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. Beta testing continues, and our testers have been quite enthusiastic. We hope to make the new site live at the beginning of January. Fingers crossed, of course (as all technology projects inevitably seem to take longer than expected!). Among the many new features on the site 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 and 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). 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. We need this as soon as possible. Thank you for your help with this -- we're sure you'll be happy with the result!

Enron and Active Investment Management

We're sure that like us, many of you have followed the stunning decline in the market value of Enron over the past two months. What most market commentators have missed,

however, is what an example like Enron says about the challenges of active investment management.

Let's start with the underlying story. How did Enron end up declaring bankruptcy this past weekend? Sadly, the underlying plot is a fairly common one: (1) Begin the story with a company that has or develops a set of relatively unique resources/capabilities that generate high returns in a given market; then (2) watch investors drive up the company's stock price through a combination of expectations based on the "best case" outcomes for all relevant pricing variables and "follow the bandwagon" momentum buying. (3) As growth slows in the company's original market, watch its leaders try to stretch the organization's capabilities into new markets to meet high investor expectations for future earnings growth (and, not incidentally, keep their stock options above); (4) Next watch the company's leaders begin to worry as they slowly discover that their capabilities aren't as well suited to these new markets as they had thought, and/or realize that these markets are growing more slowly than expected, with attendant negative consequences for growth in the company's operating cash flows; (5) To remedy this, and keep reported earnings growing and the stock price high, add leverage to the balance sheet; (6) To ensure that analysts keep their "strong buy" on your stock, try to hide some of this additional leverage off balance sheet; (7) Take a look at potential stock based acquisitions, in the hope using your high priced stock to buy assets that will keep earnings growing; (8) Finally, watch the whole thing rapidly unravel when somebody discovers -- and says -- that the emperor isn't wearing any clothes...

Sadly, we've seen the Enron story before, albeit not so big or so fast. But the important point is this: this is easy to see in hindsight. When it came to foresight, even with many, many very smart people (investors, analysts, journalists) looking at Enron's results and public filings, hardly anybody saw this coming (there was one report warning of trouble ahead published by Off Wall Street Research back in May, but its author was generally derided as being naive). In fact, it wasn't until very recently that many investment banks lowered their rating on the company to less than a "strong buy". On the one hand, some will say that this shows how big the potential rewards can be to successful active

managers; on the other hand, many will reply that this also shows just how very, very difficult it is to achieve consistent success (that is, returns above a passive benchmark) by playing the active management game. It is a story we should all keep in mind in the years ahead.

Model Portfolio Performance Update

Through the end of November, 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 substantially outperformed most others around the globe. Two further points may help put this situation in better perspective. First, by many traditional measures (e.g., Price/Earnings multiples), the U.S. equity market is now back to the high end of its historical valuation range, in contrast to other equity markets which appear to be much more reasonably valued. This valuation discrepancy suggests (based on history), that non-U.S. markets may perform relatively better than the U.S. in the years ahead. Given this, holding a diversified portfolio should lead to higher long term returns. Second, with respect to our benchmark portfolios, the returns we have seen thus far this year are for the most part about two standard deviations below their expected long term averages (based on 1/88 - 12/00 data), which makes them quite rare events, statistically speaking. On the other hand, the year-to-date returns on many of our model portfolios are currently closer to three standard deviations away from their expected long term average returns -- in other words, while the benchmark portfolios have been having very bad years, our model portfolios have been having truly terrible years in statistical terms. Surprisingly, this is quite good news, as it suggests that their current level of underperformance versus their respective benchmarks is truly exceptional.

Having said that, let's move on to the results through November.

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). Thus far this year, they have underperformed their benchmarks. 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 (10.7%), while our model portfolio is down (18.6%).

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 (7.3%), while our model portfolio is down (14.9%). The third benchmark portfolio is a mix of 20% U.S. equities and 80% U.S. bonds. Through the end of November, it is down (0.5%), while our model portfolio is down between (4.4%) and (6.2%), depending on the international bond fund used in the portfolio. In that regard, the Pimco Foreign Bond fund (which takes a more active management approach to exchange rate exposure) continues to outperform both the Fidelity International Bond and the Price International Bond funds.

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 (10.7%) year-to-date, the model portfolio is down (19.2%). The 60/40 benchmark is down (7.3%) year-to-date, while our model portfolio is down (12.6%). Finally, the 20/80 portfolio has a (0.5%) return through the end of November, while the model portfolio is down between (4.0%) and (5.9%), 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 (18.6%), our 10% target return portfolio is down (19.2%), our 8% target return portfolio is down (13.7%), and our 6% target return portfolio is down (9.5%).

Finally, our experimental actively managed portfolio is down (15.8%) year-to-date, versus (9.1%) year-to-date by its benchmark, the Fidelity Global Balanced Fund. To some extent, our active management performance has been hampered by the limitation we placed on our trading, which allowed us to switch portfolio asset allocations only four times, at the end of December, March, June, and September. That being said, we are still lagging well behind our benchmark, as rallies we had expected in foreign equities, and then U.S. equities failed to materialized when we expected them (however, just to rub salt in our wounds, the latter now seems to be staging a bit of a rally, while we are fully invested in U.S. bonds!). As we've said all along, active management is a very, very tough game to play.

Model Portfolio Rebalancing Issues

One of the most important questions in investments is easy to pose: "How should I divide my funds between different asset classes?" Answering it is a complex and challenging undertaking, because of the amount of data, uncertainty, and calculations that are potentially involved. Broadly speaking, there are three different approaches that one can take.

The first approach is to use a rule of thumb (more formally known as a "heuristic"). One rule of thumb that research suggests is commonly used by many retirement plan participants is to divide one's funds equally between all the different kinds of investment on offer. Another example would be to divide one's funds 60/40 between a broad domestic equity index and a broad domestic bond index. The advantage of using a rule of thumb to allocate one's investment funds is that this approach greatly cuts down on the amount of complexity one has to deal with. The disadvantage is that in so doing, it prevents you from making the fullest use of the information that is available. Lacking this information, you may not be able to identify allocations with the potential to deliver higher returns and/or less risk than the portfolios produced by the rule of thumb.

The second approach is to use some type of computer program to analyze the available information in order to develop a recommended asset allocation. Until recently, the most popular type of program used for this purpose employed a technique known as "Mean-Variance Optimization". MVO requires three inputs for each asset class under consideration: the expected average (or mean) annual return for the asset class, the expected standard deviation of those returns (that is, their expected dispersion around the average), and the correlation of these returns with those from the other asset classes under consideration. The advantage of the MVO approach is that it makes fuller use of the information available, and in so doing holds out the prospect of generating asset allocation recommendations that are superior to those produced by rules of thumb. However, it also suffers from some important limitations.

First, MVO solutions can be very "unstable", meaning that a slight change in one variable -- most often, the expected average annual return for a given asset class -- will result in a large change in the recommended portfolio weights for each asset class. In practice, there are a number of steps that experienced MVO users take to reduce this problem, including using broadly defined asset classes with low correlations of returns (as high correlations exacerbate the impact of slight changes in expected returns), using risk minimization as the objective of the optimization, instead of return maximization (because relative asset riskiness is more stable over time that relative asset class returns), and setting "common sense" limits on the maximum portfolio weight for any asset class. At The Index Investor, we employ all three of these techniques, and find that they substantially reduce, but do not completely eliminate the instability problem.

Second, as we discussed in last month's issue, the MVO approach uses total volatility around the average expected return as its proxy for risk. In reality, most investors are far more concerned about losses on the downside than they are about foregone gains on the upside; in other words, for most people risk means "risk of loss."

Third, the MVO approach finds it very difficult to take multiple investor objectives into account. Fourth, the MVO approach is based on a single period of time, and encounters

significant problems when it is extended to multiple periods. Consider a simple example. An asset returns 100% in year 1, and then (50%) in year 2. Its average annual return is 25%. However, ending wealth is exactly what you started with -- that is, the compound annual rate of return (also known as the geometric mean) is equal to zero. So what to do? There are two approaches that are often used to try to get around this problem. The first uses as inputs average return, standard deviation and correlation over the target time period (e.g., ten years). The problem is that given the length of the data available for forming these estimates (at best, it usually goes back only to 1926), they end up being based on less data than single period estimates, and are therefore potentially less accurate. The second approach substitutes the geometric return for the average return; however, the theoretical validity of doing this has yet to be proven.

Lastly, the MVO approach suffers from three limitations that affect all quantitative models: (a) The estimates used for asset class average expected returns and standard deviations may not be accurate (that is, the "true" distribution may diverge from the one based on the sample of returns that is used to make the estimate). (b) The relationships between the variables specified in the model (in this case, the correlations between their respective returns) may not be an accurate reflection of reality. (c) Last but not least, even if the estimates used are very close to the "true" conditions that existed in the past, they may be very inaccurate descriptions of future conditions if some aspect of the underlying system (e.g., the structure of the economy) undergoes a permanent and fundamental change. In short, any model is at best an approximation of the "true" workings of some underlying system, and, for many reasons, the accuracy of this approximation (and the conclusions that are based on it) may fall well short of perfect foresight into the future.

This brings us to the third major approach to asset allocation, which is called dynamic programming, or DP. The goal of this approach is to develop a multi-period asset allocation solution that is robust enough to satisfy an investor's objectives under a range of future scenarios. In practice, it works like this: First, you define your investment objectives. In our case, we have used two: (a) given historical data, ensure that the

investor's minimum required rate of return will be met (on a compound annual rate of return basis) over a twenty year time horizon with a minimum level of probability (e.g., "I want to be at least 90% sure my portfolio will earn a compound annual rate of return of at least 8% over the twenty years I've got left before I retire"); and (b) achieve this while taking on as little risk as possible ("…and I don't want to ride a roller coaster on the way there!"). The next step is to define the average annual return, standard deviation, and correlation assumptions for the asset classes under consideration. The third step is to define the maximum amounts that can be invested in any single asset class. So far, this sounds just like the MVO approach doesn't it? The difference lies in the way the DP approach uses these inputs. Rather than using them to construct a single period optimization solution (for the technically minded, to calculate an efficient frontier -- that is, the set of different portfolios that maximizes return for any given level of risk), DP uses these inputs as the basis for a complex Monte Carlo simulation analysis.

Here's how it works. DP starts with a "candidate" asset allocation solution (e.g., 30% domestic bonds, 40% domestic equities, and 30% international equities). It then uses the asset class inputs to calculate a scenario covering the twenty year holding period. For example, if you had seven asset classes, and a holding period of twenty hears, the scenario would contain 140 different annual returns. Once the scenario has been defined, the DP approach checks to see how well the candidate asset allocation satisfies the investor's objective. It then repeats this process many times (we use from 2,000 to 5,000 repetitions) to develop a clear picture of the range of outcomes the candidate asset allocation might produce. But it doesn't stop here. The DP program uses a search process (called "Tabu Search", that is beyond the scope of this article to describe) to identify changes in the asset allocation solution that might do a better job of achieving the investor's goals. For each alternative asset allocation solution, from 2,000 to 5,000 scenarios are generated. The goal is to identify the most robust asset allocation solution -- that is, the one that is most likely to achieve the investor's goals under a wide range of possible future scenarios.

The advantages of the DP approach are that it handles time, uncertainty, and multiple investor objectives much better than either the heuristic or MVO approaches to asset allocation. The disadvantages are that it is very computationally demanding (that is, it requires a powerful computer and a lot of time), and it is also subject to the same risks of estimation error and structural changes in the underlying economy that apply to the MVO approach.

Given this background to the different approaches, let's move on to how we use all three of them here at The Index Investor.

We use the heuristic approach to construct the benchmarks against which we track the performance of our model portfolios. Specifically, we use six different heuristic benchmarks comprised of 80/20, 60/40, and 20/80 mixes of broad domestic equity and bond indexes, and the same mixes of broad global equity and bond indexes.

We also include two sets of portfolios based on the MVO approach. The first assumes that the investor's objective in any given year is to achieve a higher rate of return than the domestic benchmark portfolio, while taking on no more risk. The second assumes that the investor wants to match the benchmark's return while taking on less risk.

Finally, this year for the first time we are utilizing the DP approach to develop our target return portfolios. These portfolios assume that an investor wants to have a certain level of confidence (e.g., 90%) that his or her portfolio will earn a compound annual rate of return of at least a certain minimum amount (in our case, either 6%, 8%, 10%, or 12%) over a twenty year holding period, while taking on as little risk as possible.

Generally speaking, all our portfolios are constructed with the same limits on the maximum amount that may be invested in any asset class. In general, no asset class may account for more than 60% of the MVO or DP portfolios, except for emerging markets equity and commodities which are each limited to no more than 10% of the total portfolio. Finally, while we have not included them in our analysis (because they violate

our correlation limits), one may be able to achieve better returns by taking different tilts within the broad asset classes we have used. We have covered this topic extensively over the past year. For example, based on history, investing in "value" equity indexes should produce higher compound returns over long periods of time than investing in broad equity market indexes; however, as recent experience has clearly shown, this may very well not be the case from year to year.

Model Portfolios for 2002

Our first set of model portfolios is based on an investor who seeks to maximize returns, while taking on the same amount of risk as the benchmark portfolio:

Benchmarks:	**80% Domestic Equities and 20% Domestic Debt.**	**60% Domestic Equities and 40% Domestic Debt**	**20% Domestic Equities and 80% Domestic Debt**
Domestic Investment Grade Bonds	0%	12%	55%
High Yield Bonds	0%	5%	3%
TIPS	0%	0%	10%
Non-U.S. Bonds	0%	5%	0%
Commodities	10%	10%	5%
U.S. Equities	55%	47%	16%
REITs	3%	6%	5%
European Equities	25%	10%	6%
Pacific Equities	0%	0%	0%
Emerging Market Equities	7%	5%	0%
Total	100%	100%	100%

Our second set of model portfolios is based on an investor who seeks to minimize risk, while matching the returns of the benchmark portfolio:

Benchmarks:	**80% Domestic Equities and 20% Domestic Debt.**	**60% Domestic Equities and 40% Domestic Debt**	**20% Domestic Equities and 80% Domestic Debt**
Domestic Investment Grade Bonds	5%	29%	40%
High Yield Bonds	0%	5%	8%
TIPS	0%	0%	25%
Non-U.S. Bonds	0%	0%	0%
Commodities	10%	5%	5%
U.S. Equities	58%	45%	10%
REITs	10%	6%	4%
European Equities	17%	10%	8%
Pacific Equities	0%	0%	0%
Emerging Market Equities	0%	0%	0%
Total	100%	100%	100%

Our third set of model portfolios is based on an investor who seeks to achieve a minimum compound rate of return over twenty years with a certain level of probability, while taking on as little risk as possible.

Goal:	**99% probability of achieving 6% CR**	**95% probability of achieving 8% CR**	**90% probability of achieving 10% CR**	**85% probability of achieving 12% CR**
Domestic Investment Grade Bonds	40%	0%	0%	0%
High Yield Bonds	5%	10%	2%	0%
TIPS	35%	28%	25%	0%
Non-U.S. Bonds	0%	15%	0%	0%
Commodities	5%	8%	6%	10%
U.S. Equities	5%	7%	39%	50%
REITs	5%	6%	2%	10%
European Equities	0%	20%	18%	25%
Pacific Equities	0%	3%	0%	0%
Emerging Market Equities	5%	3%	8%	5%
Total	100%	100%	100%	100%

Finally, we should note that it may be possible (though as we have written over the past year, it is certainly not guaranteed) to further enhance returns within some asset classes by employing different tilts. Examples of this might include a tilt toward value within any broad equity asset class, or toward intermediate maturities within a broad fixed income asset class.