The Index Investor

Why Pay More for Less?

Get Ready for Our New Site

First off this month, we'd like to bring all our readers up to date on the redesign of our site that we've been working on for much of this year. We have now progressed to the beta testing stage, and hope to take it live at the beginning of January. Quite frankly, we're very excited about it. Among its many new features will be much easier navigation, a complete library of past newsletters, downloadable versions of the newsletter in pdf format, copies of cutting edge research (into asset allocation and investing -- we haven't gone into the "hot tip" business!), and new editions of The Index Investor for our readers whose functional currencies (think of it as the currency in which your future liabilities are denominated) include Australian Dollars, Canadian Dollars, Euros, Japanese Yen, and British Pounds.

To celebrate the launch of our new site, and to thank you for your loyalty over the past four years, we're planning to give each an every one of our existing subscribers two free subscriptions, covering all of 2002 -- one for you, and one to give to a friend or organization (e.g., a library). We'll be offering more details about this next month. In the meantime, in order for us to set up your free and gift subscriptions, we need to get your most recent email address (which will also serve as your username for the new site). To help us collect an up to date list of these, please send an email to us at the following address: newsite@indexinvestor.com. Thank you for your help with this -- we're sure you'll be happy with the result!

Model Portfolio Performance Update

Through the end of October, our model portfolios have underperformed their respective benchmarks. In comparison to last year, when the diversification of our model portfolios into non-U.S. markets generated significant outperformance versus our benchmarks, this year we have seen the opposite effect, as U.S. equity, bond, and currency markets have outperformed most others around the globe.

Our first set of model portfolios are designed to deliver returns that are superior to their respective benchmarks, while taking on the same amount of risk (that is, having the same expected standard deviation of returns). Our first portfolio is benchmarked against a mix of 80% U.S. equities (as measured by the Dow Jones Total Market iShare, IYY) and 20% U.S. bonds (as measured by the Vanguard Total Bond Market Fund, VBMFX). Year to date, this benchmark is down (15.5%), while our model portfolio is down (22.3%), due to the relative underperformance of our allocations to European equities and commodities.

The second portfolio in this group is benchmarked against a mix of 60% U.S. equities and 40% U.S. bonds. Year-to-date, this benchmark portfolio is down (10.4%), while our model portfolio is down (17.2%). The third benchmark portfolio is a mix of 20% U.S. equities and 80% U.S. bonds. Through the end of October, it is down (0.2%), while our model portfolio is down between (3.8%) and (5.4%), depending on the international bond fund used in the portfolio. In this area, the Pimco Foreign Bond fund (which takes a more active management approach to exchange rate exposure) is up 5.5% year to date, and continues to substantially outperform both the Fidelity International Bond and the Price International Bond funds, which, respectively, are down 2.2% and 2.6%.

Our second set of model portfolios are designed to match the returns of their respective benchmarks, while taking on less risk. They have also underperformed. While the 80/20 benchmark is down (15.5%) year-to-date, the model portfolio is down (22.5%). The 60/40 benchmark is down (10.4%) year-to-date, while our model portfolio is down (14.4%). Finally, the 20/80 portfolio has a (0.2%) return through the end of October, while the model portfolio is down between (3.5%) and (5.2%), depending on the international bond fund used.

Our last set of model portfolios are designed differently. They assume that an investor wants to maximize the probability of achieving at least a minimum target level of return, while taking on the least amount of risk possible. Year-to-date, our 12% target return portfolio is down (22.3%), our 10% target return portfolio is down (22.5%), our 8% target return portfolio is down (15.8%), and our 6% target return portfolio is down (9.8%).

Finally, our experimental actively managed portfolio is down (12.5%) year to date, compared to (11.8%) year to date performance by its benchmark, the Fidelity Global Balanced Fund.

In Focus: Downside Risk

The core of the investment challenge is easy to state: How to allocate your savings across different investments in the face of uncertainty so as to maximize the probability of achieving your goals?

From this statement, one can derive a common sense definition of investment risk: it is the probability, and magnitude, by which I might fall short of achieving my goals.

Unfortunately, this definition is not the one that is most often used in the quantitative analysis of investment alternatives. In this case, risk is defined as standard deviation, which measures the dispersion of returns around the average return. (For those readers whose statistics are a bit rusty, recall that if the returns are "normally" distributed around the average (or mean) return, 67% of returns will fall within a range defined as the mean plus or minus one standard deviation, and 95% of returns will fall within plus or minus two standard deviations). Given our common sense definition of risk, you can see how this creates two important problems.

First, the average return may or may not be equal to the minimum rate of return I need to earn on my investments in order to achieve my goals. Second, while standard deviation measures dispersion both above and below the average return, what I'm really worried about are returns that fall below my minimum required rate of return.
Given this, why do are so many quantitative analyses of the basic investment problem stated in terms of means and standard deviations, instead of minimum required rates of return and downside risks?

The simple answer is that when Harry Markowitz published his original articles on portfolio theory in the 1950s, computing power was very inefficient and expensive, and using means and standard deviations vastly reduced the amount of it that was needed to arrive at a optimal solutions to the investment problem (technically, use of optimization algorithms to construct optimal portfolios of different assets, which maximized expected for a given level of expected risk).

In recent years, however, the mean/standard deviation (also known as mean/variance) approach has come under increasing criticism. Apart from the mismatch with common sense which we have already discussed, the other important problem has been the realization that the returns on many asset classes are not, as assumed by the mean/variance approach, normally distributed. More specifically, the return distributions for many asset classes are unequally distributed around the mean (technically, they have negative skewness if they lean left, and positive skewness if they lean right) and/or they have more of their returns at either end of the distribution than is the case for a "normal" distribution (that is, they have "fatter tails", which is technically called "positive kurtosis"). If a distribution is positively skewed, and/or has negative kurtosis, standard deviation will overestimate risk; if it is negatively skewed and/or has positive kurtosis, standard deviation will underestimate the true level of risk.

The following tables (based on monthly data from January, 1988 through December, 2000 demonstrates how this is true for a number of typical asset classes:

Asset Class	Skewness (normal distribution = 0)	Kurtosis (normal distribution = 0)
Lehman Brothers Bond Market Aggregate Index	-.07	-.16
NAREIT Index (REITS)	.31	.70
Goldman Sachs Commodity Index	.82	2.60
Salomon Brothers Non-U.S. Dollar Denominated Government Bonds	.11	.66
U.S. High Yield Bonds	-.14	6.54
Wilshire 5000 U.S. Stocks	-.68	1.59
MSCI EAFE (Non-U.S. Developed Market Stocks)	-.06	.48
MSCI Emerging Markets Index (Stocks)	-.64	2.06

Mounting criticism of the mean/variance approach has led over the past few years to the development of alternative measures and asset allocation approaches. The key building block for these is known as the semi-deviation of returns. This is similar to the standard deviation, except that the only returns taken into account are those that fall below a certain minimum level (which can be either the average return or a minimum required rate of return set by the investor).

The semi-deviation of returns is therefore a measure of downside risk that takes into account both the frequency with which returns fall below a certain level, and the size of the shortfalls involved. As important, because these differences are squared, larger shortfalls are exponentially worse than smaller ones (theoretically, shortfalls could be raised to any power, with higher powers signifying greater risk aversion. However, industry convention seems to have settled on raising them to the power of 2).

Once this downside risk measure is in place, it can then be related to some potential measure of returns. Two approaches are commonly used. The first simply relates the downside risk to the average return. The second (called the Sortino Ratio, after the person who popularized it), relates downside risk to the difference between the average return and the minimum required rate of return. For example, given an eight percent minimum required rate of return, an asset class with an average return of 17.05% and a semi-standard deviation below this target of 26.85% would have a Sortino Ratio of .34 (17.05% - 8% divided by 26.85%) -- in other words, this asset class historically generated .34% in excess return for every unit of shortfall risk taken on.

Let's take a look now at how different asset classes stack up on these measures, including our simplified Sharpe Ratio (which we define as average return divided by the standard deviation of returns). Our data covers the period from January, 1988 to December, 2000.

Asset Class	**Average Return/Standard Deviation**	**Average Return/Downside Risk (using 8% target)**	**Sortino Ratio (using 8% target)**
Lehman Brothers Bond Market Aggregate Index	1.97	.34	.02
NAREIT Index (REITS)	.86	.42	.13
Goldman Sachs Commodity Index	.64	.43	.16

Asset Class	Average Return/Standard Deviation	Average Return/Downside Risk (using 8% target)	Sortino Ratio (using 8% target)
Salomon Brothers Non-U.S$ Denominated Government Bonds	.67	.24	(.05)
U.S. High Yield Bonds	1.19	.34	.04
Wilshire 5000 U.S. Stocks	1.07	.64	.34
MSCI EAFE (Non-U.S. Developed Market Stocks)	.48	.29	.03
MSCI Emerging Markets Index (Stocks)	.57	.48	.23

Clearly, the use of downside risk measures provides a different picture from the mean/variance approach, for two reasons: First, it takes your minimum required rate of return into account. And second, it also takes into account the fact that returns on these asset classes are not normally distributed. Domestic high yield bonds are a good case in point. Using the mean/standard deviation approach, they seem quite attractive. However, when one takes into account the fact that this average reflects a few great periods, rather than lots of good ones, their attraction pales relative to other asset classes which deliver steadier performance, given your 8% target return.

Given the apparent benefits of downside risk measures, why aren't they used more often in practice? Two reasons are most often given. First, using downside risk in a portfolio optimization program is still computationally challenging. Second, downside risk measures are much more subject to what is known as "estimation risk." Let's look at each of these in turn.

In a world where all returns are assumed to be normally distributed, it is relatively easy to use optimization algorithms to identify efficient portfolios (that is, combinations of assets which maximize expected return for a given level of expected risk, or minimize risk for a given level of return). Technically, this is because there are relatively few covariances that need to be taken into account. In a world in which returns are not normally distributed, the number of covariances skyrockets, as does the complexity of the optimization algorithm and the amount of time required to compute it. Obviously, were the results from such an analysis usually substantially different from those produced by a standard mean/variance optimization, it might be well worth the time and effort needed to implement the downside approach. The problem is, in many cases the differences between the recommended asset class weightings produced by the two approaches aren't that large. Given this, many practitioners have been slow to adopt the downside approach.

However, it is the estimation risk problem that is potentially the more important objection to the wider use of the downside risk approach. The essence of this problem is the chance that any estimate that is based on a sample of historical data will fail to capture the full range of possible outcomes (technically, it will "misestimate" the variables). For many asset classes (e.g., emerging markets equity), this is a problem, because not that many years of data are available. In the case of downside risk measures, the problem is compounded because in formulating your estimate, you aren't even using the full range of the historical data you have available -- just those points that fall below some target return. Some studies have found that estimation errors are probably at least twice as large in the case of downside risk measures as they are in the case of standard deviation.

So where does that leave us? Should we simply abandon the use of downside risk measures? Or use them in spite of their potential inaccuracy? Fortunately, there is a third alternative. Recall our initial question: How to allocate your savings across different investments in the face of uncertainty so as to maximize the probability of achieving your goals? There are two approaches to answering this question. Optimization uses sophisticated mathematical techniques (e.g., non-linear programming or Monte Carlo simulation) to find the "best" solution to the asset allocation problem. However, due to estimation errors, "best" solutions that are based on historical data may not turn out to be the best solutions in the future (note that this is true of many things in life, not just asset allocation).

In contrast, the objective of "heuristic" approaches (more commonly known as "rules of thumb") is to quickly come up with a solution that is "good enough", rather than the best one available. In fact, most of the decisions we make every day are based on the heuristic approach, rather than explicit optimization calculations. With respect to asset allocation, a simple example of a heuristic approach would be to (a) choose the asset classes with the four or five highest ratios of average return to downside deviation (given your target return); (b) sum these ratios; and (c) to get your weighting for the asset class, divide its average return/downside deviation by the total you calculated in step (b). In the example above, this would result in an allocation of about 32% of the portfolio to the Wilshire 5000 (U.S. Stocks), 24% to Emerging Market Stocks, 22% to Commodities (GSCI), and 22% to Real Estate (NAREIT). But remember, heuristics aren't fool proof -- through the end of September, this allocation would have delivered year-to-date returns of (16.5%).

When all is said and done, our final take on downside risk is that it is probably best used in the context of a simulation based approach to asset allocation, which helps overcome the estimation error problem. Simulation will be one of the approaches we use in this year's portfolio rebalancing analysis. Beyond this, downside risk analysis also serves as a good reality check (remember our high yield example) that should be applied before you make the final asset allocation decision for your portfolio.

Products and Strategies: "Enhanced" Index Funds

They say that imitation is the sincerest form of flattery. This is undoubtedly the logic that lies behind the appearance in recent years of a strange creature known as the "enhanced index fund". In a nutshell, the sales pitch for the great majority of these funds comes down to something like this: "we're really an index fund, but we deliver slightly better returns than the index." Hmmm…As my three year old says, "sounds suspicious. Let's investigate."

Enhanced index funds basically take one of two approaches in their attempts to deliver "better than the index" returns. For some of them, superior performance basically comes from superior security selection. For example, some of these funds try to tilt their "indexed" portfolios towards those companies in the index which they believe to be undervalued. Others take a different approach, and try to use optimization techniques to identify a portfolio of companies that will match the volatility of their target index, while delivering superior returns.

The second major approach taken by the enhanced index funds is the use of leverage and/or derivatives. For example, an enhanced index fund could spend ten percent of its cash investing in index futures (which enable you to control $100 of the index for $10 of up front cost) while investing the remaining ninety percent in what they believe to be underpriced bonds, in the hope that their profits on the latter lead to above equity index returns. A simpler strategy would be to simply write (that is, sell) call option contracts on the equities held in the index portfolio. As long as these call options never get exercised by their holders (that is, as long as they stay "out of the money"), the profits from the option contract sales generate returns above the target index. A final technique that falls into this category is what is known as cash/futures arbitrage. In this case, the portfolio manager would invest in either futures contracts that are tied to the target index, or in the underlying shares, to take advantage of (that is, to arbitrage) any pricing discrepancies between the two. In this case, it is the resulting arbitrage profits that are expected to deliver enhanced, above index returns.

All this sounds good in theory (though perhaps a little more like active management than many index investors may like); the next logical question is how these funds have faired in practice.

Because it is one of the oldest of the enhanced index funds (and because it comes from one of the leading firms in this area) we've decided to take a close look at the Rydex Nova Fund (RYNVX). The fund's objective is to deliver returns that are 50% higher than its target index, which is the S&P 500. In exchange for this, Rydex requires a large minimum investment ($25,000, though this may be less if you invest via a registered investment advisor), and charges annual expenses of 1.34% -- very substantially more than Vanguard or an iShares ETF that tracks the same index.

Since the fund was launched in 1993, its cumulative return has been 124% of the S&P 500's. During the big downturn in 2000, its fall was 193% of the S&P 500's. Finally, year to date in 2001, its loss is 139% of the S&P 500's.

Let's look at it a slightly different way. Between January, 1995 and December, 2000, the Nova Fund had an average annual return of 24.77%, with a standard deviation of 26.01%, and a downside deviation (assuming an 8% target return) of 29.56%. In other words, if you owned this fund during that period, you received .95% of return per unit of standard deviation risk, and .84% of return per unit of downside deviation risk. During this same period, the S&P 500 had an average annual return of 22.66%, with a standard deviation of 18.03% and a downside deviation (again, using an 8% target return) of 26.30%. In other words, the S&P 500 delivered 1.26% of return per unit of standard deviation risk, and .86% of return per unit of downside risk. Considering the difference in expenses between the Nova Fund and what you would have paid for Vanguard's S&P 500 fund, and you can't escape the conclusion that the latter was the better deal -- especially when you seen how the Nova Fund actually fell by more than its stated target of 150% of the S&P 500 during the big 2000 downturn (technically, this is known as "tracking error").

So we're dead set against enhanced index funds, right? Generally, yes, but with one very important exception. Lumped into the general category of enhanced index funds are funds whose objective is to deliver returns that are the inverse of the returns on the S&P 500. In other words, their objective is to have a correlation with this index that is equal to (1.0). Regular readers of The Index Investor can imagine how our eyes twinkle when we read that. Why? Recall that the risk of a portfolio is a function not only of the riskiness of each individual asset class, but also of the extent to which their returns move together (that is, their correlation). From a portfolio diversification point of view, an asset with a correlation of (1.0) is very attractive, because its returns will tend to be positive while other's are negative. An asset with a negative correlation will also help the portfolio overcome the nasty tendency for asset class correlations to move closer to positive 1.0 when markets move down (as we have seen over the last 18 months). To put these "inverse return funds" into perspective, here are some other asset class return correlations with the S&P 500 over the January, 1995 to December, 2000 period: NAREIT Index, .24; Goldman Sachs Commodity Index, .03; Lehman Brothers Aggregate Bond Index, .22; and Salomon Brothers Non-U.S. Dollar Government Bond Index, .08. If these inverse return funds can deliver on their promise, they have the potential to really add something to your portfolio. So our next task is to see if these funds actually deliver on what they promise.

Of the few inverse return funds that exist, one of the oldest is the Rydex Ursa fund, which was launched in 1994. Until the past 18 months, it has been hard to tell whether or not this fund would be able, in a down market, to deliver on its promise of returns that were opposite those on the S&P 500 (previously, all we had was ample evidence of the fund's ability to do this in a rising market). Well, the jury is now in, and we like what we see. Since its inception, the Ursa fund has delivered cumulative returns of (10.00%) versus 15.95% for the S&P 500. During the year 2000, it delivered returns of 16.44%, compared to the S&P 500's (9.11%). And this year to date, it has 23.02% versus a loss of (7.61%) for the S&P 500. By comparison, a similar fund, the Potomac U.S. Short Fund (PSPSX) has delivered year to date returns of 18.65%.

Our conclusion: in this case, the expenses of 1.37% per year (for the Ursa Fund), are more than offset by the potential diversification benefits this inverse fund delivers. For that reason, and now that they have proven themselves in a big downturn, we are going to consider inverse funds as a separate asset class in next year's rebalancing analysis.