The Index Investor

Why Pay More for Less?

Model Portfolio Update

Before reviewing the performance of our model portfolio's through February, we'd like to take a minute to answer a question that a number of readers have asked us about what we mean when we refer to an investor's "functional currency." As we use the term, your functional currency is the currency in which your future liabilities are denominated. Here are two examples. Susan and Jim Smith live in Los Angeles, California, and are saving for their retirement. When they retire, they plan to live in North Carolina, to be close to their grandchildren. North Carolina is located in the United States, so most of their post-retirement expenses will be in U.S. dollars. Hence, their functional currency, with respect to their future retirement liability, is U.S. dollars.

Susan's cousin Rachel Knowles lives with her husband Jack in London, England. After they retire, they plan to live in Australia. Because most of their post-retirement expenses will therefore be in Australian dollars, this is their functional currency (with respect to retirement saving), even though they live today in the United Kingdom. In short, your functional currency, as we use the term, has nothing to do with where you live today; rather, it is determined by the currency in which your future liabilities (in this case, the amount of savings you will need to have the post-retirement lifestyle you desire) are denominated.

Now on to the model portfolio updates.

Our first set of model portfolios are designed to deliver returns that are superior to their respective domestic benchmarks, while taking on the same amount of risk (that is, having the same expected standard deviation of returns). Our first portfolio is benchmarked

against a mix of 80% domestic equities (as measured by the Dow Jones Total Market ETF) and 20% domestic bonds (as measured by the Vanguard Total Bond Market Fund VBMFX). Year-to-date, this benchmark is down (2.8%), while our model portfolio is down (2.9%). On the one hand, our model portfolio's performance has been hurt by its allocation to European equities, which so far this year have underperformed those in the United States. On the other hand, our performance has been helped by our allocations to Emerging Markets equities, Real Estate Investment Trusts and Commodities, which have all done well. Finally, as we have previously described, this year we are also comparing our model portfolios' performance to a set of Global Benchmarks. These are based on the same equity/debt weightings as the domestic benchmarks, but employ global rather than domestic equity and bond indexes. Given that we typically use a mix of domestic and international asset classes in our model portfolios, these global benchmarks are also a good indicator of our model portfolios' relative performance. In the case of our 80/20 portfolio, the global benchmark is down (3.3%) year to date.

The second model portfolio is benchmarked against a mix of 60% equities and 40% bonds. Year-to-date, the domestic benchmark portfolio is down (1.7%), while our model portfolio is down between (1.6%) and (1.8%), depending on the international bond fund used. As was the case last year, the two funds which actively manage their currency exposures (Fidelity International Bond FGBDX and Pimco Foreign Bond PFODX) have done better than the T. Rowe Price International Bond Fund (RPIBX) which takes a passive approach to exchange rate risk. In comparison, the global 60/40 benchmark is down (2.7%).

The third model portfolio is benchmarked against a mix of 20% equities and 80% bonds. Through February, the domestic benchmark portfolio is up 0.4%, while the model portfolio is up 0.3% and the global benchmark is down (1.6%).

Our second set of model portfolios is designed to match the returns of their respective domestic benchmarks while taking on less risk (again defined as expected standard deviation of the portfolio's returns).

The 80/20 benchmark is down (2.8%) year-to-date, while the model portfolio is down (2.8%), and the global 80/20 benchmark is down (3.3%). The 60/40 benchmark is down (1.7%) year-to-date, as is the model portfolio, while the global 60/40 benchmark is down (2.7%). Finally, the 20/80 benchmark is up 0.4%, while the model portfolio is up 0.5% and the global benchmark is down (1.6%).

Our last set of model portfolios is designed differently (for more information, see out November, 2001 issue). These portfolios assume that an investor wants to maximize the probability of achieving at least a minimum target level of return (defined on a compound, or geometric basis) over a ten year holding period, while taking on as little risk as possible. Year to date, our 12% target return portfolio has returned (2.7%), our 10% target portfolio has returned (1.3%), our 8% target portfolio has returned (0.7%), and our 6% target portfolio is up 1.6%.

Last but not least, we have also continued the "active management experiment" that we began last year. As you recall, the terms of that experiment allow us to change our portfolio's allocation between asset classes only four times per year, at the end of March, June, September, and December. For the first quarter, we are 100 percent allocated to the U.S. bond market, via the Vanguard Total Bond Market Index Fund, and our return to date is 1.5%. The benchmark against which we are comparing our performance is the Fidelity Global Balanced Fund, which so far this year has returned (1.3%). However, assuming that the good folks from Boston have been more aggressive than us in moving into U.S. equities this month, we expect the gap between our respective performances to narrow through the end of the first quarter.

Is the Average Investor a Good Asset Allocator?

As regular readers know, one of the principal reasons we started The Index Investor over four years ago was because we believed that too many people were making sub-optimal asset allocation decisions, and thereby increasing the risk that they would fail to achieve their financial goals. As the U.S. Congress begins to debate various proposals for

reforming 401(k) plans (that is, company sponsored defined contribution retirement savings plans) and Social Security (national pensions to those of you located outside the United States), this seems a good time to look again at the data that underlies this belief. In this article, we will ask, and try to answer, three important questions: (1) Has the average investor in the United States adequately diversified his or her portfolio? (2) If they have not, what are the causes of their sub-optimal behavior? And (3), what, if anything, can we do to improve the situation?

On the first question, the evidence appears quite strong that the average investor in the United States has not adequately diversified his or her portfolio. As a result, he or she is forgoing the opportunity to either earn more return for the risk being taken, or, alternatively, earn the same level of return while taking on less risk. Two widely cited examples of inadequate diversification include levels of investment in international stocks that are lower than what many experts would recommend (this is usually called the "home bias" issue), and the relatively high levels of investment of 401(k) funds in the stocks of the companies sponsoring these plans. Recently, a number of academic researchers have examined this issue in more depth.

After reviewing the extensive data on household portfolios contained in the Federal Reserve's Survey of Consumer Finances, Valery Polkovnichenko of the University of Minnesota concluded (in a working paper titled "Household Portfolio Diversification") that "many stockholders are dangerously undiversified…Many individuals invest substantially in the equities of their employers through retirement plans … Despite the existence of mutual funds, many stockholders choose to invest in undiversified portfolios of individual stocks."

Another recent paper ("Equity Portfolio Diversification" by William Goetzmann from Yale University and Alok Kumar from Cornell University) examined the portfolios of more than 40,000 equity investment accounts at a large brokerage firm over the 1991 to 1996 period. They found that more than twenty five percent of investor portfolios contained only one stock, and more than fifty percent contained three of fewer. Only five

to ten percent of portfolios (depending on the year) contained more than ten stocks. To put this in context, another recent paper ("Have Individual Stocks Become More Volatile?" by John Campbell et al in the February, 2001 Journal of Finance) found that over the 1962 to 1997 period "there has been a noticeable increase in firm level volatility relative to market volatility. As a result, the number of stocks needed to achieve a given level of diversification has increased." In the past different writers have found that a randomly selected sample of from twenty to thirty stocks were needed to approximate the volatility of the overall market (with fewer yielding a higher level of portfolio volatility, or, in other words, a riskier portfolio). As a result of increasing firm specific risk, that number is now around fifty. And only five to ten percent of the portfolios reviewed by Goetzmann and Kumar had at least ten stocks! These authors go on to speculate that "it is possible that investors who hold relatively less diversified portfolios [of individual stocks] compensate by investing in [more diversified] mutual funds." However, upon further investigation they found that "the average asset allocation to mutual funds is approximately fifteen percent of the overall portfolio…In other words, there is no evidence that investors with less diversified equity portfolios compensate by investing more in mutual funds." In short, "the vast majority of investors in our sample were under-diversified."

Perhaps prompted by increasing criticism from academics, the Investment Company Institute (the trade association for mutual fund companies) in November of last year published an extensive analysis of 401(k) plan holdings. Its data is revealing.

At year end 2000, about 42 million American workers held accounts in 327,364 plans which had total assets of 1.8 trillion dollars. The ICI database is the most comprehensive available, and covers 11.8 million participants in 35,367 plans holding $579.8 billion in assets. In plans that did not offer stock in the sponsoring company (hereafter called "company stock") as an investment option, 70.4% of plan assets were allocated to equity mutual funds. However, at plans offering company stock, the total allocation to equities was 76.4%, with 44.6% in equity funds and 31.8% in company stock.

A common feature of 401(k) plans is the provision of some matching funds by the sponsoring company, in addition to the participants' own contributions. In some companies, these matching contributions must be invested in company stock; in others, the participants themselves decide how to allocate them. In the former type plans, 26.2% of plan assets were invested in equity funds, while 52.9% was invested in company stock. Within these plans, the breakdown for the participant directed portion of the plan was 38.5% equity funds and 33.2% company stock. In contrast, in those plans that received matching company contributions but did not mandate the investment of these funds into company stock, 46.1% of plan assets were invested in equity funds, while only 22.2% of assets were invested in company stock.

Finally, it is interesting to note how people with different incomes chose to invest their 401(k) assets. At the low end of the range, people making between $20,000 and $40,000 invested 38% of their plan assets in equity funds, and 41.3% in company stock; at the other end of the spectrum, people making over $100,000 invested 46.5% of their plan assets in equity funds and only 26.4% in company stock.

Goetzmann and Kumar's study reached a similar conclusion to the ICI study: "investors in low income and non-professional categories (blue collar workers, clerical workers, and sales and service workers) held the least diversified portfolios."

Taken together, all these studies paint a compelling picture of a population of investors, both within and outside of 401(k) plans, whose portfolios have substantially lower levels of diversification than most experts would recommend. The next question we must ask is what are the possible causes of the low levels of diversification we observe?

We have identified a number of possible explanations that fall into three different categories. The main premise of the first category is that investors' behavior is completely rational. Two possible explanations of the behavior we observe fall into this class. The first is that investors realize that, given the size of their savings goal (that is, the amount of money they want to have accumulated by some date in the future), the time

remaining to the date by which they want to have reached their goal, and the amounts already saved and expected to be saved in the future, they have no choice but to make a high risk investment in the hope of earning commensurately high returns. The necessary compliment to this view is that these investors also believe that the company stock they are investing in will deliver those high returns in the years ahead. Unfortunately, the data seem to indicate that not all of the investments made in company stock are being made at those firms whose shares have been appreciating at the fastest rates. For example, in his paper titled "Excessive Extrapolation and the Allocation of 401(k) Accounts to Company Stock", Professor Shlomo Benartzi of U.C.L.A. found that while employees at the firms in his sample which had experienced the best stock price performance over the previous ten years allocated 39.7 percent of their discretionary contributions to company stock, participants at the worst performing firms still allocated 10.4 percent of their discretionary contributions to company stock. This tends to undermine our first rational explanation, although it could simply reflect the fact that many of the investors at the latter firms were expecting turnarounds to occur.

Our second rational explanation is that 401(k) investors feel compelled to allocate a substantial portion of their plan assets to company stock as an overt show of loyalty to the company that will help them to keep their jobs. The fact that the overall allocation to equity investments in plans offering company stock is higher than in plans that don't offer it could be interpreted as supportive of this view. Thus far, we have found no study which has examined this hypothesis. However, given the increase in companies' monitoring of employees' behavior in other areas (e.g., the increasing use of computer sniffing programs, drug testing, etc.), and the precarious financial position of many two income families, it certainly seems to be a hypothesis worth researching.

The main premise of the second class of explanations is that investors' high allocation of 401(k) assets to company stock is caused by information related factors. Our first information related explanation for the behavior we observe is that people are allocating so much to company stock because they simply don't know any better. In other words, they do not understand the importance of diversification, and the actions that must be

taken to successfully achieve it. The fact that allocations to company stock are higher for people earning less and in positions that one associates with lower levels of education seem to support this explanation.

An alternative information based view of what is going is based on information cascades and investor herding. In this case, plan participants interpret senior officer comments that are optimistic about the firm's prospects, and/or senior officers taking substantial amounts of compensation in the form of stock options, and/or the company's decision to allocate all of its matching contributions to company stock as being based on positive information that is unavailable to others. Under such circumstances, studies have shown that investors might rationally decide to forego undertaking their own time consuming and expensive independent analysis of a plan's investment alternatives, and simply mimic the decisions of other investors whom they believe to be better informed than themselves. Benartzi's finding that employees often "interpret the allocation of [their] employer's {401(k)] contributions [to company stock] as implicit investment advice" certainly supports this explanation. Also supportive is the finding by a number of studies that many plan participants do not deviate from their plan's "default" asset allocations when offered the opportunity to do so.

Our third category of potential explanations for the high allocations to company stock that we observe in many 401(k) programs is that investors are acting irrationally. The first explanation in this category is often referred to as "groupthink", which occurs when a group rapidly comes to a conclusion that nobody thereafter wants to question for fear of being rejected by the group. At a time when many companies, often for good business reasons, are striving to become "high commitment" organizations, it is easy to see how this might occur as an unintended side effect of such programs. Thus far, however, we have found no study that has explored this as a possible explanation for the high level of 401(k) investments in company stock.

The second explanation in this category is that plan participants are investing so much in their company's own stock because they are overconfident about its future prospects.

Benartzi's study supported this view with its finding that while the past performance of a company's shares tended to predict a plan's allocation to company stock, this allocation did not, in turn, predict future performance. Benartzi termed this "excessive extrapolation."

The third explanation for high levels of investment in company stock is that, as many studies have shown, investors are very reluctant to sell their losers -- in short, they will go to great lengths (even self-destructive ones) to avoid directly confronting the feeling of regret caused by an investment loss. While this doesn't directly explain why people initially invest their 401(k) funds in company stock, it can help to explain why they fail to reallocate their funds to other asset classes when company stock is delivering inferior investment returns.

Our third question is what can be done to improve the situation we observe. Each of our possible explanations suggests a corrective action. If investors are acting rationally, and taking on high risks because they need to earn high returns to achieve their financial goals, then we need to do more to encourage people to save more of their incomes. Up to now, most initiatives of this type have involved making saving more advantageous from a tax perspective. There is a limit, however, as to how much more savings you can induce via this route. At some point, you have to start thinking about using the tax system to make it less attractive to consume. This is the idea that has historically been behind the use of luxury taxes, and it is also the idea that is behind newer ideas such as the progressive consumption tax (the basic premise of which is the more you consume, the more tax you pay).

If, on the other hand, people simply feel compelled to invest a high proportion of their plan assets in company stock, then an effective simple solution would be to set a cap on the maximum percentage of an individual's funds that can be invested in this asset.

In the case of information based explanations, if people's behavior is caused by their being unaware of the need for diversification, and its potential benefits, then the solution

is to strengthen plan education requirements. We need to make very clear to plan participants the potentially devastating downside consequences of investing a high proportion of their financial assets in the same company that also provides their labor income. Another idea would be to mandate the offering to all plan participants of mutual funds or other investment vehicles (such as exchange traded funds) that cover a broad range of asset classes, to help ensure that adequate diversification opportunities are available.

If people's behavior is driven by an information cascade and herding, then caps on the maximum amount that can be invested in company stock seem an effective simple solution. This would also be a good way to limit the damage caused by irrational behaviors, such as group think and overconfidence.

Investors' reluctance to realize their losses presents a more thorny challenge. A cap on investment in company shares would not work, because it could be maintained even as a plan participant invested more and more money in shares whose value was rapidly falling. In this case, a better approach might be to mandate the progressive reallocation of a certain proportion of the funds invested in company stock if that stock reached certain trigger points, such as the loss of 20, 35, or 50 percent of its value over any twelve month period.

In summary, in the short term there seems to be a lot of merit to proposals that have been made to cap the maximum amount of 401(k) plan assets that can be invested in a company's own stock. There also seems to be considerable merit to further strengthening the requirement that plan sponsors provide investment education to plan participants, to ensuring that they are offered a broad range of asset classes in which to invest, and to forcing them to reduce their exposure to company stock when it substantially declines in price.

Should You Be an Active Investor?

A second fundamental reason that we started The Index Investor was because we believed that too many people were wasting their time trying to beat the market. We believed that many people could save themselves a lot of time, energy, and heartache, while increasing their chances of achieving their financial goals, if they would use indexed investment products. Recently, we looked at a new set of data to see if we needed to change our beliefs.

The data set in question is publicly available: it comes from Morningstar, and was published in Money Magazine's February, 2002 issue. It consists of 493 U.S. mutual funds that have at least ten years of historical returns data (through 2001), and which invest in either U.S. or non-U.S. equities. Forty seven funds invested in non-U.S. equities; 446 funds invested in U.S. equities. Of the latter, 253 were large cap funds (including growth, balanced, and value funds), 81 were midcap funds, 52 were small cap funds, and 60 were sector specific funds. These funds collectively had assets of $1.6 trillion at year end 2001, or about 45 percent of the total amount of assets invested in mutual funds that invest primarily in equities.

Our first question was simple: how many of these funds beat the comparable index fund over ten years, after expenses? We compared the U.S. large and midcap funds to the Vanguard S&P 500 Index Fund. The Vanguard MidCap fund hasn't been in existence for ten years, and, on top of that, the Morningstar definition of "midcap" seemed to include more than a few companies that are actually in the 500. Of the 334 funds in our comparison group, 78 beat the index over ten years after expenses, or 23.4% of the funds in our sample.

We compared the small cap funds to the Vanguard Small Cap Index Fund. Over the ten year period, 21 of them beat the index after expenses, or 40.4% of the funds in our sample.

Finally, we compared the ten European funds in our sample to the Vanguard European Index Fund. After expenses, only two actively managed funds (20% of the sample) beat the index.

To put this analysis into perspective, we need to make two further points. First, we compared the index funds to an active management all star team. The funds in our database were those that were good enough to survive for at least ten years, which is no mean feat in a business that routinely merges less successful funds into more successful ones. Second, we did not have enough data to take tax consequences into account. If we did, even fewer actively managed funds would have beaten their respective index fund competitors, because the former trade more actively, and therefor generate more taxable ordinary income and capital gains distributions for their shareholders. So our results show the actively managed funds in their most favorable light.

Our second question was about the relative importance of asset allocation versus size and style tilts. In our way of thinking, asset allocation refers to the percentage of your portfolio that is invested in different asset classes, such as U.S. and international equities. Size and style tilts refer to the way you allocate your funds within an asset class (e.g., U.S. equities) between funds that invest in large, mid, and small cap companies, and/or funds that invest in growth stocks, value stocks, or a blend of both. (Readers with good memories will remember how much we dislike that term "growth stocks" -- see our December, 2001 discussion of momentum investing for more on this).

Among the 493 actively managed funds, the average compound annual return over the ten year period was 6.5% for the funds investing in non-U.S. shares, and 12.6% for funds investing in U.S. shares. (Note that these are not weighted by the assets of each fund, because fund assets were not proportional over the ten year period). The difference of 6.1% represents the impact of asset allocation.

The table below shows the average ten year compound annual returns for funds with different combinations of size and style tilts:

	Growth	Balanced	Value	Average for Size
Large Cap	11.1%	12.4%	12.7%	12.1%
Mid Cap	11.5%	14.9%	14.7%	12.9%
Small Cap	12.0%	14.3%	14.4%	13.2%
Avg. for Style	11.3%	12.9%	13.3%	

What we found most interesting about this table was how little most size and style tilts added compared to the overall average for all actively managed funds of 12.6% (not to mention the 13.5% return for the Vanguard S&P 500 Fund and the 12.7% return for the Vanguard Small Cap Index Fund). At the level of aggregate size and style, the best improvements came from tilting toward small caps (.6%) and value (.7%). At the level of the nine possible size/style combinations, it was midcap balanced which delivered the best performance improvement, at 2.3% above the overall average, while large cap growth delivered the worst, at (1.5%) below the average. Relative to the asset allocation decision, size and style tilts were clearly less important.

Sector tilts, however, were another story. The compound annual ten year returns delivered by these actively managed funds were as follows: Technology, 20.0%; Health/Biotech, 14.4%; Financial Services, 18.5%; Natural Resources, 10.9%, and Utilities, 9.0%. At 7.4%, the difference between the return on the average Technology Sector fund and the overall return on actively managed U.S. funds was larger than the 6.1% difference between the latter and non-U.S. funds (although the difference between all the other sector fund averages and the U.S. average was less than 6.1%). These data further confirm what we wrote last year: sectors tilts have the potential to deliver greater return increments than size or style tilts (of course, because they also represent more concentration than either size or style tilts, they also carry with them higher risk, as any

heavy investor in a technology sector fund can tell you these days). That being said, you also have to remember that today you can also invest in sector index vehicles in addition to actively managed funds (unfortunately, sector index funds haven't been around for ten years, so we can't compare their performance to the active funds). Given the relatively high expenses charged by active sector funds, this is an important consideration.

Okay, so what we've found so far is that a few actively managed funds have the potential to deliver higher returns, after expenses (but maybe not after taxes) than index funds. The next logical question is how do you go about identifying future index beaters?

This is where the real problems arise. In a nutshell, we found no way to use historical data to identify, ahead of time, actively managed funds that later beat their respective index fund competitors.

We started by testing the (somewhat reasonable, we thought) assumption that a fund that charged higher expenses (to pay all those extra bright people, no doubt) should deliver higher returns. Unfortunately, what we found was that the correlation between an actively managed fund's ten year compound annual return and its expense charges was (.17). That's right, it was weakly negative, implying that higher expenses resulted in relatively worse performance. Imagine that…Hold that double mocha latte! A recent study by John Chalmers et al ("Fund Returns and Trading Expenses") drives home this point. Its premise was that the relationship between fund returns and their trading expenses (as opposed to their overall expenses) provides a more powerful test of the value of active fund management. Chalmers found that the average fund incurred trading expenses equal to 0.75% of assets, and that the level of these expenses was negatively correlated with returns. In short, "the level of [additional] returns generated by [active] fund managers' trading activity fell short of the expenses they incurred [in making those trades]."

Our next step was to see if a fund's relative performance over the first five years of our data (that is, between 1992-1996) could be used to predict its relative performance over

the next five years (1997-2001). What we found was very, very discouraging (if you're an active investor). We found that only one fund whose performance was in the top ten percent of all actively managed U.S. funds over the first five years was also in the top ten percent after the second five years. One: The Fidelity Select Electronics Fund. Only one fund that was in the first decile in the first five years was in the second decile in the second five years: the Seligman Communications and Information Fund. And only two funds that were in the top ten percent of funds during the first five years were in the third decile in the second five years: the Fidelity Select Computers and Fidelity Select Banking Funds. That's only four out of 446 funds -- less than one percent -- that managed to end up in the top 30% of funds in the second five years after having been in the top ten percent after the first five years. In short, they're not kidding when they tell you that past performance is no guarantee of future performance.

We're not the only ones who have reached this conclusion. The Financial Services Authority in the U.K. last year published a report ("Past Imperfect: The Performance of UK Equity Managed Funds") that concluded that "retail investors could not usefully exploit information on past performance." Others have reached the conclusion that, at best, while good performance does not persist, bad performance does, and can be used as a signal to sell a fund. In no case has any study reached the conclusion that past good performance can be successfully used as a way to pick future winners.

Not only does good performance by an actively managed fund in one period not predict good performance in the next one, you cannot even be sure what it means. As Stan Beckers described in his Journal of Portfolio Management article ("Manager Skill and Investor Performance"), luck alone plays a great, and usually unacknowledged role (what a surprise!). He performed a simulation experiment similar to the one we described in last month's issue. What he found was that, strictly because of luck, "some [active] managers could easily stay in business for ten years without having any skill whatsoever…The natural conclusion is that manager selection on the basis of returns data (and league rankings) is extremely hazardous…This has very sobering consequences for

the mutual fund industry, where typically the only information in the public domain is the returns data track record."

In summary, what emerges (yet again) from this analysis is confirmation of the basic beliefs that have guided us since we launched The Index Investor:

- Most investors do not spend enough time thinking about how to allocate their investments across different asset classes, despite the fact that this decision has by far the largest impact on the long term returns they will earn on their portfolios.

- Most investors are overconfident. We believe that active investment management -- trying to "beat the benchmark" through superior manager selection, market timing and/or security selection -- is exceptionally difficult, and very, very few people consistently do it profitably year after year.

- Most investors would be better off if they used low cost, tax efficient index investment vehicles (mutual funds or exchange traded funds) to implement their asset allocation decisions.