The Index Investor

Why Pay More for Less?

Global Asset Class Returns

Year to Date	**In USD**	**In AUD**	**In CAD**	**In EURO**	**In JPY**	**In GBP**
US Equity	16.70%	1.38%	3.21%	12.03%	15.07%	18.79%
US Bonds	1.00%	(14.32%)	(12.49%)	(3.67%)	(0.63%)	3.09%
AUS Equity	25.00%	9.68%	11.51%	20.33%	23.37%	27.09%
AUS Bonds	11.33%	(3.99%)	(2.16%)	6.66%	9.70%	13.42%
CAN Equity	30.50%	15.18%	17.01%	25.83%	28.87%	32.59%
CAN Bonds	17.47%	2.15%	3.98%	12.80%	15.84%	19.56%
Euroland Equity	15.40%	0.08%	1.91%	10.73%	13.77%	17.49%
Euroland Bonds	7.29%	(8.03%)	(6.20%)	2.62%	5.66%	9.38%
Japan Equity	19.40%	4.08%	5.91%	14.73%	17.77%	21.49%
Japan Bonds	(0.14%)	(15.46%)	(13.63%)	(4.81%)	(1.77%)	1.95%
UK Equity	7.10%	(8.22%)	(6.39%)	2.43%	5.47%	9.19%
UK Bonds	(0.57%)	(15.89%)	(14.06%)	(5.24%)	(2.20%)	1.52%
World Equity	16.60%	1.28%	3.11%	11.93%	14.97%	18.69%
World Bonds	3.10%	(12.22%)	(10.39%)	(1.57%)	1.47%	5.19%
Commodities	14.40%	(0.92%)	0.91%	9.73%	12.77%	16.49%
XR Chng v. USD	0.00%	15.32%	13.49%	4.67%	1.63%	-2.09%

Model Portfolio Update

The objective of our first set of model portfolios is to deliver higher returns than their respective benchmarks, while taking on no more risk. The benchmark for the first portfolio in this group is an aggressive mix of 80% domestic equities, and 20% domestic bonds. Through the end of August, this benchmark had returned 13.6%, while our model portfolio had returned 16.7%. We have also compared our model portfolios to a set of global benchmarks. In this case, the global benchmark is a mix of 80% global equities, and 20% global bonds. Through the end of last month, it had returned 13.9%.

The benchmark for the second portfolio in this group is a mix of 60% domestic equities and 40% domestic bonds. Through the end of last month, it had returned 10.4%, while our model portfolio had returned 14.2%, and the global benchmark had returned 11.2%.

The benchmark for the third portfolio in this group is a conservative mix of 20% domestic equities and 80% domestic bonds. Through the end of last month, it had returned 4.1%, while our model portfolio had returned 6.3% and the global benchmark 5.8%.

The objective of our second set of model portfolios is to deliver less risk than their respective benchmarks, while delivering at least as much return. The benchmark for the first portfolio in this group is an aggressive mix of 80% domestic equities, and 20% domestic bonds. Through the end of last month, this benchmark had returned 13.6%, while our model portfolio had returned 15.4%. We have also compared our model portfolios to a set of global benchmarks. In this case, the global benchmark is a mix of 80% global equities, and 20% global bonds. Through the end of last month, it had returned 13.9%.

The benchmark for the second portfolio in this group is a mix of 60% domestic equities and 40% domestic bonds. Through the end of last month, it had returned 10.4%, while our model portfolio had returned 11.5%, and the global benchmark had returned 11.2%.

The benchmark for the third portfolio in this group is a conservative mix of 20% domestic equities and 80% domestic bonds. Through the end of last month, it had returned 4.1%, while our model portfolio had returned 6.2% and the global benchmark 5.8%.

The objective of our third set of model portfolios is not to outperform a benchmark index, but rather to deliver a minimum level of compound annual nominal return over a ten-year period.

Through last month, our 12% target return portfolio has returned 16.6% year-to-date, our 10% target return portfolio has returned 13.6% our 8% target return portfolio has returned 10.3%, and our 6% target return portfolio has returned 6.1%.

Last month, the active portfolio was allocated as follows: 60% to the Vanguard Inflation Protected Securities Fund, 15% each to the Oppenheimer Real Assets Fund and the T. Rowe Price International Bond Fund, and 10% to the U.K. Equity Market iShare. These will not change next month. Year-to-date, our actively managed portfolio has returned 7.6%.

Equity Market Valuation Update

Our valuation analysis rests on two fundamental assumptions. The first is that the long term real equity risk premium is 4.0% per year. The second is the rate of productivity growth the economy will achieve. As described in our June, 2003 issue, we use both high and a low productivity growth assumptions. Given these assumptions, here is our updated market valuation analysis at the end of last month:

Country	Real Risk Free Rate	Equity Risk Premium	Required Real Return on Equities	Expected Real Growth Rate*	Div Yield	Expected Real Equity Return**
Australia	3.25%	4.00%	7.25%	4.90%	3.80%	8.70%
Canada	3.09%	4.00%	7.09%	2.10%	1.90%	4.00%
Eurozone	1.92%	4.00%	5.92%	2.50%	2.60%	5.10%
Japan	1.99%	4.00%	5.99%	2.80%	1.00%	3.80%
U.K.	2.08%	4.00%	6.08%	2.50%	3.40%	5.90%
U.S.A.	2.78%	4.00%	6.78%	4.50%	1.70%	6.20%

**High Productivity Growth Scenario. See Asset Class Review, in our June 2003 Issue.*

*** Expected real growth rate plus current dividend yield*

Country	Implied Index Value*	Current Index Value	Current/Implied (high productivity growth)	Current/Implied (low productivity growth)
Australia	397.40	245.76	62%	88%
Canada	99.66	261.73	263%	315%
Eurozone	99.11	130.37	132%	189%
Japan	27.93	89.10	319%	419%
U.K.	255.11	268.62	105%	149%
U.S.A.	305.31	409.47	134%	193%

* *High productivity growth scenario.*

Product and Strategy Notes

We begin this month with a number of interesting product and strategy notes.

Vanguard has announced the launch of a new set of target retirement index funds. These will be available in the fourth quarter of this year, and will be comprised of different mixes of Vanguard equity and bond index funds, based on an investor's expected retirement age. These funds will include the Total (U.S.) Stock Market Fund, the European and Pacific Stock Index Funds, the Total (U.S.) Bond Market Fund, and the Inflation Protected Securities Fund. For example, the Target Retirement 2035 Fund will allocate 80% of its assets to equities, and 20% to bonds, while the Target Retirement 2005 Fund will allocate 35% to equities and 65% to bonds. Over time, as an investor draws closer to his or her retirement, these funds will gradually adjust their asset allocations, lowering the amount invested in equity, and raising the portion in bonds. Our initial reaction to these products is mixed. On the positive side, their expected expense ratios (of only .21% to .23% of net assets) is appealing, as is (for some investors) the "one stop shopping" approach. On the other hand, we wish they had included a broader range of funds (e.g., including non-U.S. bonds, real estate, and commodities). We plan on taking a more detailed look at these products in a future issue.

Another interesting product development this month was progress toward the launch of more gold-based index fund products. The World Gold Council (a trade association made up of many of the world's largest gold mining companies) is reportedly planning to list Gold Bullion plc in the U.K. later this year, as well as a gold-backed exchange traded fund (tentatively called Equity Gold Trust) in the United States, assuming some rather sticky intellectual property ownership and investor taxation issues can be worked out. The proposed U.K. and U.S. vehicles will make it much easier for investors to gain exposure to gold without having to purchase and store the physical commodity. Similar vehicles are already available in Canada (Central Gold Trust) and Australia (Gold Bullion Ltd.). Once they become available, we plan to take a closer look at them, and to more fully examine the pros and cons of gold as an asset class separate and apart from commodities.

On the fixed income front, Barclays Global Investors has received preliminary Securities and Exchange Commission approval for two new Exchange Traded Funds. One will track the Lehman Brothers Aggregate Bond Market Index, while the other will track the Lehman Brothers Inflation Protected Securities Index (i.e., U.S. real return bonds). We very much look forward to these two new funds becoming available.

With respect to the indexes themselves, the FTSE Group has announced that as of September 22nd, it will expand the coverage of its All-World Index, which will now include 90% of the equity market capitalization in its constituent countries. The top 70% of companies by market capitalization will comprise the large cap index, while the next 20% will comprise the mid-cap index. The FTSE Group also plans to launch a global small capitalization index product, which will track the performance of those companies in the top 8% of the last 10% of market capitalization in each country. This moves the FTSE methodology much closer to the one used by Dow Jones, which we have long favored.

On the economic front, the International Monetary Fund has released a very comprehensive new report titled: "Deflation: Determinants, Risks and Policy Options." It is available from www.imf.org.

Finally, Philip Coggan had a particularly good column recently in the Financial Times. Commenting on a Watson Wyatt (an actuarial firm) report ("How Have Older Workers Responded to Scary Markets?") on the impact of the recent equity market decline on the timing of retirement, Coggan noted that there was a substantial difference between workers covered by defined benefit plans and those covered by defined contribution plans. Because DB plan sponsors take the risk of poor investment performance (which force them to increase their annual contributions to the plan), the equity market downturn had little or not impact on the workers covered by these schemes. This was not the case in DC plans (e.g., a 401k in the U.S.), where the employee bears the risk of poor investment performance. Many workers covered by these plans have been forced to delay their planned retirement dates because of poor investment performance. As Coggan notes, with DC plans growing more popular all the time (at least with employers), "this is a glimpse of the future…Although the government is keen to transfer the burden of pensions from taxation to the private sector, it has made very little effort to provide workers with the education they need to cope with this responsibility…The danger [is that] people who are saving for their pensions now…are unaware of the risks involved, and unaware of the returns or the level of savings needed to generate a decent retirement income…If the pensions crisis is to be solved in the long run…workers will simply have to stop regarding pensions as 'boring' and become much better informed."

The Confusing World of Factor Models

As economists use the term, a "factor" is a variable that can help explain the returns on different securities. "Factor models" are very popular, and are used to both explain past returns and to forecast future returns. In a typical factor model, the return on a security is described as a function of (1) its weight (also called its "loading" or "beta") on different factors, times (2) the return premium associated with those factors, plus (3) residual return not associated with any factor. Depending on the model, factor loadings and their associated return premia may be stable or vary over time. Thus the return on a given security can be decomposed into "factor-based" return (also known as systematic or market related return) and "company-specific" return (also known as non-systematic or non-market return, or "alpha"). As risk is the other side of return, these are also known as market risk and company specific risk. For example, a "market neutral" hedge fund would use long and short positions, and/or derivative instruments to eliminate its exposure to common market risk factors, and expect to derive all of its return from its exposure to company specific risk.

As this example, shows, factor models allow the return earned by a portfolio manager to be broken down into three elements: (a) market timing, that is, the decision to hold cash or invest in securities; (2) return related to bearing different levels of market (factor-based) risk; and (3) return related to bearing company specific risk.

Different schools of thought interpret the existence of factors that can help explain or predict returns in different ways. One school sees the market as basically efficient, and interprets factors as different sources of risk, which investors are rewarded for bearing. In contrast, those who don't believe in the basic efficiency of financial markets see the existence of factors as evidence of investor bias and misvaluation.

Before taking a stand on this controversy, we first need to develop a better understanding of the confusing world of factor models.

Let's start with two of the most famous factor models. The first is known as the Capital Asset Pricing Model, or simply CAPM ("cap-m"). It is a one factor model, and describes market risk as being a function of the extent to which the returns on a given security co-vary with the returns on the market as a whole. The CAPM is the basis for most single period mean/variance optimization models. Over the years, CAPM has been subject to a lot of criticism, as well as spirited defenses.

Among the leading critics of CAPM have been professors Eugene Fama and Kenneth French, who developed a three factor model which statistically does a better job of explaining historical equity returns. In addition to the market risk factor used in the CAPM, the Fama

French model uses a factor based on the difference in returns between stocks with high book to market ratios and low book to market ratios (often referred to as the “value” factor), and another factor based on the difference in returns between companies with small market values and big market values (often referred to as the “size” factor). These three "Fama French factors" are commonly described as MKT, HML (high minus low), and SMB (small minus big). Fama and French theorized that both HML and SMB are proxies for a single so-called "distress" factor -- that is, they believe that because small firms and high book/market firms have greater risks of encountering financial distress, investors demand (and receive) higher returns for holding their shares. From a portfolio optimization perspective, Fama and French's model implies that investors must optimize expected return relative to not just one factor, but to three of them (needless to say, this makes the calculations involved rather complex). However, the Fama French model also has been criticized on a number of grounds.

Before describing the criticisms, however, we should detour briefly to clarify a point that has caused no end of confusion since Fama and French published their three factor model in 1992 (in their article titled "The Cross Section of Expected Stock Returns"). HML and SMB are not equivalent to portfolios of "value" and "small cap" stocks. Rather, they are constructed by substracting the return on the top ten or twenty percent (depending on who is doing the analysis) of stocks with the lowest book to market ratio from the return on a similar group of stocks with the highest book to market ratio. This same "long/short" approach is used to construct the SMB factor. It is therefore difficult, if not impossible for investors to simply replicate these factors in their portfolios, because they assume the maintenance of constant short positions in low book/market (e.g., growth or glamour) and high capitalization stocks. In practice, the closest you can come is to invest in value or small cap index funds, accepting that these only replicate the "long" position in the Fama French factors, but not the factors themselves.

As we said, the Fama French Three Factor Model has been challenged on a number of grounds. First, the economic meaning of the HML and SMB factors has been widely questioned. In their paper "Evidence on the Characteristics of Cross Sectional Variation in Stock Returns", Daniel and Titman (note that all of the authors we refer to are respected academics, and their papers can usually be found at either google.com or ssrn.com) basically accuse Fama and French of confusing correlation with causation. Their analysis finds that rather than having a high loading on a unique "distress factor", firms with high book/market ratios and/or small capitalizations have a lot of characteristics in common (e.g., membership in the same industry sector) that better explain the extent to which their returns tend to move together. More specifically, they find that the returns on firms with common characteristics, but different loadings on the HML and SMB factors still move together (in contrast to what the Fama French model would predict). They note that "firms which become distressed should logically load on the same factor or factors, and become distressed when these factors turn negative. They don't have negative results because they load on a unique distress factor."

Along a similar line, in their paper "Commonality in the Determinants of Expected Stock Returns", Haugen and Baker find that the highest return deciles of stocks are not relatively riskier based on the Fama French factors, and that the differences in realized returns between different stocks are actually based on other factors, including differing growth potential, relative pricing (e.g., comparable price/cash flow ratios), liquidity, and momentum. The existence of the latter two factors have both been credibly supported in other papers, including "Liquidity Risk and Expected Stock Return" by Pastor and Stambaugh and "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency" by Jegadeesh and Titman.

Another criticism of the Fama French model has been made by Berk, Green and Naik in their paper "Optimal Investment, Growth Options, and Security Returns". They note that over time, companies typically encounter a stream of potential projects in which they can invest. Moreover, these projects have differing degrees of risk -- some are higher than the current average level of risk for the firm as a whole, and some are lower. When interest rates (and therefore the firm's cost of capital) are high, the present value of the cash flows these projects are expected to produce will be quite low (as is the present value of future cash flows from the firm's existing assets). The firm will therefore find its market value declining relative to its book value. However, while relatively few projects will be undertaken when rates are high, the ones that are will have lower than average risk. As a result, future returns on the market value of the firm's equity will be quite high (because the initial market value is lower, and the projects undertaken less risky). The opposite will happen when interest rates fall. The market value of future cash flows from current assets will rise. More projects will have positive net present values, and be undertaken (increasing the size of the firm), including a greater proportion of risky projects. The subsequent return on the market value of the firm's equity will be lower (because the initial value is high and the projects undertaken are, on average, riskier). In short, rather than being proxies for a unique "distress factor", the HML and SMB effects identified by Fama and French reflect the normal workings of a dynamic economy.

In another paper "An Empirical Re-Examination of the Relation Between Firm Size and Return", Berk also finds that other factors that measure the size of a company (e.g., its book value, sales revenue, or number of employees) have no relationship with return. This leads him to conclude, again, that the Fama French SMB factor "merely reflects the inverse relationship between market value and a firm's cost of capital." Berk, Green, and Naik's findings are in line with those of other studies which have analyzed the dynamic behavior of equity returns over time (e.g., see "Risk and Return: Some New Evidence" by Guo and Whitelaw). A key assumption of this group is that, to some extent, investors purchase equities to hedge their consumption over time. Given this, assets whose value tends to increase when consumption (or labor income) declines will be more attractive (and have lower required returns) than those assets whose values track changes in consumption. Moreover, these hedging related return premia will tend to vary over time in line with the state of the business cycle. The total rate of return on an asset will therefore be a function not only of its risk (that is, the extent to which its returns covary with the market porfolio) but also of its attractiveness as a hedging vehicle. Interestingly, the asset classes with the lowest correlation of returns with the change in real personal consumption expenditure in the U.S. are real return, domestic investment grade, and foreign bonds bonds, and emerging market equities.

In defense of the Fama French model, professor Maria Vassalou of Columbia University has shown that HML and SMB are good proxies for news about future economic growth, and that the premia associated with them are only present in the group of stocks with the highest credit default risk. In light of this, she also concludes that the return premia on HML and SMB are compensation for bearing business cycle risk.

Given the controversy over the use of market based factors like the Fama French model, one would expect that people would look to fundamental economic variables to explain and predict asset returns. A good overview of efforts in this area is contained in "Asset Returns and Economic Risk" by Cesare Robotti of the Federal Reserve Bank of Atlanta. He notes that beginning with a paper by Chen, Roll, and Ross ("Economic Factors and the Stock Market"), "the view that the view that macroeconomic risks systematically affecting asset returns should also be significantly priced has attracted widespread attention and generated a large body of empirical work…Chen, Roll, and Ross suggested that in selecting factors one should consider forces that will explain changes in the rate used to discount future expected cash flows, and influence these cash flows themselves."

In other words, investors should be willing to pay higher prices, and accept lower returns, for those assets that best hedge against macroeconomic risks. Unfortunately, Robotti concludes that "logical though they may be, the evidence supporting macroeconomic factors has been weaker than one would have expected", though the real rate of interest, the maturity premium (the difference between short and long term interest rates) and the default premium (the difference between rates on government and corporate bonds) have shown much promise in different factor models.

Rather than using market or economic factors, another group of academics have recently taken a third approach to the construction of factor models. Their starting point is a key assumption that underlies the CAPM: the returns on the assets in which one can invest are normally distributed (that is, when plotted, they look like the familiar "bell curve"). If this is true, then the only risk an investor should care about is the extent to which returns on different assets covary with each other. In reality, however, the returns on most assets are not normally distributed. In fact, the distribution of most assets' returns tends to be more skewed than the normal distribution (that is, more tilted to the left or right of the average), and to have "fatter tails" (that is, to have a greater proportion of returns located farther away from the average than would be the case in a normal distribution). This means that, in addition to the covariance of returns (e.g., the extent to which the returns on a given asset vary with the returns on the market as a whole), an investor should also be concerned with the coskewness of returns (the extent to which the returns on a given asset tilt in the same direction as the market as a whole) and their cokurtosis (the extent to which an extreme return on a given asset will occur at the same time as the market as a whole has an extreme return). This insight has given rise to the "four moment" capital asset pricing model, which is well described in two papers: "CAPM and Higher Comoment and Factor Models of UK Stock Returns" by Hung, Shackleton, and Xu, and "Portfolio Selection With Higher Moments" by Harvey, Liechty, Liechty, and Muller.

In sum, anybody trying to either explain historical returns, or forecast future returns is faced with a lot of uncertainty about which factor model to use. Practically, this has important implications. It makes it much more difficult to evaluate active investment managers' performance, and it also makes asset allocation more difficult, as we will see in our next article.

Asset Allocation Review: Dealing with Uncertainty

We began our asset allocation analysis back in May, with a review of historical and expected future real returns across various asset classes and currencies. Last month, we used our future risk and return estimates as inputs into a simulation optimization process, to produce portfolios with the highest probability of achieving compound target real returns of 3%, 5%, and 7% over a twenty year period.

This month, we will begin by using the historical data for each asset class as inputs into the same simulation optimization process. We will then discuss the limitations of our approach, and various proposals for how they might be overcome. We will conclude this month with our final recommended twenty year target return portfolios. Next month, we will conclude our asset allocation analysis with a discussion of how our recommendations change when the investment horizon is shortened to ten years.

Portfolios Based on Historical Data

The inputs for this month's target return portfolios were estimated from 32 years of monthly real returns data, covering the period from 1971 to 2002. They are summarized in the following table:

	Avg Return	*Std Dev*	Real Bonds	Dom Bonds	For Bond	Comm Prop	Commodities	Dom Equity	For Equity	Emrg Equity
Real Bonds	2.30%	2.50%	1.00	0.27	0.02	-0.02	0.15	-0.15	-0.13	-0.14
Dom Bonds	3.80%	5.40%		1.00	0.20	0.18	-0.06	0.16	0.06	-0.07
For Bond	9.50%	11.20%			1.00	-0.01	0.03	0.06	0.45	0.00
Comm Prop	7.90%	9.80%				1.00	-0.05	0.41	0.24	0.28
Commodities	8.10%	18.30%					1.00	-0.05	0.00	0.02
Dom Equity	7.30%	16.30%						1.00	0.61	0.62
For Equity	7.00%	17.20%							1.00	0.55
Emrg Equity	9.60%	24.00%								1.00

We used these historical inputs to generate the model target return portfolios using our simulation optimization process. As we did last month (when we generated target return portfolios using our future asset class rate of return forecasts), we set limits on the maximum amount that could be allocated to certain asset classes, including 40% to foreign bonds, and 20% each to commercial property, commodities, and emerging markets equity. The model portfolios are as follows:

U.S. Dollar	**3% Target**	**5% Target**	**7% Target**
Real Return Bonds	45%	0%	5%
Domestic Bonds	25%	20%	5%
Foreign Bonds	10%	25%	40%
Commercial Property	10%	20%	15%
Commodities	5%	0%	10%
Domestic Equity	0%	20%	5%
Foreign Equity	5%	5%	0%
EM Equity	0%	10%	20%
Total	100%	100%	100%
Probability of Achieving Target Return	98.0%	91.0%	74.0%
Expected Average Annual Return	4.5%	7.5%	8.4%
Standard Deviation of Expected Returns	2.6%	5.5%	7.0%

A couple of items in this table might need a bit of further explanation. The "Probability of Achieving Target Return" refers to the target compound annual rate of return at the end of the twenty year holding period. The "Expected Average Annual Return" is the expected rate of return for the portfolio in any single year, rather than over the full twenty years. Generally speaking, the average annual (also known as arithmetic) return will be higher than the average compound (also known as geometric) return, because the latter reflects the full impact of poor annual returns on the achievement of your long term financial goal. A simple example can help make this clear. Consider an investment which, over three years, has annual returns of 20%, (30%), and 20%. If you invested 100 at the beginning of the period, at the end of three years you would only have 100.8 in your account, for a compound return of only 0.27% over the three year period. Your average annual return, however, would be much higher: 3.33%!

So while the Expected Average Annual Return and Standard Deviation give you an indication of the performance a model portfolio is likely to deliver in any given year, what really counts is the portfolio's probability of achieving your long term compound annual return target.

Another interesting point about these portfolios is the relatively heavy weights they give to bonds, property, and commodities, and the relatively low weights (at least compared to what you often see in the media) they give to equities. The logic behind these allocations is the same as in our return example above. When you are trying to achieve a long term compound annual rate of return goal, avoiding big losses is just as important as achieving big gains. Our multiyear simulation optimization model takes this into account, and generates higher weightings to relatively less risky non-equity asset classes (to minimize the chances of suffering a big loss) than other asset allocation models (e.g., single period mean/variance optimization models).

Long time readers will also notice two other differences from previous year's asset allocation reviews. First, we have included three, rather than four target return portfolios. In the past, we have developed model portfolios for target nominal returns of 6%, 8%, 10% and 12%. Assuming a long term rate of inflation of 3%, these are equivalent to target real return portfolios of 3%, 5%, 7%, and 9%. However, if you believe that our forecasts for future real rates of return on different asset classes are on target (see last month's as well as the May and June issues for these forecasts and the logic behind them), it is going to be extremely difficult to achieve a compound annual rate of return of 9% in the years ahead. When we used our simulation optimization model to generate portfolios with the highest probability of achieving this goal, we usually ended up with results very similar to our 7% target return portfolios (though with a lower probability of achieving the target rate of return). Given this, it seemed prudent not to include the 9% target real rate of return portfolios in this year's analysis.

The second major difference is the absence of any review of our "benchmark beating" portfolios this year. As you know, these are model portfolios whose objective is to either deliver higher returns than a domestic benchmark (e.g., 80% equity/20% bonds; 60% equity/40% bonds, or 20% equity/80% bonds) while taking on the same amount of risk, or to deliver the same returns while taking on less risk. Basically, this comes down to a matter of philosophy. In a nutshell, we believe that the purpose of good investment management is to ensure that over the long term, the value of your assets matches or exceeds the value of your liabilities (e.g., the amount of capital you need to achieve your retirement income goal). Our target return portfolios are based on this belief. However, our benchmark beating portfolios are not -- they focus on relative annual returns, rather than achieving one's long term goals. When we started The Index Investor back in 1997, we felt we had to include the latter, because we sensed that relative annual performance was important to a substantial percentage of our potential readers. However, following the substantial bear market we experienced in the last two years, we now believe that this percentage has grown smaller, as more and more people have realized that it is long term performance (and not bragging rights at the pub about last quarter's performance) that really counts. So while we will continue to publish our "benchmark beating" portfolios (there is some value in tradition, after all), we do not plan to update them as frequently as we have in the past.

Dealing with Uncertainty

As we have discussed in previous articles, the main problem with using historical sample data in an asset allocation analysis is estimation uncertainty. You simply can't be sure that the inputs derived from your sample reflect the "true" value for the population as a whole. And even if they do, you can't be sure that the process which generated the historical return data will remain unchanged ("stationary" in stats-speak) in the future. This problem is most acute for average returns; by increasing the frequency of the data collection within a given year (e.g., from quarterly to monthly), the estimation uncertainty associated with the standard deviation and correlation of returns can be substantially reduced. With average returns, however, the only way to reduce estimation uncertainty is to use a longer historical sample (assuming, as noted above, that the underlying returns generating process doesn't change). Generally speaking, the impact of estimation errors (that is, differences between your sample estimate and the true values of a variable) is a function of two factors: the length of your sample period and of your forecasting horizon. As we have seen, using a longer sample period improves the accuracy of your estimate. In contrast, the longer your forecasting horizon, the greater the potential impact of an estimation error, due to the compounding effect. When the forecast horizon is longer than the sample period, expected returns will tend to be biased upwards (note, this is not the case in our analysis, where the forecast horizon is 20 years, but the sample period is 32 years). More broadly, as a rule of thumb, the greater the ratio of forecasting horizon to sample period, the more an investor should suspect assets with high expected returns, due to the potential impact of estimation errors.

Four different approaches have been proposed to reduce the potential impact of estimation errors.

First, you can impose constraints, and limit the maximum weight that can be given to certain asset classes. Statistically, the most logical asset classes to constrain are those whose returns and/or risk are high relative to the average for all the asset classes being used. These are the asset classes that are most likely to be subject to estimation error (e.g., see "Risk Reduction in Large Portfolios" by Jagannathan and Ma, which can usually be found by searching on either google.com or ssrn.com). We have used this approach in our analysis.

Second, if you suspect that the returns generating process may not be stable over time, you can weight your sample data so that more recent data points count more heavily (e.g., see "Time Weighted Portfolio Optimization" by Lee and Stevenson). While we recognize that the means, standard deviations, and correlations in our sample change over time (e.g., many of the latter move closer to 1.0 when markets are falling, then fall back when markets are rising), we are not convinced that this in fact represents a changing in the underlying returns generating process (indeed, there is evidence on both sides of this question). In our analysis, we have assumed that it is a stable process with ups and downs that are likely to average out over our twenty year holding period.

A third approach to minimizing the impact of estimation uncertainty is to explicitly recognize it in your optimization process. This is known as resampling or bootstrapping. To apply it, one starts with the historical distribution of returns for a given asset class (that is, with its mean and standard deviation), repeatedly draws a new sample of returns from it (i.e., one "resamples" it), and then calculates the mean and standard deviation of the new sample. By repeating this process many times, you develop a probability distribution for the mean and standard deviation statistics themselves -- the shape of these probability distributions reflecting the degree of uncertainty about the "true" values of these statistics. Practically, explicitly recognizing estimation uncertainty in this manner makes it clear that portfolios with very different asset weights may be, in the statistical sense, equivalents (and in so doing confirm the point that when it comes to asset allocation, some judgement is required!). On the other hand, an important criticism of this approach is that it is still only based on your sample distribution -- actual reduction of estimation uncertainty (as opposed to making its impact explicit) requires the introduction of additional information (see "Portfolio Choice and Estimation Risk" by Herold and Maruer). Our simulation optimization process uses the bootstrapping approach, as it resamples a distribution (based on either historical or forecast data) to develop a probability distribution of the likely results from a given asset allocation.

The fourth approach to minimizing the impact of estimation uncertainty combines the inputs derived from a historical sample with inputs derived from some prior view of the returns generating process (e.g., combine .67 times your sample asset class weights with .33 times the weights in your prior view). This is also known as the "shrinkage" approach to managing estimation error (e.g., see Jorion, "Bayes-Stein Estimation for Portfolio Analysis") . However, using this approach introduces another source of uncertainty, about the correctness of the model that you use to form your prior view (e.g., see our article on factor models in this issue, as well as "Stock Return Predictability and Model Uncertainty" by Dov Amarov, "Comparing Asset Pricing Models" by Pastor and Stambaugh, and/or "A Shrinkage Approach to Model Uncertainty and Asset Allocation" by Zhenyu Wang).

Different authors have suggested different models that could be used to form your prior view of expected asset class returns, risks and correlations.

If you do not believe that your historical sample data provide any useful information about future returns or risks, you should use an equally weighted portfolio (EWP) as your prior view. If you believe that your historical sample provides more information about risk than it does about returns, you should use what is known as the Minimum Variance Portfolio (MVP) as your prior view. The MVP is the combination of asset weights that minimizes the standard deviation of your portfolio (taking correlations into account). In other words, it is the least risky portfolio you can form given the asset classes you have decided to include in your analysis. The following table shows the minimum variance portfolios for our six different currencies (based on historical real standard deviations and correlations from our 1971-2002 sample):

Minimum Variance Portfolio Weights

	US$	A$	C$	GBP	EURO	JPY
Real Return Bonds	71.3%	83.3%	75.3%	79.8%	69.2%	N/A
Domestic Bonds	15.3%	5.3%	8.1%	5.3%	16.6%	53.5%
Foreign Bonds	3.5%	1.9%	6.2%	5.9%	5.2%	17.5%
Commercial Property	4.6%	3.2%	4.2%	3.1%	4.4%	12.9%
Commodities	1.3%	1.2%	1.4%	2.2%	1.0%	3.6%
Domestic Equity	1.7%	1.2%	1.5%	1.1%	1.6%	4.7%
Foreign Equity	1.5%	3.3%	2.4%	1.9%	1.4%	5.4%
Emerging Markets Equity	0.8%	0.6%	0.9%	0.7%	0.6%	2.4%
	100%	**100%**	**100%**	**100%**	**100%**	**100%**
Std Deviation of Annual Returns	2.11%	2.28%	2.17%	2.23%	2.08%	4.10%

Finally, if you believe that your sample provides useful information about both risk and return, you should use the global market portfolio as your prior. Theoretically, (assuming markets are efficient and in equilibrium) this is the portfolio which maximizes expected return per unit of expected risk. Very closely related to this is another common approach to asset allocation, known as the Black-Litterman model (as described in their paper "Global Portfolio Optimization"), which shrinks the optimal portfolio that results from an investor's future return forecast toward the global market portfolio. Black and Litterman's underlying assumption is that markets are generally in equilibrium, with the expected return on each asset class balancing the available supply and demand of securities. To the extent that a forecast of future returns implies an optimal portfolio that has different weights from the market portfolio, it represents a departure from this equilibrium condition. To adjust for the possibility that this is a result of estimation error rather than actual disequilibrium, Black and Litterman shrink the optimal portfolio back toward the global market portfolio.

However, the global market portfolio itself is not without problems. First, there is the question of what asset classes belong in it. While debt and equity surely belong (though people argue about whether and how to include privately placed issues), a number of writers have argued against including commercial real estate and commodities. While the former is undoubtedly a large asset class, most of it is illiquid, and its returns data tend to be compromised by the fact that valuations are based on subjective opinions rather than objective market prices. The reason commodities aren't included is to avoid double counting -- for example, a substantial portion of the value of the world's physical oil is already implicitly included in the value of the bonds and shares issued by the world's oil companies.

A second issue with the global market portfolio is that it is not at all easy to obtain timely information about the size and composition of the global bond market. Unlike equities, which trade on exchanges, bonds all trade privately ("over the counter"), which makes data about them much more difficult to collect. For example, in 2001 the world's three leading references on the global bond market (produced by Merrill Lynch, the Bank for International Settlements, and the International Monetary Fund) produced three different estimates for its market value (in billions of U.S. dollars): $32,972, $35,327, and $41,792. The good news, however, was that they were all more or less in agreement with each other as to the share of outstanding bonds that had been issued in different currencies.

Finally, the asset class weights in the global market portfolio are constantly changing over time, as investors change their required rates of return on different asset classes. Unfortunately, as data about the market values of different asset classes only becomes available with a significant time lag (and even then is still subject to significant uncertainty), you can never, as a practical matter, be completely sure of the current weights in the global market portfolio.

As a point of reference, we have calculated the following global market portfolio weights for different currencies and assets based on 2002 data, which is the most recent available:

Estimated Global Bond and Equity Market Weights in 2002

Currency in Which Asset is Denominated	**Percent of Global Bond Market Capitalization**	**Percent of Global Equity Market Capitalization**
A$	0.6%	1.8%
C$	1.5%	2.2%
US $	47.4%	54.2%
Euro	22.8%	13.5%
GBP	3.9%	10.7%
Yen	16.7%	7.9%
Other	7.1%	9.7%
-- of which, emerging markets		4.1%
Weight in Global Market Portfolio*	70%	30%

* We used the BIS market size estimate for bonds. Our equity market size information came from the FTSE All World Index (adjusted for market coverage).

Most people reading this table will have the same reaction: I never realized the bond market was bigger than the equity market! Undoubtedly, this results from the latter getting much more publicity (and having more readily available data) than the former. However, when you think about it, it makes sense. Most corporations employ more debt than equity on their balance sheets. Then add to this the debt issued by governments (which don't issue equity), and by various asset backed security vehicles (e.g., mortgage backed bonds) which also have very little equity supporting them. When you consider all these factors, it is no surprise that the bond market is bigger. But that much bigger? Keep in mind that at the end of last year, equity market values were well below their 2000 peaks, while very low interest rates had caused bond market valuations to be very high. On a long term basis, a 60/40 split seems more realistic than the 70/30 split at the end of 2002.

Using the 2002 data, we have calculated what the global market portfolio looked like from the perspective of our six different currency regions:

Global Market Portfolio Weights in 2002
(rounded)

Asset Class	Australia	Canada	Eurozone	Japan	UK	USA
Domestic Bonds*	1%	1%	16%	12%	3%	33%
Foreign Bonds	69%	69%	54%	58%	67%	37%
Domestic Equity	1%	1%	4%	3%	4%	16%
Foreign Equity	28%	28%	25%	26%	25%	13%
Emrg. Mkt Equity	1%	1%	1%	1%	1%	1%
Total	100%	100%	100%	100%	100%	100%

· Includes real return bonds, where available

Obviously, the most striking thing about this table is the heavy weight given to foreign bonds in the global market portfolios. However, we also need to keep in mind that most investors have an aversion not only to risk, but also to regret. Studies have shown that most people are willing to move away from the optimal return/risk portfolio (which, theoretically, is the global market portfolio) if such a move reduces the chances of them underperforming a popular benchmark or peer group (which is usually based on domestic market results), or suffering substantial negative returns during certain periods of time (e.g., see Chow, "Portfolio Selection Based on Risk, Return, and Relative Performance" or "How Much Foreign Stock?" by Herold and Maurer). Given this, as a practical matter we do not believe that many investors are comfortable holding allocation to foreign bonds suggested by the global market portfolio. Hence, in our simulation optimization model, we have limited the maximum allocation to foreign bonds to 40% of our model portfolios.

So where does this leave us?

To derive our final model portfolio allocations, we chose to employ a shrinkage approach by combining the portfolios we created using data from our historical sample (which is subject to estimation uncertainty) with the portfolios based on our forecast of future returns (which is subject to model uncertainty). Given the length of our sample period, and the rather dodgy track record of market forecasters, we decided to give the portfolio based on the sample data a weighting of .67, and the portfolio derived from our forecasts a weighting of .33. These weights are purely subjective, and you should feel free to change them when forming your own portfolio. You may also find it interesting to replace either of our portfolios with the global market portfolio, again using relative weights of your choosing. Frankly, there is no settled theory on the "right" way to make this decision; at best, looking at the results of a variety of approaches can help to inform one's judgement.

The resulting model target real return portfolios are shown in the following tables. Please note that the asset class weights in the first column are based on historical sample data, the weights in the second column are based on our estimates of future asset class returns, and the weights in the last column reflect 2/3 of the historical and 1/3 of the future weights.

	3% Target Return (Historical Return Assumptions)	3% Target Return (Future Return Assumptions)	3% Target Return (Combined Weights)
Real Return Bonds	45%	25%	38%
Domestic Bonds	25%	30%	27%
Foreign Bonds	10%	25%	15%
Commercial Property	10%	5%	8%
Commodities	5%	10%	7%
Domestic Equity	0%	5%	2%
Foreign Equity	5%	0%	3%
EM Equity	0%	0%	0%
Total	100%	100%	100%
Probability of Achieving Target	98.0%	97.0%	N/A
Expected Average Annual Return	4.5%	4.8%	N/A
Standard Deviation of Expected Returns	2.6%	1.0%	N/A

	5% Target Return (Historical Return Assumptions)	5% Target Return (Future Return Assumptions)	5% Target Return (Combined Weights)
Real Return Bonds	0%	5%	2%
Domestic Bonds	20%	15%	18%
Foreign Bonds	25%	25%	25%
Commercial Property	20%	0%	13%
Commodities	0%	15%	5%
Domestic Equity	20%	25%	22%

	5% Target Return (Historical Return Assumptions)	5% Target Return (Future Return Assumptions)	5% Target Return (Combined Weights)
Foreign Equity	5%	0%	3%
EM Equity	10%	15%	12%
Total	100%	100%	100%
Probability of Achieving Target	91.0%	71.0%	N/A
Expected Average Annual Return	7.3%	6.1%	N/A
Standard Deviation of Expected Returns	1.7%	1.9%	N/A

	7% Target Return (Historical Return Assumptions)	7% Target Return (Future Return Assumptions)	7% Target Return (Combined Weights)
Real Return Bonds	5%	0%	3%
Domestic Bonds	5%	0%	3%
Foreign Bonds	40%	40%	40%
Property	15%	0%	10%
Commodities	10%	20%	13%
Domestic Equity	5%	20%	10%
Foreign Equity	0%	0%	0%
EM Equity	20%	20%	20%
Total	100%	100%	100%
Probability of Achieving Target	74.0%	46.0%	N/A
Expected Average Annual Return	8.0%	6.8%	N/A
Standard Deviation of Expected Returns	1.7%	2.1%	N/A

Next month, we plan to conclude our asset allocation review with a look at how shortening the investment horizon from twenty to ten years changes our model portfolio weights. We also plan to compare our model portfolios with those actually employed by a number of successful institutional investors. We also know that many of our readers are both keenly interested and very knowledgeable about the challenges involved in asset allocation. So we close with a final request: if there are other asset allocation topics you'd like us to include in next month's issue, please don't hesitate to get in touch!