Retired Investor

Invest Wisely…Get an Impartial Second Opinion.

This Month's Issue: Key Points

Here at *Retired Investor* we are strong believers in the value of intellectual honesty. When it comes to investing (as with so many other things in life), wishful thinking can get you into trouble. For that reason, this month we are going to take an in depth look at the argument in favor of active investment management. As evidenced by the dominant market share enjoyed by active management, many people agree with this argument (either explicitly or implicitly), so it behooves us to make sure we clearly understand its logic and assumptions.

The cocktail party version of the active management argument goes like this: "Investors don't have equal access to information. And investors don't have equal skills. Therefore, some investors are going to beat the indexes." Unfortunately, this vastly oversimplified version of the case for active management too often wins the point. The critical dimension it leaves out is time. To summarize (and simplify) the article in this issue, consistently successful active management (that delivers better than index fund returns after sales loads, expenses, and taxes have been taken into account) comes down to consistently successful forecasting. This must be based on either access to superior information or use of a superior model. Changes in regulation (e.g., S.E.C. Regulation FD, which sharply limits companies' ability to selectively disclose sensitive information) and the internet's growth (which speeds the flow of information to all investors) have made it much more difficult for active managers to obtain a consistent information advantage. Models suffer from two limitations: changes in the underlying dynamics of the real economy (e.g., due to changing consumer tastes, or the entry of new competitors from China) invalidate their assumptions, while copying by other managers eliminates the edge they provide. Given this, we would, prior to looking at any historical data, expect that only a tiny proportion of active managers would be able to generate statistically significant alpha (a fancy term for returns above the index) over the long time horizons (say, ten years or more) faced by many investors. We also look at comparative performance data, and find that it confirms our prediction.

Our second article this month looks at yet another reason find it so hard to consistently pick winning companies: from a C.E.O.'s perspective, it is very hard to be one year after year. Across a wide range of companies and industries, we see the same cycle repeated. Successful innovation generates high returns and managerial overconfidence. Competition soon follows, prompting increased focus on cutting costs by doing things better. In many cases, organizational resistance to significant change leaves an opening that is exploited by a newer, more innovative competitor. This produces a sharp decline in performance, which triggers a major strategic change. The ones we typically read about are the few that succeed. Most, however, fail, and yesterday's top performers disappear from view (anybody who doesn't believe this need only look at how many of the firms in any large company equity index change from decade to decade).

This Month's Letter to the Editor

How often do you plan to change the asset allocations in your model portfolios?

Our baseline position is that we change our model portfolio asset allocations as infrequently as possible. If we reviewed them every year, we would confront the normal human tendency to sell last year's worst performing asset class, and buy more of last year's best performer. This is a temptation best avoided, because "performance chasing" has been shown to reduce long term performance (e.g., because too often you end up buying something at the top, while selling something else at the bottom – and buy high, sell low isn't a recipe for long term success). As we have repeatedly written, superior long term performance results from identifying the asset allocation that will maximize the probability of achieving the portfolio rate of return you need to reach your long term goals, and then rebalancing over time to maintain these portfolio weights.

On the other hand, there are two circumstances that will trigger a review of our model portfolio asset allocations. The first is the introduction of new index investment products which make it possible for individuals to invest in an asset class that had previously been available only to institutional investors. In recent years, these innovations have included the introduction of commodity index funds, and, more recently, the first hedge fund index

products. Looking to the future, there are at least four potential new product offerings that, if they were introduced, would trigger an asset allocation review. The first would be an index product tracking the private equity asset class. The second would be an index product tracking residential real estate. The third would be the launch of index products that are linked to foreign exchange returns. And the fourth would be index products that are based on the underlying volatility of an asset class. Some have asked if the introduction of an index product that track global commercial property would also trigger a new asset allocation review. Our answer is that it might, though we would first have to see by how much the risk/return characteristics of such an instrument differed from national or region-specific commercial property indexes.

The second trigger for an asset allocation review would be the publication of significant new research findings that cause us to re-think some aspect of our underlying asset allocation methodology. For example, last year we decided that the application of new Bayesian statistical techniques could help us significantly improve the way we handle estimation errors in the inputs we use in our asset allocation models. Looking forward on the methodology front, we are monitoring research in a number of areas (some of which we'll write about later this year). These include (a) asset allocation using higher statistical moments (i.e., coskewness and cokurtosis), (b) advances in combinatorial optimization (the second part of the "simulation optimization" methodology we use to develop our long term target return portfolios), (c) asset allocation using Conditional Value at Risk, and (d) the application of extreme value theory to long term asset allocation decisions.

However, the bottom line is that unless any of these developments (or something similarly important) comes to pass, we aren't going to be changing the target return portfolios' asset class weights.

Global Asset Class Returns

YTD 30Jan04	In USD	In AUD	In CAD	In EURO	In JPY	In GBP
US Bonds	0.80%	-0.66%	2.99%	1.63%	-0.71%	-1.40%
US Prop.	4.30%	2.84%	6.49%	5.13%	2.79%	2.10%
US Equity	2.20%	0.74%	4.39%	3.03%	0.69%	0.00%
AUS Bonds	-0.78%	-2.24%	1.41%	0.05%	-2.29%	-2.98%
AUS Prop.	1.49%	0.03%	3.68%	2.32%	-0.03%	-0.71%
AUS Equity	-0.70%	-2.16%	1.49%	0.13%	-2.21%	-2.90%
CAN Bonds	-1.77%	-3.23%	0.42%	-0.94%	-3.28%	-3.97%
CAN Prop.	1.86%	0.40%	4.05%	2.68%	0.34%	-0.34%
CAN Equity	0.40%	-1.06%	2.59%	1.23%	-1.11%	-1.80%
Euro Bonds	-0.94%	-2.40%	1.25%	-0.11%	-2.45%	-3.14%
Euro Prop.	2.48%	1.02%	4.68%	3.31%	0.97%	0.29%
Euro Equity	2.00%	0.54%	4.19%	2.83%	0.49%	-0.20%
Japan Bonds	1.65%	0.19%	3.84%	2.48%	0.14%	-0.55%
Japan Prop.	14.53%	13.07%	16.72%	15.36%	13.02%	12.33%
Japan Equity	0.20%	-1.26%	2.39%	1.03%	-1.31%	-2.00%
UK Bonds	1.71%	0.25%	3.90%	2.54%	0.20%	-0.49%
UK Prop.	5.05%	3.59%	7.24%	5.87%	3.53%	2.85%
UK Equity	-0.50%	-1.96%	1.69%	0.33%	-2.01%	-2.70%
World Bonds	0.45%	-1.01%	2.64%	1.28%	-1.06%	-1.75%
World Prop.	6.20%	4.74%	8.39%	7.03%	4.69%	4.00%
World Equity	1.90%	0.44%	4.09%	2.73%	0.39%	-0.30%
Commodities	0.60%	-0.86%	2.79%	1.43%	-0.91%	-1.60%
A$	-1.46%	0.00%	-3.65%	-2.29%	0.06%	0.74%
C$	2.19%	3.65%	0.00%	1.37%	3.71%	4.39%
Euro	0.83%	2.29%	-1.37%	0.00%	2.34%	3.02%
Yen	-1.51%	-0.06%	-3.71%	-2.34%	0.00%	0.68%
UK£	-2.20%	-0.74%	-4.39%	-3.02%	-0.68%	0.00%
US$	0.00%	1.46%	-2.19%	-0.83%	1.51%	2.20%

Equity Market Valuation Update

Our equity market 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 average rate of productivity growth an economy will achieve in the future. As described on our website (see the green button in the members' section labeled "domestic equity"), we use both high and a low productivity growth scenarios. Given these assumptions, here is our updated market valuation analysis at the end of last month:

Country	Real Risk Free Rate Plus	Equity Risk Premium Equals	Required Real Return on Equities	Expected Real Growth Rate* plus	Dividend Yield Equals	Expected Real Equity Return**
Australia	3.43%	4.00%	7.43%	4.90%	3.77%	8.67%
Canada	2.53%	4.00%	6.53%	2.10%	1.78%	3.88%
Eurozone	1.66%	4.00%	5.66%	2.50%	1.90%	4.40%
Japan	1.62%	4.00%	5.62%	2.70%	0.90%	3.70%
U.K.	1.96%	4.00%	5.96%	2.50%	3.20%	5.70%
U.S.A.	2.23%	4.00%	6.23%	4.50%	1.60%	6.10%

**High Productivity Growth Scenario. See our website (green button, "domestic equity"), for assumptions used in both productivity growth scenarios for each region.*

*** When required real equity return is greater than expected real equity return, theoretical index value will be less than actual index value – i.e., the market will appear to be overvalued.*

Country	Implied Index Value*	Current Index Value	(Under) or Overvaluation in High Growth Scenario	(Under) or Overvaluation in LowGrowth Scenario
Australia	149.01	100.00	(49%)	(7%)
Canada	40.18	100.00	60%	67%
Eurozone	60.13	100.00	40%	59%
Japan	31.91	100.00	68%	76%
U.K.	92.49	100.00	8%	35%
U.S.A.	92.49	100.00	8%	41%

** High productivity growth scenario.*

The Case for Active Management

Here at *Retired Investor*, we believe in intellectual honesty. For that reason, this month we are going to take an in depth look at the argument in favor of active investment management. As evidenced by the dominant market share enjoyed by active management, many people agree with this argument (either explicitly or implicitly), so it behooves us to make sure we clearly understand its logic and assumptions.

The cocktail party version of the active management argument goes like this: "Investors don't have equal access to information. And investors don't have equal skills. Therefore, some investors are going to beat the indexes." Unfortunately, this vastly oversimplified version of the case for active management too often wins the point. Unless, of course, there's an index Jedi present (like you), who asks our active investor two questions. "Does your argument also mean that your active manager will be able to persistently outperform the index, year after year, after taking sales loads, expenses, and taxes into account? And, if it does, how can I identify these superior managers in advance?" At this point one of three things will happen. Your questioner might decide that he urgently needs to chat with his Aunt Emily across the room. Or, he might impulsively reply "Of course it means they'll persistently outperform -- and you can spot them by looking at the returns they've delivered in the past", and then hurry over to Aunt Emily before you can ask him to name the great fund managers he spotted five years ago, and how their returns have compared to your index fund's since then. Finally, he might, if he's a sensible man, acknowledge that you have asked the two questions that lie at the heart of the active management versus indexing debate. If you receive this third response, you need to be prepared for a rather longer discussion. To varying degrees, it will proceed as follows:

First of all, we need to clearly acknowledge that unlike indexing, active management is a zero sum game. When you index, the only risk you are taking is related to the variability (i.e., the volatility) of the returns on the asset class as a whole. This is variously (and confusingly) called "market" or "systematic" or "beta" risk. You are not taking any risk that is specific and unique to a member of the asset class (e.g., a company issuing a bond or a share), because in a diversified portfolio (e.g., the one that comprises the asset class your

index fund tracks) all those risks will cancel each other out. On the other hand, the additional returns that are compensation for taking those company specific risks will also cancel out (e.g., in any given year, the returns for holding some company specific risks will be positive, while others will be negative). When you decide to hire an active manager (whether you invest in an actively managed mutual fund, or do the investing yourself), you (explicitly or implicitly) believe that he or she has some advantage that will enable him or her to come out ahead (that is, deliver positive returns above what your index fund earns) in this less-than-zero sum game over some time period.

Why is active investing a less-than-zero sum game, you ask? Because collectively, all active managers must underperform index funds by the weighted difference between their respective costs (which can be defined to include some combination of sales loads, operating expenses, trading costs and taxes). But not your active manger. She has an edge (right?) that will enable her to invest in a mix of securities (e.g., stocks) that is different from the mix in the index fund, and that will, as a result, generate returns (after expenses, etc.) that are above those on the index fund. Otherwise you wouldn't be investing in her fund instead of an index fund, right? So what is her edge? Will it last? And will it be sufficient to deliver the returns above the index fund that you expect? At this point, your active manager advocate has just left to go get another drink, with a slightly worried look on his face. You might want to do that too at this point, since what follows is unavoidably going to get a bit technical (we'll do our best to be gentle!).

In the world of professional investors, the additional return you expect your active manager to earn is known as "alpha". This name comes from the way it is typically identified, by regressing an active fund's returns in a given period against the returns on one or more index funds. The weights on each index fund (which must sum to 1.0) are known as "betas", and the portion of the active manager's return which cannot be explained by the index fund return is known as "alpha". Think of "alpha" as the amount your active manager would be expected to earn if the return on the index fund was zero. In other words, alpha is the return you earn for taking company specific risk.

As we noted before, in different periods of time, an active manager may have positive or negative alphas. And just as the variability of market risk and return give rise to a risk/return ratio, so too is the case with our new friend alpha. Again, using investment-speak,

the ratio of the average alpha earned over a period to the variability of alpha (i.e., the standard deviation of alpha) is known as the "information ratio", or simply "IR" for short. While we're on the subject, we should also throw in another investment-speak term: "tracking error" (since it is a measure of how closely our active managers returns track the returns on the index fund benchmark). Don't let this confuse you: tracking error is only a fancy name for the standard deviation of alpha. But back to our story. IR is a very useful tool. For example, if my active manager delivers average alpha of 2% per year, with a standard deviation ("tracking error") of 4%, she has an IR of .5. Alternatively, you are willing to take on 6% "tracking error" or risk above and beyond the risk of the index fund. How much should you expect to earn in return? If you know your active manager has an average IR of .5, you also know your average expected additional return (alpha): 3% (.5 x 6%). Similarly, say you want to earn 3% above the index fund benchmark, and you know an active fund manager has an IR of .5, and an average tracking error of 4%. Armed with this information, you can reasonably conclude that this manager probably won't achieve your return goal, since her expected alpha of 2% (.5 x 4%) is less than the 5% alpha you want. To get that 5% alpha, you'll either have to find a manager with a higher IR (e.g., .75 for a tracking error of 4%), or take on more risk (e.g., 6% tracking error if IR is .5).

Well, that was easy enough, wasn't it? However, there are still a few questions we haven't answered yet. Let's start with a big one: where does alpha come from? What is the nature of your active manager's competitive advantage? And will it be enough to deliver the IR you seek?

The essence of the answer to this question is simple: successful active management comes down to successful forecasting. So, intuitively, the maximum amount of alpha (or IR) an active manager can create (for a given level of tracking error) is a function of two variables: (1) the accuracy of her forecasts, and (2) the number of forecasts that she makes.

In their book Active Portfolio Management, Grinold and Kahn quantify this intuition into what they call "the fundamental law of active management": Maximum IR equals the "Information Coefficient" times "Breadth". The information coefficient (or IC) is the correlation between a manager's forecasts (e.g., for the alpha a stock is expected to produce in a given period) and their actual outcomes. Breadth is the number of independent forecasts that are made in a given period. Let's look at each of these a little more closely.

First of all, what is a reasonable value for IC? Let's start with one extreme: if returns follow a random walk, as they would in a perfectly efficient market (where all information is instantly reflected in an asset's price), then IC would equal zero, and forecasting would be impossible. So, right away, we know that, given the amount of money that has been invested in actively managed funds, a lot of people must believe that financial markets are not, to varying degrees, perfectly efficient. The good news is that a growing body of research says they are right. While this is still a hotly contested area, more and more studies are finding that returns are, (in retrospect at least) slightly predictable, so that a positive IC is theoretically possible (see, for example, the paper "Model Uncertainty, Thick Modeling, and the Predictability of Stock Returns" by Aiolfi and Favero). Still, most estimates of active managers' average IC is quite low -- e.g., on the order of .05 to .15, which implies being right only slightly more than 50% of the time.

What might give rise to a positive IC? What, in other words, might be the source of an active manager's "edge"? Basically, there are only two potential answers to this question. The first is superior information, and the second is a superior model. An information advantage typically can come from either of two sources. The first is private information. This is not the same as "inside information." All inside information is private, but not all private information is inside information. For example, a resourceful active manager who wants to learn how well a company's new product is selling can either (a) wait for the company's quarterly financial report to be released (at which time the information will be public, and widely available), or (b) go count the trucks coming out of the company's plant (which will produce information that is private, but not insider). The second source of information advantage is the fact that even public information takes a while to reach all investors (although with the internet that time has been cut quite sharply).

Broadly speaking, there are two types of superior model. The first gives you a better understanding of the value of an asset than other investors, and the second gives you a better understanding of the way other investors are likely to behave. Superior valuation models generally fall into three classes: (a) bottom up models, where you forecast the future cash flows for individual assets, (b) top down models (such as the factor models discussed in our August, 2003 issue), in which you forecast the returns on different assets based on their

loadings (that is, there betas in a regression model) on different factors, and the returns you earn for holding these factor risks, and (c) market based models, where you back out other investors consensus bottom-up or top-down based views on an asset, and decide whether or not you think they are reasonable. All three of these modeling approaches have the same goal: to help the active manager decide whether or not an asset's current market price is above or below its true (also known as its "fundamental") value.

The second modeling approach isn't concerned with value, but rather with how other investors are likely to behave in the future. Will they be buying an asset (and driving its price up), or selling it? Typically, these models are either based on technical indicators (e.g., moving price averages, trading volume, and the like) or on theories of human behavior (i.e., systematic over or underreaction that can be profitably exploited). Both of these modeling approaches have the same goal: to determine the direction in which momentum will move an asset's price over some future period.

Let's move on now to the other part of the fundamental law of active management. As you recall, this states that maximum IR is limited by both forecasting ability and the breadth with which that ability is applied. Unfortunately, the breadth of an active manager's strategy is harder to measure than his or her forecasting ability. We can, however, make a good approximation of it. Most important, breadth refers to the number of independent forecasts made by an active manager during the period over which the IR is measured (say, one year). Consider, for example, an active manager who focuses on a group of 100 stocks. If that manager uses a bottom up approach, and values each stock once per year, breadth would equal, at most, 100, assuming no common valuation assumptions were used. However, suppose the manager used a top-down model, which valued the 100 stocks using four different factors. In this case, breadth would equal only 4. On the other hand, the top-down manager could increase her breadth by doing her valuation analyses more than once per year. For example, if she did them quarterly, breadth would equal 16. This example makes a very important point: because different active management strategies involve differing numbers of forecasts per period, they should produce (for a given level of forecasting ability) differing levels of IR or alpha. Consider three examples. The first strategy (call it market timing, or tactical asset allocation) involves estimating the returns on eight different asset classes once per year. The second strategy involves quarterly switching between four different equity

styles (e.g. large and small cap growth and value), eight equity sectors, and two bond styles (long and short maturity, and high and low credit risk), based on two forecasts (for interest rates and economic growth), as shown in the following table:

Economy:	**Recession**	**Strengthening**	**Strong**	**Weakening**
Interest Rates:	**Falling**	**Bottom**	**Rising**	**Peak**
Broad Equity Index	Bottoming	Increasing	Peaking	Declining
Highest Relative Return from Growth v. Value (one period ahead):	Growth	Value	Value	Growth
Highest Relative Return from Large v. Small Cap (one period ahead):	Small	Small	Large	Large
Highest Relative Return from Sectors (one period ahead):	Cyclicals and Technology	Basic Materials, Industrials	Energy, Staples	Utilities, Financials
Bond Investments (one period ahead):	Higher Risk Issuers	Shorter Maturity (Duration)	Lower Risk Issuers	Longer Maturity (Duration)

The third strategy involves independent monthly evaluations of the likely returns on fifty different stocks (classic stock picking). Which strategy should produce the highest IR (assuming the same IC)? The breadth of the first one is 8; the breadth of the second (contrary to what you might first think) is also 8 (4 x 2); and the breath of the last one is 600 (12 x 50). No contest: in this case, stock picking, rather than asset allocation or style rotation, should theoretically produce the highest IR, given a constant IC. (Note: for more on this point, see "Asset Allocation Versus Security Selection" by Kritzman and Page, and "Value of Skill in Security Selection Versus Asset Allocation in Credit Markets" by Dynkin, Hyman, and Wu).

Technically, (and for reasons that are to convoluted to explain here), the actual relationship among these variables is that IR equals the product of the Information Coefficient (IC) times the square root of breadth. For the full explanation, you need to read Grinold and Kahn's book. For now, please trust me on this. To carry on with our example, the maximum potential IR from the first two strategies, given an IC of .10, is only .28 (.10 times 8 to the 1/2 power), while the potential IR from the latter is an astounding 2.45.

At this point, I can imagine what you're thinking: Wait a minute! How do you reconcile that last sentence with the results of historical studies that have found that asset allocation has a much bigger impact on returns than stock selection? Glad you asked the question.

There are a number of possible answers to it. One is that those fifty stocks really weren't independently evaluated. For example, if the valuation analyses used a common assumption for future economic growth or interest rates, the actual number of independent forecasts would have been much lower than fifty. The second possible answer is that the IC isn't constant -- for example, perhaps it is higher for some types of decision (e.g., asset allocation), or perhaps it varied from month to month. Finally, there is a third explanation, to which we'll now turn.

Remember back at our cocktail party, when active management was justified by the claim that investors had unequal skills and access to information? By now, you realize that all that referred to was forecasting ability, or an active manager's IC. However, in order to turn forecasting ability into actual returns, it has to be implemented via the allocation of real money to real investments in a real portfolio. Ideally, there is a perfect correspondence between the forecast outcome and the resulting portfolio weights. What our cocktail party friend failed to mention is how often this doesn't happen in practice. To begin with, many investors, for very good governance reasons, place constraints on the portfolio positions an active manager can take (for a very good article on this, see "Why Constrain Your Mutual Fund Manager" by Almazon, Brown, Carlson, and Chapman). For example, these constraints might include limits on the maximum investment that can be made in one company, country, or industry; limits on turnover per period (to minimize trading costs), or prohibitions on using leverage, or derivatives, or selling short. In their paper "Portfolio Constraints and the

Fundamental Law of Active Management", Clarke, de Silva and Thorley created what they call the "Transfer Coefficient" (you got it, TC), to measure the correlation between the portfolio recommended by the forecasts, and the one that could actually be implemented, given the constraints placed on the manager's action.

In their research, the authors found that "TC values of as low as .3 may be common among long-only U.S. equity managers" (e.g., at typical equity mutual funds). TC is a very interesting statistic. Because it is a correlation, when you square it you get a figure that describes the percentage of variation in portfolio returns that is actually attributable to forecasting ability. For example, at a fund with a TC of .3, only 9% of the variation in the fund's returns is due to the manager's forecasting ability -- the rest is random noise, or, put another way, luck. As the authors note, "managers with low transfer coefficients will experience frequent periods when [their forecast] works, but performance is poor, and periods where performance is good even though the return forecasting process failed." The other nice thing about TC is that mathematically, it fits right into our IR equation, which now looks like this: IR equals IC times TC times the square root of breadth (BR).

So, to go back to our previous example, let's put a TC of .3 on our stock picking manager. This reduces her previous maximum IR of 2.45 to a still very impressive .75. This has a number of very important (and disturbing) implications beyond the substantial reduction in potential IR and alpha. First, it makes performance attribution extremely difficult. Given a low TC, what is a manager's alpha (and IR) really telling you about their true skill (that is, their forecasting ability)? A lot less than a lot of people would like you to believe. Second, because a low TC reduces IR and alpha, it also reduces a fund manager's potential compensation, because it limits their ability to fully exploit whatever forecasting skill they have. Is it any wonder why so many good mutual fund managers have left to run hedge funds, where the TCs are much higher (and where compensation is often a hefty percentage of alpha)? To put it delicately, this raises awkward questions about the quality of those mutual fund managers who have not left to run a hedge fund.

While we're on the subject of dirty little secrets about active management (or, to look at it another way, a fully developed theory of it), we also need to consider transaction costs. Let's consider again our intrepid mutual fund investor, who is willing to take on 5% more risk above her index fund benchmark in the pursuit of higher returns. With our stock picking

active manager, whose IR is .75, those additional returns (alpha) should amount to 3.75% (5% x .75) per year, right? Maybe. There are three important (that is, potentially costly) uncertainties here.

First, it isn't clear (in our analysis) whether or not the manager's TC included a limit on turnover. And second, that alpha number does not include the mutual fund manager's expenses. Its easy to see how this could lead to a very disappointed mutual fund investor. Assume the fund has annual turnover of 100% (that is, it sells all its holdings once per year and buys new ones), and that the one way cost of a trade is .78% (that is, 78 basis points, which is a conservative estimate -- see our March, 2003 article on mutual funds' costs). The funds actual trading costs (as a percentage of its assets) will be about 1.5% per year (2 x .78). And lets assume that the fund's annual expense ratio is 1.25%. This means that we have to reduce that gross alpha of 3.75% by 2.75% (1.50% + 1.25%). In other words, after taking transaction costs and expenses into account, our mutual fund investor receives a net alpha of only 1.00% in exchange for taking on 5% more risk than the index fund benchmark. Finally, we also have to consider the tax impact on our mutual fund investor of all that turnover. This could easily (depending on whether or not those transaction costs are factored into the TC) reduce the final after tax alpha realized by our investor to less than one percent (or, put another way, it could reduce the realized IR to less than .20).

Up to now, we've only talked about what goes on in a single year at an actively managed fund. However, as you recall from our initial cocktail party conversation, a critical question about active management is whether or not successful fund managers can persistently deliver superior returns (e.g., positive alpha or IR) year after year. At this point, we're just going to look at what theory says about this (the data comes in the next section!). In a nutshell, there are good reasons to believe that it should be very difficult for an active fund manager to persistently generate positive alpha or IR.

The assumption that a manager can sustain a superior model from year to year seems has good arguments against it. First, the underlying economic process that generates returns probably isn't stationary -- that is, the variables that are important in the manager's return forecasting model, and/or the relationships between them tend to change over time. One ironic aspect of this is that as a successful manager's forecasting model becomes well known (e.g., think of the value and size effects), they themselves become part of the return generating

process! For a fascinating discussion of this, see "Predicting the Stock Market" by Hellstrom and Holmstrom. To put it another way, profitable investing strategies tend to be self-destructive. As more capital is used to exploit them, they tend to move market prices against themselves, while also becoming more visible and thereby making it easier for other investors to copy them and compete away their alphas.

Cognitive psychology provides the second set of arguments against the persistence of superior forecasting skills. To sum up a vast amount of literature in a few sentences (see Heuristics and Biases: The Psychology of Intuitive Judgement by Gilovich, Griffin and Kahneman for a full discussion), first impressions have a stronger impact on us than later information. This means that it takes less information for us to form a view than it does to change it. Moreover, once we have formed a view about something, two things happen. First, it becomes "affectively charged" -- we make an emotional investment in it. This makes it even more difficult to change an initial view, as you must overcome not only rational but also emotional hurdles to do so. Second, once formed, an opinion affects the way we look at and process information. In contrast to the scientific method (which, as you recall, is based on disproving theories, rather than supporting them), we tend to pay more attention to information which supports our existing views, and attach less importance to any information we receive which contradicts them (for a fuller description of this, see "First Impressions Matter" by Rabin and Schrag). This "confirmation bias" tends to engender overconfidence in us about the correctness of our views. Moreover, this overconfidence is compounded by our natural tendency toward "biased self-attribution" -- the tendency to credit ourselves with skill when our forecasts are proven correct, while blaming adverse outcomes on bad luck. Taken together, these factors make it more likely that a previously successful active manager will continue to use a forecasting model even after its effectiveness (that is, the resulting IC) has declined.

With respect to superior information, we have already noted that the idea that superior forecasting ability is linked to the slow diffusion of public information is rapidly being eclipsed by technology. That leaves us with superior access to private information. At first glance, this seems like it could be a sustainable basis for a persistently positive IR. However, it is hard to disentangle this from the underlying model which determines the nature of the

private information which the active manager will seek out. As such, it seems subject to the same limitations of the model itself.

Thus far, we have only looked at the theory of active management. Our examination leads us to the conclusion that persistent positive alphas and high information ratios are likely to be quite rare. The question to which we will now turn is whether studies that have used historical active manager performance data have found this to be the case.

A number of studies over the years have looked at the information ratios actually achieved by active managers. We should begin by noting the limitations of these studies. First, they estimate alpha after regressing fund returns on various factors (e.g., the market return, as well as value, size, and momentum). To the extent that the fund loadings and return premia for these factors vary over time, this will cause (sometimes large) errors in the estimated fund alphas (for more information on this, see "Estimating the Dynamics of Mutual Fund Alphas and Betas" by Mamaysky, Spiegel, and Zhang, "A Matter of Style" by Russ Wermers, and "In Search of True Performance: Testing Benchmark Model Validity" by Allen and Soucik). Second, as we have seen, a low transfer coefficient (TC) causes the percentage of realized alphas that are due to luck to rise in proportion to those that are due to skill.

We should also note that, with a few exceptions, most of these are based on gross alphas, which don't take transaction costs, operating expenses, or taxes into account. Most of these studies have focused on the information ratios achieved by the top quartile of active fund managers in different asset classes. Why the focus on top quartile managers? First, as Grinold and Kahn noted, "overall, there is no evidence for average active management's producing exceptional returns." This was echoed by the Bank for International Settlements in its September, 2003 Quarterly Review, which noted the "widespread recognition that, at least in the largest and most informationally efficient markets, actively managed funds do not, on average, earn returns sufficient to offset their costs." As we noted at the outset, in the aggregate, active management is a less-than-zero sum game.

The following table shows the results of various studies which have directly estimated the information ratios and alphas achieved by top quartile managers in different asset classes:

Asset Class	**William Mercer, 1995 to 1999, Top Quartile IR Gross**	**Rogers Casey, 1991 to 1996 Top Quartile IR Gross**	**Gupta, Projogi, Stubbs 1992 -1997 Top Quartile Gross IR**
U.S. Equtiy	.36 (large cap) to .94 (small cap)	.13 (Large Value) to 1.17 (Small Growth)	.51 (Large Cap) to .88 (Small Cap)
European Equity	.52	.55 (EAFE)	.68 (EAFE)
Pacific Equity	.54 (ex Japan)		
Japan Equity	.52		
Emerging Markets Equity	.50	.39	.73
U.S. Fixed Income		1.02	.76
International Fixed Income		.29	.53

In their book, Grinold and Kahn used data from the early 1990s to estimate top quartile alphas for U.S. mutual funds, after fees, but not taxes. For bond funds, the top quartile IR was (.22), for equity funds it was .58. Two other studies are worth looking at. The Frank Russell Company looked at the actual gross alphas that were earned by different active managers in different asset classes. Taking a different approach, the University of Texas Investment Management Company compared the average returns earned by top quartile managers with those earned by third quartile managers between 1980 and 1997. These results are shown in the following table:

Asset Class	**Frank Russell Company, 1992 - 2000 Top Quartile Manager Average Gross Alpha**	**University of Texas Investment Management Company; 80 to 97 Average Top Quartile Gross Return (not alpha) less Average Third Quartile Return**
U.S. Equity	1.8%	1.70%

Asset Class	**Frank Russell Company, 1992 - 2000 Top Quartile Manager Average Gross Alpha**	**University of Texas Investment Management Company; 80 to 97 Average Top Quartile Gross Return (not alpha) less Average Third Quartile Return**
Non U.S. Equity	4.0%	2.10%
Emerging Market Equity	3.6%	
Japan Equity	6.3%	
UK Equity	1.8%	
Europe Equity	4.6%	
U.S. Fixed Income	0.6%	2.1%
Global Fixed Income	0.7%	4.8%
Real Estate		4.0%
Venture Capital		16.5%
Leveraged Buyouts		18.4%
Hedge Funds		22.7%

These tables make some very important points. First, information ratios above .50 are quite rare, even on a gross basis. And even an IR at this level presupposes a relatively high degree of skill in manager selection on the part of the investor. Second, the really high alphas seem most likely to be found in the very asset classes (venture capital, leveraged buyouts and hedge funds) that are out of reach for most retail investors, and to which the most talented active managers have flocked.

It is, to put it mildly, a very sobering picture for anyone considering investing in an actively managed retail mutual fund. What about persistence, or the ability of active fund managers to deliver positive alpha and IR year after year? Two of the most comprehensive recent studies on this subject came to the same conclusion: there is very little evidence that positive past performance persists (in other words, they find that superior past performance is not a useful predictor of superior future performance). See "A Review of Research on the Past Performance of Managed Funds" by the Australian Securities and Investment Commission, and "Past Imperfect", published by the U.K. Financial Services Authority. Even

Grinold and Kahn (authors of the book, Active Portfolio Management) could only note that "the conclusion of these studies is that even if performance does persist, it doesn't persist at an impressively high rate." And these are two of the leading advocates of active management!

On the other hand, there are a few studies that seem to suggest that some managers can, in fact, consistently deliver positive alpha. However, they all suffer from the same limitation: the short length of their data series makes it impossible to reach a statistically significant conclusion as to whether the observed alphas were due to luck or skill. The following table shows how big this problem is:

Information Ratio	**Years of Data Needed for Statistically Significant Alpha (T Ratio > 2)**
1.0	4
.75	7
.67	9
.50	16
.33	36
.25	64
.20	100
.10	400

Given that top quartile managers typically have Information Ratios of .5 or less, and performance histories of less than sixteen years, it is basically impossible to tell from their performance whether they are truly skilful or just plain lucky. As Grinold and Kahn note, "it is a fact of investment management life that proof of investment management prowess will remain elusive."

However, three creative analyses of the active management issue have recently moved us closer to a definitive answer. In "Mutual Fund Flows and Performance in Rational Markets" by Berk and Green, the authors suggest a reason why superior performance is not likely to persist. Their thesis is that "the fact that investments with active managers do not outperform passive benchmarks is a consequence of the competitiveness in the market for

capital investment. If investors compete with each other for superior returns, they end up ensuring that none exist." However, and this is a very good insight, they also note that "this lack of persistence, however, does not imply that differential ability across managers is unrewarded." How could this be?

The authors ask us to "imagine an economy [in which everyone has complete] information. Skilled investment mangers exist who can generate positive, risk-adjusted returns [in excess of their benchmark indexes]. Managers and investors alike know who these superior managers are. What would the returns these managers provide to investors look like? In equilibrium, investors who choose to invest with active managers cannot expect to receive positive excess returns on a risk-adjusted basis. If they did, there would be an excess supply of capital to those managers who achieved superior returns. Every investor in the economy who held asset of equivalent risk would want to sell those assets and invest with the superior active mangers instead. Markets can only clear when the expected return to investors in these funds equals the expected return in alternative investment opportunities."

"If skill or superior ability in active portfolio management could be deployed on an unlimited scale without dissipating its effectiveness, then in a given risk class, all funds in this hypothetical world would flow to the manager with the highest ability. However, it seems reasonable to assume that managerial ability to generate excess returns cannot be effectively employed on an unlimited scale. If there are decreasing returns to scale in the use of [investment management] ability, funds will be invested with skilled managers only up to the point where the manager provides investors with expected returns equal to those available in passive alternatives."

"This also suggests the mechanism the skilled manager can use to capture a substantial share of the value created by his or her skills. He or she can charge a fee that is proportional to the assets under management. With this incentive scheme, investment will flow into the fund until it is so large that its expected excess return is zero. Highly skilled managers will manage larger funds, earning more income than less skilled peers."

In a more realistic economy in which people lack perfect information (that is, one in which uncertainty exists), investors will need to infer fund managers' relative ability from their past returns. In this case, the same process will play out, but over a longer period of time. The authors' key conclusion is that the process they describe (which essentially says

that investment capital will flow into its most productive uses, bidding up their price, and reducing their returns to levels in line with other assets of similar risk) "necessarily implies that investors cannot expect to make positive excess returns going forward, which also implies that superior performance cannot be predictable in advance."

The second important analysis is "Can Mutual Fund Stars Really Pick Stocks?", by Kosowski, Timmerman, White and Wermers. These authors used U.S. equity mutual fund performance data from 1962 to 1994 (net of expenses, but not taxes), and an innovative modeling approach (bootstrapping) to simulate a much longer data series. They begin by noting that "in the huge universe of funds, it is natural to expect that some funds will outperform market indexes by a large amount simply by chance." Their study attempts to distinguish between those managers whose superior performance over time is due to skill, and those for whom it is due to luck. They find that "superior funds that beat their benchmarks (net of expenses) by an economically and statistically significant amount do exist. [However] we also find strong evidence of inferior funds. We do not find it surprising that large numbers of inferior managers exist in our sample, since performance measurement is a difficult task requiring for precision a long fund lifespan. This evidence of inferior fund management is consistent with consumers who have difficulty in identifying the few fund managers that can beat the market, and especially in terms of judging the skills of managers of relatively new funds." Specifically, they found that truly talented managers accounted for only five percent of their sample. However, they did not take taxes into account. Were these included, the percentage of fund managers who beat their respective indexes on the basis of their superior skill would have been even lower than five percent. With respect to performance persistence, the authors found some evidence that the top ten percent of funds ranked by their three year performance continued to deliver top performance for another year. However, they noted that this aspect of their findings needed further study, and that most of their evidence was consistent with the predictions of Berk and Green.

In his subsequent paper “Is Money Really Smart?”, Professor Russ Wermers focused squarely on the issue of performance persistence. Using a set of data covering U.S. equity fund holdings and performance between 1976 and 1994, and including expenses and trading costs (but not taxes), he found that a complicated process causes style adjusted superior performance to persist for up to two years. First, consumers disproportionately invest their

savings into the previous year's top performing funds. This flow continues for the next two years, due to what Wermers calls a "reputation effect". Fund managers "invest these cash flows into high past return stocks to refresh the momentum in their portfolio returns." Moreover, because top performing funds tend to have correlated stock holdings, their investment of their new cash inflows into these securities further pushes up their prices. However, once the abnormally high [cash inflows] cease, the prices of these stocks tend to decrease. Wermers notes that his "finding that [fund] performance does not seem to persist after controlling for cash inflows casts doubt on previous studies that found that managers have talents in choosing stocks that beat their benchmarks."

So, where does this leave us? Our prior theory-based view suggests that persistent positive alpha (or a high IR) should be very difficult for an active fund manager to achieve, particularly on a net basis. However, we also noted that the measurement of alpha is itself problematic. Our examination of studies based on historical and simulated data confirmed that consistent positive alphas and high information ratios are rare, and usually not statistically reliable. Ideally, we should try to combine these two perspectives to help us reach a more definitive conclusion.

A number of very recent papers have used some advanced methods (basically, Bayesian statistics and simulation) to do exactly this. The best one we have seen is "Mutual Fund Performance With Learning Across Funds" by Jones and Shanken. They utilized a sample of more than five thousand U.S. mutual funds, with an average life of about six years. They include prior uncertainty about both the factor model used to generate estimated alphas, as well as the true extent of fund manager skill. They calculate alphas after expenses and trading costs (though they use a relatively low estimate for the latter) but before taxes. The authors' analysis combines different prior assumptions about the alpha estimating model and the likelihood of persistent managerial skill with the actual fund results to produce a combined (technically, a posterior) view about likely fund alphas. Their findings are interesting. Given a highly skeptical prior view about the likelihood of persistent manager skill, the authors find (depending on the factor model used) an average posterior expected alpha of between (.69%) to (.74%). In other words, like many others they find that on average active management doesn't generate positive alpha. On the other hand using these same assumptions but looking at the extremes of the distribution of fund alphas instead of their

average, one finds a maximum expected alpha of between 1.86% to 4.22% (before taxes). Basically, Jones and Shanken reach a familiar conclusion: while managers with persistent superior active investment management skill probably exist, they also appear to be very rare.

Unfortunately, there is no easy way to identify these future winners in advance. Past performance has been shown to be of no help with this task. Nor have the attempts by different fund rating services to predict future superior performance proved to be useful (see "Morningstar Ratings and Mutual Fund Performance" by Blake and Morey). Institutional investment consultants suggest interviewing managers, and choosing those with high quality people and investment processes. However, this is a luxury unavailable to most individual investors in actively managed mutual funds. In their paper "The Dimensions of Active Management", Waring and Siegel succinctly sum up the situation: "If we cannot usually rely on past performance to select active managers, then how can we select them at all? We don't have a recipe, and we know there aren't any recipes. If there were, everyone would be following it, and of course, then it wouldn't work…Each investor has to develop his or her own methodology for forecasting manager alphas [and building portfolios of active managers]…If you don't think you can do this, maybe you shouldn't hire active managers." The logic behind this conclusion is quite clear. As Kritzman and Page noted in their study, "as beneficial as it is for skilful investors to focus on activities [which have a high dispersion of potential returns], it is equally important for unskilled investors to avoid them." Peter Bernstein (the founding editor of the Journal of Portfolio Management) made the same point when he noted that investors "who cannot identify skilled managers would do well to index…[because] indexing should do better than unskilled managers." So, unless you or your financial advisor have an information or model based edge in forecasting the future performance of investment managers (i.e., an IC greater than zero), index funds will logically have higher long term expected returns (after loads, fees, trading costs and taxes) than their actively managed peers in the same asset class.

Delivering Superior Returns: the CEO's Perspective

As we have seen, it is very difficult for active managers to consistently deliver returns above those on index funds, after taking sales loads, expenses, trading costs, and taxes into

account. The previous article explored three reasons this is the case. First, most active managers' forecasting skills probably aren't consistently accurate over time. Sometimes this skill depends on access to private information, the volume and quality of which usually varies over time. In other cases, superior forecasting skill is based on a unique model for making sense of the information available to the manager. However, models can be copied, or made obsolete by changes in the underlying economic processes which generate returns on financial assets.

The second reason it is hard for active managers to deliver better returns than index funds over the long term is that the former often face significant constraints on their ability to translate their forecasts into portfolio allocations. The third reason for most active managers' underperformance is the fact that they trade more often than index funds do, which not only increases their transaction costs, but also generates capital gains on which their investors have to pay taxes.

There is, however, one more reason why active managers find it difficult to consistently pick winning stocks and thereby deliver higher returns than index funds. Simply put, from a C.E.O.'s perspective, it is very, very hard to be a winning stock in the first place. Or, to put it differently, it is very hard for a company to consistently deliver returns that are significantly above those of its peers. This month, we'll take a closer look at why this is the case.

Broadly speaking, superior shareholder returns can result from two factors: either above average corporate performance, and/or above average investor enthusiasm. We'll leave the latter until the end of this article, and concentrate on the former. In theory, it results from creating and efficiently implementing a strategy that is superior to those of competing firms in your industry. Let's look at the challenges that are involved in practice.

A strategy is, in essence, an organization's answers to a series of interrelated questions. To simplify matters, assume that these answers are of the simple "yes or no" variety. Further assume that these decisions cover, at minimum, marketing, production, finance, human resources, systems, mergers and acquisitions and research and development issues. Assuming just five decisions in each policy area (an unrealistically low number, given our yes/no format), the company's challenge is to choose a strategy from 2^{35} possible options. And this probably underestimates the number of choices available, as in the real world the

total number of variables that define the landscape containing all possible strategies open to a manager can never be known with certainty.

Moreover, the shareholder value created by a strategy usually depends, to varying degrees, on the extent to which the decisions made in various areas are consistent with each other. For example, a marketing decision to offer consumers low cost products would not be consistent with a decision to use a high cost precision manufacturing technology, or to locate that operation in a place with relatively high costs. Strategies in which many decisions are related to each other have a big advantage: they are hard for competitors to copy (because they are often hard to understand, and harder still to implement completely). However, they also have a big disadvantage: they are hard to change, because so many people have to agree to do things differently. As long as the environment in which the company competes remains unchanged, highly integrated strategies can produce high shareholder returns. But if competitive conditions change, the organizational rigidity they create can also cause those returns to quickly fall.

Assuming that the company wants to employ even a modestly integrated strategy (e.g., one with four or more interrelated decisions), the estimation of potential shareholder value from different strategies is highly nonlinear (computationally), and not susceptible to optimization using any type of mathematical model. In short, when defining their strategy (for which searching for the highest possible peak in a rugged landscape while starting out in one of its valleys seems a very good metaphor), managers cannot and should not spend their time trying to identify the "optimal" strategy, and instead typically look for one that is "good enough" and then try to improve it over time to avoid being taken over or going out of business.

Broadly speaking, there are two approaches to achieving these improvements. The first can be called "doing things better." It basically focuses on learning by doing, and making incremental improvements to existing strategies. It is the most popular approach when a firm's performance is judged by its managers to be in the acceptable range. The problem with this approach is that each successive improvement tends to double the amount of time required to achieve another one of the same magnitude (for more on this subject, see the book The Origins of Order, by Stuart Kauffman).

The second approach, which tends to be undertaken when performance has been judged unacceptable (which, we should note, is sometimes due to the expectations created by irrational investors running up the company's stock price), can be described as "doing different things", or attempting to change many strategic decisions all at once. The problem with this approach, as any experienced manager can tell you, is the high level of risk involved. For example, the majority of acquisitions fail to create value for the acquiring company, and many new product, process and business development projects fall well short of their expected results. One or two causes account for these problems. Either the potential consequences of the different strategic options weren't accurately understood (resulting in one or more incorrect decisions), or, even if the correct decisions were made, they were imperfectly implemented.

By this time, it should be apparent why few companies are able to deliver above average returns year after year. But wait – it gets worse.

Before we assumed that the shareholder value created by a set of strategy decisions only depended on a subset of the decisions themselves (e.g., the value created by a marketing decision depends not only on other marketing decisions, but also on decisions taken in manufacturing and finance). In reality, the value created by an internal decision also depends on decisions made by external parties to which the company is linked, which could include customers, competitors, suppliers and/or regulators. For example, adding new functions to a product (e.g., web browsing, email, and a digital camera to a mobile phone) will only create additional shareholder value if customers' needs have also changed in this direction, and if competitors' have not yet introduced superior offerings, and if suppliers' prices for critical components aren't higher than expected.

In other words, not only is it impossible (except by luck) to identify an optimal strategy in advance, but the landscape of possibilities itself is also constantly changing.

Finally, most companies face significant organizational limitations on their ability to navigate this landscape. As individuals, corporate leaders are subject to the same cognitive limitations as the rest of us. It takes more information to change their views than it does to initially form them. And once formed, leaders' existing views affect the information they pay attention to, and the weight they give to it, particularly when those views have generated above average results. Under these circumstances, information that conflicts with current

views is often either overlooked or not taken as seriously as it should be. Moreover, these individual cognitive shortcomings are usually reinforced by group processes which promote conformity and discourage dissent. In many cases, all these factors cause companies to become overconfident, miss important changes in their environment, make poor investments, and eventually to end up on the receiving end of nasty business surprises that lead to substantial performance declines.

Across a wide range of companies and industries, these dynamics produce recurring patterns in shareholder returns. Successful innovation generates high returns and overconfidence. Competition soon follows, prompting increased focus on cutting costs by doing things better. In many cases, resistance to significant change leaves an opening that is exploited by a newer, more innovative competitor. This produces a sharp decline in performance, which triggers a major strategic change. The ones we typically read about are the few that succeed. Most, however, fail, and yesterday's top performers disappear from view (anybody who doesn't believe this need only look at how many of the firms in any large company equity index change from decade to decade).

Consequently, very, very few companies consistently deliver above average shareholder returns over long periods of time. For example, in their paper "Industry, Corporate, and Segment Effects and Business Performance", Ruefli and Wiggins found that between 1984 and 1996 only 18.8% of a sample of almost 5,000 companies were able to deliver superior performance (as measured by return on assets) in any given year, while only 3.5% were able to deliver superior performance in every year. Similarly, in their paper "Is Performance Driven by Industry or Firm Specific Factors?" Hawawini, Subramanian, and Verdin found that for most firms, industry factors had the biggest impact on performance. Only a few firms were either well above or well below the industry average. While most firms are able to institute the basic management practices needed to avoid poor performance, they find it exceptionally difficult to identify and successfully implement a set of strategic choices that will deliver sustained superior performance. A simulation based study by Philip Auerswald from Harvard ("The Complexity of Production, Technological Volatility, and Inter-Industry Differences in the Persistence of Profits Above the Norm") covers the same issues, and reaches a similar conclusion. Finally, the existence of these dynamics also makes it very hard to predict a company's future rate of growth (which, for most firms, is the major

determinant of market value and shareholder returns). As Chan, Karceski, and Lakonishok document (in their paper "The Level and Persistence of Growth Rates"), company growth rates tend not to persist over time beyond what would be expected due to luck alone.

As you can see, the challenges facing a C.E.O. who is trying to consistently deliver superior shareholder returns are extremely daunting, and the probabilities of long term success are very low. By definition, the probability that an active investment manager will be able to consistently identify these few winning firms in advance therefore must be lower still – the probability of corporate success must be higher than the joint probability of corporate success and accurate foresight by our investment manager. Seen in this light, the heavy dependence on momentum (that is, on accurately forecasting other investors' behavior) by successful active managers should come as no surprise. As Professor Russ Wermers concluded in his recent paper "Is Money Smart?" the "finding that [active fund] performance does not seem to persist after controlling for cash inflows [and the momentum buying they finance] casts doubt on previous studies that found that [active] managers have talents in choosing stocks that beat their benchmarks." So, once again, we reach a now familiar conclusion: it is extraordinarily difficult for most mortals to be consistently successful at active investment management over long periods of time, and the great majority of investors would be much better off investing in a diversified portfolio of index funds.

Model Portfolios Year-to-Date Nominal Returns

We offer over 2,000 model portfolio solutions for subscribers whose functional currencies (that is, the currency in which their target income and bequest/savings are denominated) include Australian, Canadian, and U.S. Dollars, Euro, Yen, and Pounds-Sterling. In addition to currency, each solution is based on input values for three other variables:

1. The target annual income an investor wants her or his portfolio to produce, expressed as a percentage of the starting capital. There are eight options for this input, ranging from 3 to 10 percent.
2. The investor's desired savings and/or bequest goal. This is defined as the multiple of starting capital that one wants to end up with at the end of the chosen expected life. There are five options for this input, ranging from zero (effectively equivalent to converting one's starting capital into a self-managed annuity) to two.
3. The investor's expected remaining years of life. There are nine possible values for this input, ranging from 10 to 50 years.

We use a simulation optimization process to produce our model portfolio solutions. A detailed explanation of this methodology can be found on our website. To briefly summarize its key points, in order to limit the impact of estimation error, our assumptions about future asset class rates of return, risk, and correlation are based on a combination of historical data (from 1971 to 2002) and the outputs of a forward looking asset pricing model. For the same reason, we also constrain the maximum weight that can be given to certain asset classes in a portfolio. These maximums include 20% for foreign bonds and foreign equities, and 10% each for commercial property, commodities, and emerging markets equities. There are no limits on the weight that can be given to real return and domestic bonds, and to domestic equities.

Each model portfolio solution includes the following information: (a) The minimum real (after inflation) compound annual rate of return the portfolio must earn in order to achieve the specified income and savings/bequest objectives over the specified expected lifetime. (b) The long-term asset allocation strategy that will maximize the probability of achieving this return,

given our assumptions and constraints. (c) The recommended rebalancing strategy for the portfolio. And (d) The probability that the solution will achieve the specified income and savings/bequest goals over the specified time frame.

The following table shows how asset allocations with different target compound annual rate of return objectives have performed year-to-date:

	YTD 30Jan04	**Weight**	**Weighted Return**
	In US$		In US$
7% Target Real Return	*YTD Returns are Nominal*		
Asset Classes			
Real Return Bonds	1.2%	0%	**0.0%**
U.S. Bonds	0.8%	0%	**0.0%**
Non-U.S. Bonds	0.1%	20%	**0.0%**
Commercial Property	4.3%	10%	**0.4%**
Commodities	2.7%	10%	**0.3%**
U.S. Equity	2.2%	50%	**1.1%**
Foreign Equity (EAFE)	1.4%	0%	**0.0%**
Emerging Mkt. Equity	3.1%	10%	**0.3%**
		100%	**2.1%**

.

	YTD 30Jan04	**Weight**	**Weighted Return**
	In US$		In US$
6% Target Real Return	*YTD Returns are Nominal*		
Asset Classes			
Real Return Bonds	1.2%	0%	**0.0%**
U.S. Bonds	0.8%	0%	**0.0%**
Non-U.S. Bonds	0.1%	20%	**0.0%**
Commercial Property	4.3%	10%	**0.4%**
Commodities	2.7%	10%	**0.3%**
U.S. Equity	2.2%	45%	**1.0%**
Foreign Equity (EAFE)	1.4%	5%	**0.1%**
Emerging Mkt. Equity	3.1%	10%	**0.3%**
		100%	**2.1%**

.

	YTD 30Jan04	Weight	Weighted Return
	In US$		In US$
5% Target Real Return	*YTD Returns are Nominal*		
Asset Classes			
Real Return Bonds	1.2%	0%	**0.0%**
U.S. Bonds	0.8%	0%	**0.0%**
Non-U.S. Bonds	0.1%	20%	**0.0%**
Commercial Property	4.3%	10%	**0.4%**
Commodities	2.7%	10%	**0.3%**
U.S. Equity	2.2%	30%	**0.7%**
Foreign Equity (EAFE)	1.4%	20%	**0.3%**
Emerging Mkt. Equity	3.1%	10%	**0.3%**
		100%	**2.0%**

.

	YTD 30Jan04	Weight	Weighted Return
	In US$		In US$
4% Target Real Return	*YTD Returns are Nominal*		
Asset Classes			
Real Return Bonds	1.2%	5%	**0.1%**
U.S. Bonds	0.8%	35%	**0.3%**
Non-U.S. Bonds	0.1%	20%	**0.0%**
Commercial Property	4.3%	10%	**0.4%**
Commodities	2.7%	10%	**0.3%**
U.S. Equity	2.2%	5%	**0.1%**
Foreign Equity (EAFE)	1.4%	10%	**0.1%**
Emerging Mkt. Equity	3.1%	5%	**0.2%**
		100%	**1.5%**

.

	YTD 30Jan04	Weight	Weighted Return
	In US$		In US$
3% Target Real Return	*YTD Returns are Nominal*		
Asset Classes			
Real Return Bonds	1.2%	75%	**0.9%**
U.S. Bonds	0.8%	0%	**0.0%**
Non-U.S. Bonds	0.1%	10%	**0.0%**
Commercial Property	4.3%	10%	**0.4%**
Commodities	2.7%	5%	**0.1%**
U.S. Equity	2.2%	0%	**0.0%**
Foreign Equity (EAFE)	1.4%	0%	**0.0%**
Emerging Mkt. Equity	3.1%	0%	**0.0%**
		100%	**1.5%**

.

	YTD 30Jan04	Weight	Weighted Return
	In US$		In US$
2% Target Real Return	*YTD Returns are Nominal*		
Asset Classes			
Real Return Bonds	1.2%	85%	**1.0%**
U.S. Bonds	0.8%	0%	**0.0%**
Non-U.S. Bonds	0.1%	10%	**0.0%**
Commercial Property	4.3%	5%	**0.2%**
Commodities	2.7%	0%	**0.0%**
U.S. Equity	2.2%	0%	**0.0%**
Foreign Equity (EAFE)	1.4%	0%	**0.0%**
Emerging Mkt. Equity	3.1%	0%	**0.0%**
		100%	**1.2%**