

The Index Investor

Why Pay More for Less?

Global Asset Class Returns

<i>YTD 31Oct03</i>	<u>In USD</u>	<u>In AUD</u>	<u>In CAD</u>	<u>In EURO</u>	<u>In JPY</u>	<u>In GBP</u>
US Equity	21.90%	-4.22%	2.67%	11.58%	13.87%	16.64%
US Bonds	2.80%	-23.32%	-16.43%	-7.52%	-5.23%	-2.46%
AUS Equity	39.90%	13.78%	20.67%	29.58%	31.87%	34.64%
AUS Bonds	20.02%	-6.10%	0.79%	9.70%	11.99%	14.76%
CAN Equity	40.00%	13.88%	20.77%	29.68%	31.97%	34.74%
CAN Bonds	24.64%	-1.48%	5.41%	14.32%	16.61%	19.38%
Euroland Equity	24.00%	-2.12%	4.77%	13.68%	15.97%	18.74%
Euroland Bonds	13.85%	-12.27%	-5.38%	3.53%	5.82%	8.59%
Japan Equity	32.90%	6.78%	13.67%	22.58%	24.87%	27.64%
Japan Bonds	6.31%	-19.81%	-12.92%	-4.01%	-1.72%	1.05%
UK Equity	19.10%	-7.02%	-0.13%	8.78%	11.07%	13.84%
UK Bonds	5.13%	-20.99%	-14.10%	-5.19%	-2.90%	-0.13%
World Equity	24.70%	-1.42%	5.47%	14.38%	16.67%	19.44%
World Bonds	7.00%	-19.12%	-12.23%	-3.32%	-1.03%	1.74%
Commodities	12.30%	-13.82%	-6.93%	1.98%	4.27%	7.04%
XR Chng v. USD	0.00%	26.12%	19.23%	10.32%	8.03%	5.26%

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 October, this benchmark had returned 18.1%, while our model portfolio had

returned 22.9%. 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 21.2%.

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 14.3%, while our model portfolio had returned 19.2%, and the global benchmark had returned 17.6%.

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 6.6%, while our model portfolio had returned 9.5% and the global benchmark 10.5%.

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 18.1%, while our model portfolio had returned 20.6%. 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 21.2%.

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 14.3%, while our model portfolio had returned 15.8%, and the global benchmark had returned 17.6%.

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 6.6%, while our model portfolio had returned 9.6% and the global benchmark 10.5%.

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 22.8% year-to-date, our 10%

target return portfolio has returned 19.5% our 8% target return portfolio has returned 16.0%, and our 6% target return portfolio has returned 9.5%.

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 15.9%.

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	3.32%	4.00%	7.32%	4.90%	3.80%	8.70%
Canada	2.94%	4.00%	6.94%	2.10%	2.00%	4.10%
Eurozone	1.77%	4.00%	5.77%	2.50%	2.60%	5.10%
Japan	2.07%	4.00%	6.07%	2.70%	1.00%	3.70%
U.K.	2.21%	4.00%	6.21%	2.50%	3.20%	5.70%
U.S.A.	2.51%	4.00%	6.51%	4.50%	1.60%	6.10%

**High Productivity Growth Scenario. See Asset Class Review, in our June 2003 Issue, for details of both productivity growth scenarios for each region.*

*** When expected return is greater than required return, theoretical index value will be greater than actual index value.*

Country	Implied Index Value*	Current Index Value	Current/Implied (high productivity growth)	Current/Implied (low productivity growth)
Australia	436.48	277.97	64%	90%
Canada	118.79	287.46	242%	292%
Eurozone	111.82	140.63	126%	183%
Japan	28.17	102.34	363%	474%
U.K.	254.53	295.09	116%	163%
U.S.A.	341.37	428.84	126%	188%

* *High productivity growth scenario.*

The Case for Active Management

Here at The Index 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 manager. 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 3% alpha you want. To get that 3% 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, their 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 breadth 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 too 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% ($5\% \times .75$) 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 ($2 \times .78$). 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. Equity	.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%
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 managers 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 managers 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.

Portfolio Tilts and Home Bias

The previous article provides a perfect introduction to an important asset allocation question: should you invest in broad asset classes, or should you "tilt" your investment toward one or more subgroups within a given asset class in order to earn higher return or take on less risk? To answer this question, you have to ask two further ones. First, does the benefit you expect to receive arise from rational or irrational causes? For example, does the extra return from a subgroup result from additional risk (a rational explanation) or from systematic investor decision making mistakes (a behavioral explanation). And if you believe the latter to be the case, why do you expect this behavior to remain unchanged? In general, we think the extra returns from the tilts people typically consider represent rational compensation for additional risks. Let's look at the bond and equity markets to see what we mean.

In the bond market, people often consider two tilts. The first is based on the average maturity of the bonds in a given fund (technically, these tilts are based on average duration, which we will somewhat inaccurately call maturity for simplicity's sake). Longer maturity bonds are much more sensitive to interest rates, and so command a higher rate of return than short term bonds. Most broad investment grade bond indexes have intermediate maturities, which historically have usually provided the best risk/return trade-off. The other bond market tilt that people often consider is between bonds with high creditworthiness (e.g., those issued by governments) and bonds with low creditworthiness. As in the case of maturity tilts, bonds with lower credit ratings pay higher rates of return for a good reason: they have a higher risk of defaulting. Again, most broad investment grade bond indexes tend to be in between these two extremes, with a mixture of government and investment grade corporate bonds or similarly high rated securities (e.g., mortgage backed bonds). A final bond market tilt that one can take in some countries is toward bonds whose interest payments are exempt from some or all taxes. For example, in the United States these are called "municipal bonds" and are issued by states, cities, and other local government units. Frequently, the yield on these bonds (when converted to a taxable equivalent) is higher than that on an equivalently rated taxable bond. Is this a free lunch? I'm afraid not. The usual cause for this discrepancy is the markedly lower size of most municipal bond issues, which necessitates the payment of higher yields to compensate investors for the additional illiquidity risk they are taking on.

Moving on to the equity market, let's look at three possible tilts. The first one that people often consider is toward "growth" or "value" funds. Historically, the data show that higher long term returns have been earned on value funds in many different markets around the world. Why this is so continues to be a source of much discussion (and disagreement) in financial circles. The rational school of thought believes the value premium to be compensation for bearing additional risk, in that the returns on value companies are thought to be more sensitive to the business cycle. In contrast, the other school of thought believes the value premium reflects some very difficult to change aspects of investor behavior, such as under and over-reaction to news (see "Contrarian Investment, Extrapolation, and Risk" by Lakonishok, Shleifer and Vishny), or mental accounting a loss aversion (see "Mental Accounting, Loss Aversion and Individual Stock Returns" by Barberis and Huang).

The second equity market tilt that people often consider is toward small capitalization stocks. However, in this case the evidence in favor of the existence of a consistent “small cap stock premium” is much less convincing than the evidence backing the value premium. For example, in their paper “On the Robustness of Size and Book to Market in Cross Sectional Regressions”, Knez and Ready show how the removal of a few extreme cases from a typical sample reverses the size premium, so that smaller companies earn lower rather than higher average returns than large companies. In “Robust Measurement of Size and Book to Market Premia”, Chou and Hsu reach the same conclusion. Finally, in a recent research newsletter, Lipper (a mutual fund research service) concluded that “the size premium is sporadic and not something that can be reliably priced.”

So, to sum up what we’ve said so far: in the bond market, we can understand tax driven tilts, but prefer to avoid maturity and credit driven tilts. In the equity market, we recognize the long term higher returns that a value tilt is likely to generate (assuming history is a valid guide to the future), but raise two cautions: (a) a value tilt also involves additional economic risks; as well as (b) additional risk of regret, as we saw in the nineties when value focused investors suffered through many years of returns that lagged behind their peers. As for size based tilts, we don’t believe in taking them, as the additional return seems too uncertain, while the risks seem quite real (small firms are more vulnerable to failure during economic slowdowns).

A third tilt that some people will consider in their equity portfolios is toward a different mix of industry sector weightings than the one contained in the broad market index. Again, while this can be an effective way to adjust the risk and return profile of one’s equity holdings (e.g., to limit exposure to the industry which already generates your labor income), it seems just as unlikely as other asset class tilts to permanently offer “something for nothing”. As always, taking on more risk tends to generate higher returns, and vice versa. However, there is also another dimension to sector tilts that can make them even more complicated.

The underlying question is the extent to which the variation in stock returns across countries is being driven by industry factors compared to country factors. Traditionally, the latter have

been viewed as the more important influence, but in recent years a number of research studies found that industry factors had equaled or exceeded country factors in relative importance. These studies usually concluded that rather than initially dividing equities into different geographic groupings, more effective asset allocation processes would henceforth either begin with global industry weights (as shown in the table below), or optimize them simultaneously with geography.

Industry Sector	Approximate Weight in MSCI World Index, Spring, 2003
Financial Services	22%
Healthcare	14%
Consumer Cyclicals	13%
Information Technology	11%
Industrials	10%
Consumer Staples	9%
Energy	8%
Telecommunications	5%
Basic Materials	4%
Utilities	4%

More recently, however, a new study from the International Monetary Fund (“International Stock Returns and Market Integration” by Brooks and Del Negro) has found that the sharp increase in the importance of global industry factors was heavily influenced by the technology bubble. After eliminating this effect, they find that countries are still more important than industries, except in Europe where, following the launch of the Euro, industries have become more important factors than countries in determining stock returns.

The continued importance of countries in the determination of equity returns raises another very important issue: home bias. As shown in the following table, most investors are already making a sizeable tilt in their equity portfolios:

Country	Domestic Equity Market Capitalization as a Percentage of the Global Equity Market in 2000	Percent of Total Equity Allocation to Domestic Equity in 2000
Australia	2%	89%
Canada	3%	88%
Germany	4%	80%
France	4%	83%
United Kingdom	8%	78%
Japan	13%	91%
United States	48%	89%

Source: "Equity Home Bias: Can Information Cost Explain the Puzzle?" by Karsten Jeske, Federal Reserve Bank of Atlanta

A number of alternative explanations have been offered for the size of these tilts. As with so many other financial markets analyses, they divide into rational and behavior explanations. Historically, cross border capital controls provide a good rational explanation for these extreme home country tilts. Unfortunately, even after they were removed, the tilts remained. A subsequent set of analyses examined the hypothesis that higher foreign transaction costs were to blame. This explanation has also been debunked (see "Home Bias and High Turnover Reconsidered" by Francis Warnock). A third rational explanation was that higher information acquisition costs in foreign markets was the cause of the large home market tilts. However, as Jeske recently showed (see "Equity Home Bias: Can Information Costs Explain the Puzzle?"), the implicit additional costs above domestic information related costs (as a percentage of asset value) needed to justify the observed home bias ranged from a low of 1.48% in the United States to 9.79% in Japan. All of these are far above the operating expenses charged on exchange traded fund which allow domestic investors to easily invest in foreign equity markets.

The lack of plausible rational explanations for the observed home market tilt has led other researchers to explore potential behavioral explanations. Three different ones have been offered. The first is that investors prefer local equities because information on them is relatively more available, recent, vivid and familiar than information on foreign equities. This

results in a substantial overestimation of the relative riskiness of foreign equities (see “The Local Bias of Individual Investors” by Zhu Ning, and “Confidence in the Familiar” by Kai Li).

The second behavior explanation that has been offered is that domestic investors are unrealistically and excessively optimistic about the future prospects for home country shares (see “Understanding the Equity Home Bias” by Strong and Xu, and “The Equity Home Bias: Contrasting an Institutional With a Behavioral Explanation” by Fellner and Maciejovsky).

A final explanation spans the rational and behavioral camps. It notes that investors not only have an aversion to risk, but also to uncertainty (e.g., in alpha estimation model parameters) and to regret (e.g., the feeling one would have if one had invested substantially more in foreign equities than one’s domestic peers, and then seen that investment underperform the domestic benchmark). While foreign and domestic equities may have roughly equal risks (and in fact, may be highly complimentary in a quantitative asset allocation model), they are perceived by domestic investors to be more uncertain and to have greater potential to cause regret. Because of these two latter concerns, only very small percentages of a portfolio are invested in them.

Given the lack of supporting evidence for the rational hypotheses that have been put forward to explain home country bias in equities, we tend to believe that the causal factors are the behavioral ones. Given that, investors who can overcome them seem likely to enjoy a better long term risk/return tradeoff than their excessively home bound counterparts.

And On to Foreign Currency Bonds (Again)

A recent study (“Diversification, Original Sin, and International Bond Portfolios” by Berger and Warnock of the Federal Reserve) has noted that the home bias in bond portfolios is often even more severe than in the case of equities. Consider the following shares of foreign bonds in the Global Market Portfolio (i.e., the market capitalization weighted portfolio that contains all the world’s liquid bonds and equities), as seen from different countries’ perspective:

Country or Region	Share of Foreign Bonds in Global Market Portfolio
Australia	56%
Canada	56%
Eurozone	45%
Japan	47%
United Kingdom	55%
United States	28%

While one suspects that the same set of behavioral factors contributes to home bias in both the equity and bond markets, the latter may have one additional consideration at work: the currency risk hedging issue. In the case of foreign equities, because the future value and holding period is relatively hard to predict, foreign exchange risk hedging is done much less frequently than it is with foreign bonds. This seems to be something of a blessing, because hedging seems to create no end of confusion.

Let's start with the arguments in favor of including foreign currency bonds in one's portfolio. Three of them stand out.

First, because local currency returns on foreign bonds have a low correlation with the home currency returns on domestic bonds, there is a potential diversification/risk reduction benefit from holding them. In fact, a recent study ("Credit Risk Diversification: Evidence From the Eurobond Market" by Simone Varotto) found that international diversification caused a bigger reduction in portfolio credit risk than domestic diversification across industries, maturities, or credit ratings.

Second, foreign currency bonds are attractive because of the way their correlation with other asset classes tends to vary over time. Specifically, the correlation between domestic equities and foreign currency bonds tends to be lowest during periods of equity market distress (due to

the flight to quality phenomenon). This is exactly the opposite of what happens to the correlation between domestic and foreign equities during periods of distress. In other words, foreign currency bonds tend to provide diversification benefits when they are most needed, while foreign equities tend to do the opposite. (This point is elaborated on at length in the recent paper "Asymmetric Dynamics in the Correlation of Global Equity and Bond Returns" by Cappiello, Engle, and Sheppart).

Third, foreign currency bonds provide a natural inflation hedge, because an increase in domestic inflation relative to the rest of the world will result in a depreciation of one's home currency and an appreciation in the value of foreign currency denominated assets. Of course, other asset classes also provide a similar benefit, including real return bonds, and, to a lesser extent, commodities and property.

Unfortunately, the currency hedging issue seems to overshadow these arguments. The key issues involved include (a) whether or not an investor should hedge the exchange risk on his or her foreign currency bonds, and, if so, (b) how extensive that hedging needs to be. On the one hand, there is an argument that completely hedging foreign exchange risk (at a cost of between .25% and .50% per year) reduces the volatility of the foreign bonds asset class much more than its return, and, in the findings from the Berger and Warnock study, results in "the optimally allocated bond portfolio [assuming a one year holding period] being comprised almost entirely of non-U.S. bonds."

However, there are also strong counterarguments to this view. In "Currency Hedging Over Long Horizons", Kenneth Froot finds that "the ability of hedging to reduce portfolio risk holds at short horizons, but not at long horizons." As he explains it, "at short horizons, exchange rate changes in excess of the forward rate [that is, the rate at which the hedge is entered into when the foreign currency bond is purchased] average about zero, and have virtually no correlation with any variable, including local currency asset returns. Given this, taking currency risk offers no additional reward, and therefore should be hedged away...However, currency hedges have very different properties at long horizons compared to short horizons. The data show that while over short horizons hedging reduces risk

substantially, over long horizons hedging often does not reduce risk at all. In fact, at long horizons, many fully hedged international investments actually have greater return variance [i.e., volatility or risk] than their unhedged counterparts.”

“Hedge returns at different horizons are driven by very different factors. At relatively short horizons, hedge returns are dominated by changes in real exchange rates (i.e., the purchasing power of one currency compared with another). However, mean reversion in real exchange rates implies that these purchasing powers tend towards parity, so that real exchange rates over time remain roughly constant. At long horizons, hedge returns are instead dominated by fluctuations in cross country differences in expected inflation and real interest rate differentials.” Given this, “the optimal portfolio hedging strategy will depend on the investment horizon...[and]... investors with longer horizons may want to hedge much less than 100% of their exposure.”

Similarly, Campbell, Viceira and White (in their paper “Foreign Currency and Long Term Investors”) note that the “conventional wisdom holds that investors should avoid exposure to foreign currency risk.” They go on to argue that “the conventional wisdom may be wrong for long term investors.” They note that any domestic bond with a maturity (duration) that is shorter than the investor’s horizon is risky, because it must be rolled over (that is, reinvested) at an uncertain future rate. This risk can be hedged by holding foreign currency bonds if the domestic currency tends to depreciate when the domestic real interest rate falls, as implied by the theory of Uncovered Interest Rate Parity. They go on to note that while uncovered interest rate parity doesn’t work well in the short run (as a forecaster of future exchange rates), recent studies find support for it over longer periods. As a result, the authors find much higher foreign bond holdings for long term as opposed to short term investors.

As our model portfolios have a twenty year horizon, we have included unhedged foreign currency bonds as an asset class, rather than hedged foreign currency bonds.

That being said, we are still faced with the question as to whether or not foreign currency bonds should get the weight that they did in the model portfolio optimizations that were done

using both historical and anticipated future returns. There are three points with respect to our specific assumptions where we recognize that reasonable people can disagree.

First, our historical data series for foreign currency bonds was based on the combination of the Salomon Brothers World Government Bond Index and different IMF Long Term Government Bond indexes (for those years before the SB index existed, combined using current SB weights, which is admittedly an approximation). As the SB index is an intermediate maturity (duration) index, we arguably could have used an average of the IMF money market and long term government bond yields to extend the SB data series back to 1971 (unfortunately, the IMF doesn't have an intermediate term series). However, this would, at most, have changed our historical average real yields by less than one percent, which would not have had a big impact on our optimization results.

Second, the correlations we used were based on the full 1971-2002 historical sample period. Using a more recent period would have raised the correlation between the returns on foreign currency bonds and other asset classes, but also would have excluded data from periods of time during which the bond and foreign exchange markets were much less benign than they have been in recent years. We found, however, that this had no impact on our optimization results – even the more recent correlations (e.g., from 1988-2002, which were significantly higher than the ones from 71-02) were still relatively low enough to produce high portfolio weights for foreign bonds.

This brings us to the third point, which is our future return assumptions for foreign currency bonds. Our estimate was based on a simple average of two different estimation approaches: (a) the historical spread over domestic bond yields and (b) the combination of current real yields in different countries, plus the future exchange rate changes implied by the difference in current nominal ten year government bond yields, weighted by countries' relative weight in the SB World Government Bond Index (i.e., the "uncovered interest parity" or UIP approach). Arguably, we could have stuck with the latter approach, which is theoretically more defensible (at least at longer horizons) and which would have produced lower return estimates for foreign currency bonds. On the other hand, some readers have argued that the difference

in ten year yields at the time we wrote the article did not fully reflect the future exchange rate disruptions that may be associated with the unwinding of the U.S. current account deficit. In short, there seems to be no easy answer to the questions raised by foreign bonds.

In light of these concerns, and in the interest of conservatism (granted, a dying virtue in the world of finance), we have recalculated our “forward looking” model portfolios using return assumptions for foreign bonds and foreign equities that are solely based on the use of the Uncovered Interest Parity to estimate future exchange rate changes. The following tables present these new model portfolios, along with model portfolios based on historical data. In both cases, we limited foreign currency bonds and equity to a maximum of forty percent each of a model portfolio. As you can see, the use of the UIP approach reduces the expected return on these asset classes to the point that the forty percent constraint is no longer binding in most of the forward looking model portfolios. Finally, as we did the last time, we have combined these two portfolios (that is, the one based on historical data and the one based on forward looking assumptions) by giving a weight of .67 to the former and .33 to the latter. The resulting portfolios are the ones we will be using in next year’s calculations, starting on January 1st. We will also be updating our “benchmark relative” portfolios next month. However, as this is much less computationally intensive than the update of our target return portfolios, we wanted to finish the latter before moving on to the former.

	3% Target Return (Historical Return Assumptions)	3% Target Return (Future Return Assumptions)	3% Target Return (Combined Weights)
Real Return Bonds	45%	30%	40%
Domestic Bonds	25%	25%	25%
Foreign Bonds	10%	5%	8%
Property	10%	5%	8%
Commodities	5%	10%	7%
Domestic Equity	0%	20%	7%
Foreign Equity	5%	0%	3%
EM Equity	0%	5%	2%
Total	100%	100%	100%
Expected Average Annual Return	4.5%	4.6%	
Standard Deviation of Expected Returns	2.6%	5.0%	
Probability of Achieving Target	98.0%	89.0%	

	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%	5%	15%
Foreign Bonds	25%	15%	22%
Property	20%	0%	13%
Commodities	0%	20%	7%
Domestic Equity	20%	40%	27%
Foreign Equity	5%	5%	5%
EM Equity	10%	10%	10%
Total	100%	100%	100%
Expected Average Annual Return	7.5%	6.0%	
Standard Deviation	5.5%	9.8%	
Target Probability	91.0%	61.0%	

	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%	5%	28%
Property	15%	0%	10%
Commodities	10%	20%	13%
Domestic Equity	5%	65%	25%
Foreign Equity	0%	0%	0%
EM Equity	20%	10%	17%
Total	100%	100%	100%
Expected Average Annual Return	8.4%	7.1%	
Standard Deviation	7.0%	12.4%	
Target Probability	74.0%	36.0%	

Asset Allocation Comparisons

As a final check in our model portfolio construction process, we have collected information about the allocations being used by a number of institutions which have been widely recognized for their investment management skills. These include Harvard, Yale, and Stanford universities, the Ontario Teachers Pension Fund, and the California Public Employees Retirement System (“CALPERS”). Our assumption is that all of them are aiming to generate annual real returns of between 5% and 7%.

Asset Class	Harvard	Yale	Stanford	Ontario Teachers	CALPERS
Real Return Bonds	6%			9%	
Domestic Bonds	11%	10% All Bonds	12% All Bonds	24 % All Nominal Bonds	26% All Bonds
Foreign Bonds	5%				
Real Estate	10%	0%	16%	11%	9%
Natural Resources and Commodities	13%	20%	7%	3%	0%
Domestic Public Equities	15%	15%		21%	39%
Foreign Public Equities	10%	12.5% Includes Emerging	40% (Domestic, Foreign, Emerging)	28% Includes Emerging Markets	19% Includes Emerging Markets
Emerging Market Equities	5%				
Private Equity	12%	17.5%	10%	4% (and Hedge Funds)	7 % (and Hedge Funds)
Hedge Funds	13%	25%	15%		

As you can see, our model portfolio allocations are reasonably close to those currently being used by these five highly respected investors (after adjusting for the fact that we have not included private equity and hedge funds in our universe of asset classes, because they are not easily accessed by retail investors). To be honest, this is reassuring, and strengthens our belief that we're on the right track.