation. We continue to ask the question: “Is there a better combination of managers than the one currently in place?” We answer this question by re-running our optimization process frequently (as often as monthly) to see if we can find a better combination of managers than the one we have in place. This may lead us to replace a perfectly good fund with another fund if that change improves the portfolio. We may also include a fund in the portfolio that has not been performing well recently, if it has the potential to help the portfolio when conditions change down the road. We don’t really care how a fund has performed relative to a particular index. If its unique investment approach improves the portfolio, we will use the fund.

We do not rebalance our portfolios in the traditional sense. As market movements change the allocation of funds in the portfolio, the portfolio will exhibit different future expected return characteristics. If those characteristics do not appear as good as those of our straw man, the current Strategic Allocation, we will make changes to the portfolio. For instance, after a long period of small value stocks outperforming large growth stocks, the portfolio’s expected return pattern may show that it is more likely to underperform the Strategic Allocation when large growth is performing well. If so, we may sell a fund with a small value orientation and buy one with more of a large growth orientation to improve the portfolio’s future prospects. Rebalancing takes place automatically, in this manner.

A Constant Search for the Best Combination. The Complementary Genius process is highly dynamic. We continuously search for new “genius” managers to add to our approved list. We delete managers from the list if they fail to meet our quantitative or qualitative screens. Our Strategic Allocations are adjusted as our expectations about the future change. But the ultimate question remains the same: “What combination of ‘genius’ managers is most likely to beat the current Strategic Allocation in all markets?”

Conclusion

The work of Harry Markowitz provides a strong conceptual framework for managing portfolios. However, many of the practices that have evolved from his theories are keeping financial advisors from producing the best results for their clients.

The beliefs and practices incorporated into Complementary Genius are designed to help overcome some of those shortcomings. We believe that some managers truly are more skilled than others, and that *how* you combine managers in a portfolio is as important as *which* managers you combine. Asset allocation is not an end in itself, but it is possible to improve its usefulness by continuously re-thinking allocation estimates for the future.

One question you may ask is: “How well does Complementary Genius work in real life?” We think it works pretty well. We have been managing mutual fund portfolios for over 20 years and have used some form of Complementary Genius since 1999. As of the end of 2008, our core equity-risk portfolio, Long-Term Growth, had beaten its benchmark, the S&P 500, every single calendar year for the past 10 years, after fees. Only 7 out of over 3,000 mutual funds with ten-year histories in Morningstar’s database have done so. We think that’s pretty strong proof that the theory works in the real world.