

# The Index Investor

*Why Pay More for Less?*

## Global Asset Class Returns

<i>YTD 28Nov03</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	23.30%	-5.64%	2.25%	9.10%	14.95%	16.35%
US Bonds	3.00%	-25.94%	-18.05%	-11.20%	-5.35%	-3.95%
AUS Equity	39.80%	10.86%	18.75%	25.60%	31.45%	32.85%
AUS Bonds	21.55%	-7.39%	0.50%	7.35%	13.20%	14.60%
CAN Equity	45.90%	16.96%	24.85%	31.70%	37.55%	38.95%
CAN Bonds	26.90%	-2.04%	5.85%	12.70%	18.55%	19.95%
Euroland Equity	31.10%	2.16%	10.05%	16.90%	22.75%	24.15%
Euroland Bonds	16.96%	-11.98%	-4.09%	2.76%	8.61%	10.01%
Japan Equity	28.90%	-0.04%	7.85%	14.70%	20.55%	21.95%
Japan Bonds	7.51%	-21.43%	-13.54%	-6.69%	-0.84%	0.56%
UK Equity	21.90%	-7.04%	0.85%	7.70%	13.55%	14.95%
UK Bonds	6.82%	-22.12%	-14.23%	-7.38%	-1.53%	-0.13%
World Equity	26.75%	-2.19%	5.70%	12.55%	18.40%	19.80%
World Bonds	8.25%	-20.69%	-12.80%	-5.95%	-0.10%	1.30%
Commodities	15.20%	-13.74%	-5.85%	1.00%	6.85%	8.25%
XR Chng v. USD	0.00%	28.94%	21.05%	14.20%	8.35%	6.95%

## 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 November, this benchmark had returned 19.2%, while our model portfolio had

returned 25.6%. 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 23.1%.

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 15.2%, while our model portfolio had returned 21.3%, and the global benchmark had returned 19.4%.

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 7.1%, while our model portfolio had returned 10.6% and the global benchmark 12.0%.

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 19.2%, while our model portfolio had returned 23.1%. 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 23.1%.

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 15.2%, while our model portfolio had returned 17.6%, and the global benchmark had returned 19.4%.

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 7.1%, while our model portfolio had returned 10.7% and the global benchmark 12.0%.

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 25.7% year-to-date, our 10%

target return portfolio has returned 21.5% our 8% target return portfolio has returned 18.1%, and our 6% target return portfolio has returned 10.2%.

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 17.0%.

## 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	3.56%	4.00%	7.56%	4.90%	3.88%	8.78%
Canada	2.85%	4.00%	6.85%	2.10%	1.88%	3.98%
Eurozone	1.84%	4.00%	5.84%	2.50%	2.54%	5.04%
Japan	1.66%	4.00%	5.66%	2.70%	0.95%	3.75%
U.K.	2.18%	4.00%	6.18%	2.50%	3.30%	5.80%
U.S.A.	2.42%	4.00%	6.42%	4.50%	1.65%	6.15%

*\*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	400.84	274.80	69%	94%
Canada	116.22	293.63	253%	306%
Eurozone	112.41	147.82	131%	191%
Japan	32.09	96.61	301%	406%
U.K.	271.35	302.60	112%	157%
U.S.A.	371.84	432.69	116%	177%

\* *High productivity growth scenario.*

### **This Month's Feature Articles: Key Points**

In this month's first feature article, we look at how often you should review your portfolio's performance, and rebalance it back to your target asset class weights. We tested a number of different strategies, and concluded that for our 3% target return portfolio, rebalancing annually back to our target asset class weights added the most to the probability of achieving our long term return goal. For the 5% and 7% target return portfolios, the best approach turned out to be different. In these cases, we found that it was best to rebalance only when the actual portfolio weight of at least one asset class was more than 20% above or below its target. We also found that rebalancing your most overweight asset class to 2.5% below its target weight, and your most underweight asset class to 2.5% over its target weight (with the others to their target weights) maximized the probability of achieving our long term return goal. We also review why it is psychologically difficult to implement a disciplined rebalancing strategy.

In our second feature article we look at yet another reason why it is hard for active managers to consistently pick winning stocks: from a C.E.O.'s perspective, it is very hard to consistently be a winning stock! We explore the reasons for this in depth.

Finally, in this month's issue we also note the arrival (after a year's delay) of hedge fund index products for individual investors. We review the pros and cons of investing in hedge funds in general, and take a closer look at the new products that have been launched and the underlying indexes they are based on.

## **How Often Should You Review and Rebalance Your Portfolio?**

Reviewing your portfolio's performance and rebalancing it back to your target asset class weights is one of those ideas that sounds great in theory, but in practice can get you into trouble if you're not careful.

Consider two extreme approaches. In the first case, the portfolio is rebalanced to equal asset class weights on a daily or weekly basis. In the second case, the portfolio is never rebalanced; its asset class weights simply vary over time in line with changes in their respective market capitalizations (a strategy that research suggests is implicitly used all-too-often in many 401k accounts!). In effect, the essence of the first strategy is to constantly sell winners and buy losers (call it “extreme value”), while in the second it is to do just the opposite (call it “extreme momentum”). Intuitively, most people sense that the best approach is probably somewhere in between these two. To put it a bit more formally, most people sense that both value, (i.e., the idea that asset returns tend to revert over time to their long term average), and momentum (i.e., the idea that returns tend to continue in the same direction) both contribute to long-term asset class returns, though to varying degrees from year to year.

Academic research supports this view. For example, in their recent paper “Momentum and Mean Reversion Across National Equity Markets”, Balvers and Wu find considerable evidence that both mean reversion and momentum affect security returns, though over different time periods. Typically, mean reversion (also known as contrarian or value strategies) operates over time horizons of roughly three years, while momentum is a shorter term phenomenon, working over a one year horizon.

When you think about it, these findings intuitively make sense. Mean reversion in financial markets reflects the operation of real economic processes at the firm level, where companies often earn high returns following innovation, which are subsequently eroded as competitors either copy them at lower cost or introduce even newer offerings themselves. In contrast, momentum in financial markets is solidly grounded in individual investor psychology. It is widely recognized that it takes far less information to form an initial opinion than it does to change it. Moreover, once formed, an initial opinion affects not only the information to which we pay attention, but also the weight we give to it. Both of these factors cause us to be overconfident about the accuracy of our forecasts for the future (e.g., a belief that a security or market will continue to deliver high returns).

Given these findings, portfolio rebalancing is essentially a risk management tool, in which you are trying to manage this balance between value and momentum so as to maximize your return per unit of risk. Unfortunately, rebalancing tends to be costly, in terms of both the transaction fees involved and tax payments it can trigger (if your investments are held in a taxable account). So not only will very frequent rebalancing limit one's exposure to momentum driven gains, but the taxes and transaction fees involved will also further reduce returns. On the other hand, very infrequent rebalancing will minimize transaction fees and taxes and maximize your exposure to the momentum factor, but will also minimize your exposure to the value factor (that is, to the gains from mean reversion).

Given the relative unattractiveness of these two extremes, the real challenge is how to maximize your exposure to gains from both momentum and mean reversion, while limiting transaction costs and taxes. To answer this question, we analyzed the potential benefits from a number of different rebalancing strategies.

These strategies are based on an investor's decisions on a number of sub-issues. The first is whether rebalancing should be triggered by the calendar (e.g., do it every December), or by the extent to which actual portfolio weights have diverged from their targets. In the case of the former, the second issue is how often one should rebalance – e.g., every year, every other

year, or some other period. Where rebalancing is determined by portfolio weights, the second issue is the level at which one should set the trigger – e.g., when one or more asset class is five percent above or below its target weight, twenty percent away from it, or somewhere in-between. The third issue is what one should do when one rebalances – should you simply return your portfolio to its target weights, or should you try to take advantage of mean reversion by rebalancing overweight asset classes to slightly below their target weight, while rebalancing underweight asset classes to slightly above their target weights?

We tested all of these questions in our analysis, as well as the impact of transaction costs. We evaluated time based rebalancing using one, two, and three year frequencies. We evaluated portfolio weight strategies that were triggered when the actual weight for one or more asset classes was either 5%, 10%, or 20% above or below their target weights. We also looked at the impact of rebalancing the two most “misaligned” asset classes back to 2.5% and 5.0% above or below their respective target weights (above target weights for the worst performing asset class, and an equivalent amount below the target weight for the best performing asset class). Finally, we varied the transaction cost involved in executing the rebalancing, from a low of .35% of the value of the transactions, to a high of 1.35% (to reflect both trading commissions and bid/ask spreads).

We did not, in our initial analysis, include the cost of capital gains taxes potentially triggered by the sale of above target weight assets held in taxable accounts. In a subsequent analysis, we took them into account by raising the effective transaction costs involved. Doing this did not significantly affect our conclusions. Moreover, in many cases capital gains triggered by rebalancing can be offset by other transactions. For example, assume you have to realize \$100 in capital gains due to your rebalancing driven sale of a European equity index fund. Further assume that you will be redeploying these funds into U.S. equities, where the market value of your current position (an index fund that tracks the Wilshire 5000 index) is \$100 below your target weight. One way to offset your capital gain on the sale of the European index fund shares could be to sell your \$100 holding of the Wilshire 5000 index fund (which would generate a capital loss), and then invest \$200 into another U.S. equity index fund that tracks a different broad based U.S. equity index, such as the Dow Jones Total Market Index or the

Russell 3000 index. A number of writers have noted that the change in the underlying index that is tracked appears to exempt this type of transaction from the "wash sale" rule, and therefore make it an attractive tax management tool. However, given that few regulations seem to remain constant in the world of taxes, one should always confirm with a tax advisor that this interpretation still applies before executing this type of transaction.

We ran our rebalancing analysis on our U.S. dollar 3%, 5%, and 7% target real return portfolios, and their respective asset allocation weights. Our criterion for evaluating the impact of different rebalancing strategies was the extent to which they increased the probability of achieving the portfolio's target rate of return over a twenty year holding period. Our analysis was based on a simulation optimization approach. We first started with a candidate rebalancing strategy (e.g., time or trigger based, time period or trigger level, and adjustment factor, if any), and then simulated how well it performed over 10,000 simulations of different annual returns each year for each asset class and different costs for rebalancing transactions. We used two different distributions for future asset class returns. The first (which we used in 67% of the simulations) was based on historical asset class returns. The second (which we used in 33% of the simulations) was based on our estimated future asset class returns, as described in our recent asset allocation review. To be consistent with our previous work, the 67%/33% split was the same one we used when formulating our recommended model portfolios. The 10,000 simulations produced a distribution of expected outcomes for the candidate rebalancing strategy. We then switched to a new rebalancing strategy, and repeated the process. After we had evaluated all the possible rebalancing strategies, we chose the one which maximized the probability of achieving the target real compound annual rate of return over the twenty year holding period. To put it slightly differently, our goal was to identify the most "robust" rebalancing strategy – that is, the one which would maximize the probability of achieving the target rate of portfolio return under a wide range of possible future asset class return and transaction cost scenarios.

Our results were very interesting. Let's start with the 7% target real return portfolio. Our base case was rebalancing every year back to the portfolio's target weights. We found that the optimum rebalancing strategy for this portfolio (that is, the one which maximized the

probability of achieving the real target rate of return over twenty years) was to set the trigger at 20%, and, when one or more asset classes reached this point (i.e., were either 20% above or below their target levels), to rebalance the asset class furthest above its target to 2.5% below it, the asset class furthest below its target to 2.5% above it, and all the others to their target weights. Compared to the base case, this rebalancing strategy raised the probability of achieving the target real portfolio return by just under three percent.

Our analysis of the 5% target real return portfolio found that the same rebalancing strategy produced the best results – again increasing the probability of achieving the target real portfolio return by just under three percent. As was the case with the 7% target return portfolio, we observed that, due to the impact of diversification, the simulated asset class weights in the portfolio rarely reached the 20% trigger point over 10,000 simulations. This minimized transaction costs, while also allowing the maximum gains from momentum effects. Moreover, given the relative rarity with which the 20% over or under rebalancing trigger level was reached in these simulations, the use of the 2.5% adjustment factor usually enhanced the portfolio benefit when asset class returns reverted toward their long term means.

The 3% target real return portfolio proved to be a different case. Here the incremental benefits of different rebalancing strategies (compared to the baseline strategy of rebalancing every asset class every year back to its target weight) were very slight – delivering at most, a .50% (that is, one half of one percent) improvement in the probability of achieving the target real rate of return over twenty years. The reason for this was the fact that most of the 3% target real return portfolio is allocated to asset classes with relatively low standard deviations of returns. In other words, if the returns on the asset classes in which you've invested aren't likely to vary much from year to year, the additional benefit from not using the base case rebalancing strategy is likely to be minimal, and probably not worth the time involved.

Finally, we note that these rebalancing strategies generally agree with another important finding from asset allocation research. As we have frequently noted, because of unavoidable estimation errors in their assumptions about future asset class returns, risks, and correlations, the results produced by all asset allocation models are themselves only statistical estimates of

the "true" optimal weights for different asset classes in a portfolio. This means that portfolios whose asset class weights differ by five to ten percent from each other could, statistically speaking, be equally optimal. On the other hand, the transaction costs and taxes incurred by rebalancing are very real. In the past, this led us to reach the qualitative conclusion that frequent rebalancing was, at best, of dubious value (the exception being those very rare situations when an asset class is obviously and substantially either over or undervalued). We now have a quantitative analysis which supports this view.

But there is more to the rebalancing story than statistics. We also have to recognize that psychologically, it can be a very difficult strategy to apply in practice.

Rebalancing asks you to sell assets that have performed relatively well in the past (and about whose performance you may have bragged a bit), and to buy those that have done just the opposite. Not surprisingly, being human, many people find this very difficult to do. Consider this example. Last year, after much careful research and consideration, your friend Susan decided on the asset allocation that was right for her, and implemented it through a mix of mutual and exchange traded funds. However, when reviewing her portfolio's performance twelve months later, she saw that while her U.S. equity and U.S. bond asset class allocations had increased in value, her international equity and bond allocations had suffered a loss. What do you think she feels, and how do you think she will react?

Psychologists have identified the concept of "prospect theory", which helps to answer these questions. Prospect theory basically has three elements. The first is the observation that for many people, the emotional pain of a loss is about twice as severe as the enjoyment one experiences after an equivalent gain. The second is that losses and gains are usually determined relative to some reference point. And the third is that following losses, people become risk seekers, willing to take chances to relieve the painful feelings they are experiencing, while following gains people become risk averse, and less willing to take chances that might result in the pain and disappointment associated with losses.

Now let's apply this to our friend Susan. Prospect theory suggests that a key issue is the nature of her reference point. Unfortunately, there is no obvious answer to this question. Let's look at two alternatives. The first reference point is the best investment performance achieved by Susan's friends over the past year (assuming they're telling the truth about their portfolios' performance, which isn't always the case!). As consumer research has shown, people increasingly don't want to be "just average." The disappearance of traditional "mass market" products and stores, and their replacement by "mass luxury" items and outlets is testimony to the power of this trend. Given this, it is easy to see how Susan could experience feelings of loss if the performance of her portfolio appeared to lag behind that of one or more of one's peers.

In this situation, Susan's feelings of loss could also be compounded by something called "hindsight bias", which is a normal human tendency to believe that the events that occurred in the past were more likely than our previous foresight originally had estimated them to be. Consider a the findings from a typical experiment in this area. On January 1st, you may estimate that there is a 50% probability that U.S. bonds will be the best performing asset class this year, a 25% chance that it will be U.S. stocks, a 20% chance that it will be foreign stocks, and a 5% chance that it will be foreign bonds. Assume that by year end, it turns out that foreign stocks have delivered the best performance. If you are then asked to recreate from memory your January probability estimates, it is almost certain that the probability you assign to foreign stocks will be higher than your original 20% estimate. This is the essence of hindsight bias. In practice, hindsight bias probably adds to Susan's feelings of frustration at not having matched her best performing peer's investment results, and tempts her to adopt a riskier asset allocation in the hope of outperforming her peers the following year.

But what if Susan's reference point isn't the performance of her friends' portfolios? What if it is simply whether or not at year end her investments are worth more or less than they were twelve months earlier? Frankly, there isn't much difference at all. Prospect theory suggests that if Susan has lost money, she'll be tempted to change her asset allocation toward a riskier stance, while if she has made money, she will be tempted to adopt a more conservative one.

The moral of the story is this: the whole purpose of diversifying your portfolio across different asset classes is to reduce its overall risk, and to maximize your long term returns by avoiding big short term losses. A disciplined approach to rebalancing is a key part of your diversification strategy. In practice, diversification means that from year to year, different asset classes will have positive and negative returns. Susan needs to keep in mind that the benefits of her portfolio diversification strategy will only become apparent (and get bigger, to boot) over time. Given this, the biggest challenge for Susan is to avoid overreacting when one of the asset classes in her portfolio, or even the portfolio itself has a bad year. As we have seen, human nature makes this level of self-control a very difficult challenge to master. And the more frequently we review our portfolio's performance, the greater the psychological pressure we put on ourselves to deviate from our rebalancing strategy. Besides resisting the urge to evaluate your performance too often, it also helps to be explicit about the reference point you are using when you do evaluate it, and to choose one that is long-term and internally focused (e.g., am I on track to achieve my long term goals), and not short term and externally focused (e.g., did I outperform my Uncle Carl last year?).

Finally, we should cover one special rebalancing situation that is occasionally encountered: how to implement a major change in your portfolio's asset allocation. In many cases, the investor considering such a change has substantial amounts of unrealized gains and losses that could trigger significant tax consequences were they to substantially change either their asset allocation and/or their holdings within different asset classes (e.g. switching from actively managed to index funds). Under these circumstances, professional tax advice is a critical input when deciding whether, when, and how much change to undertake. To be sure, one should keep in mind some general guidelines on asset location and tax efficiency (see, for example, the article on this subject in our March, 2003 issue, which can also be accessed via a green button on our home page). And one should also keep in mind the trade-offs between index mutual funds and exchange traded index funds (ETFs). For example, the former are usually much better for people who are gradually changing their position, because the commission costs charged on ETF purchases are avoided. On the other hand, index mutual funds offer less control over the realization of future taxable capital gains distributions than do ETFs. However, to reiterate the main point that applies in these situations, because of the

complexity of many investors' tax situations, these general principles must be integrated with (and occasionally traded off against) the specifics of an individual's tax situation.

So, to sum up, it is important to keep these points in mind:

- Be clear about the reference point you will use to evaluate your portfolio's performance. Remember that short-term, externally focused reference points can easily get you into trouble. Keep in mind that the key question is whether or not you are still on track to achieve your long term goals.
- Unless you are confronted with an obviously and substantially overvalued asset class (which you don't see very often), or an obviously and substantially undervalued one (which you almost never see), don't rebalance any more frequently than once per year.
- If your target long term real rate of return is 3%, rebalancing your portfolio once per year back to your target weights is a good strategy to follow. However, if your long term target rate of return is 5% or 7%, you have more to gain from carefully balancing your potential exposure to momentum and mean reversion. Our analysis concludes that these portfolios should be rebalanced only when the actual weight on one or more of their asset classes deviates by 20% or more from their target weights. Rebalancing more frequently risks adding real transaction and tax costs for what may be statistically insignificant benefits.
- When you do rebalance, rather than simply going back to your target portfolio weights, reset your most overweight asset class to 2.5% below its target weight, and your most underweight asset class to 2.5% above its target weight, to take advantage of the tendency of asset class returns to revert back to their long term average over time. Rebalance the other asset classes back to their target weights.

## **Delivering Superior Returns: the CEO's Perspective**

Last month we discussed some of the reasons why it is difficult for active managers to consistently deliver returns above those on index funds, after taking sales loads, expenses, trading costs, and taxes into account. First, most active managers' forecasting skills probably aren't consistently accurate over time. Sometimes this skill depends on access to private information, the volume and quality of which usually varies over time. In other cases, superior forecasting skill is based on a unique model for making sense of the information available to the manager. However, models can be copied, or made obsolete by changes in the underlying economic processes which generate returns on financial assets.

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

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

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

A strategy is, in essence, an organization's answers to a series of interrelated questions. To simplify matters, assume that these answers are of the simple "yes or no" variety. Further assume that these decisions cover, at minimum, marketing, production, finance, human resources, systems, mergers and acquisitions and research and development issues. Assuming just five decisions in each policy area (an unrealistically low number, given our yes/no format), the company's challenge is to choose a strategy from  $2^{35}$  possible options. And this probably underestimates the number of choices available, as in the real world the total number of variables that define the landscape containing all possible strategies open to a manager can never be known with certainty.

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

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

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

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

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

Before we assumed that the shareholder value created by a set of strategy decisions only depended on a subset of the decisions themselves (e.g., the value created by a marketing decision depends not only on other marketing decisions, but also on decisions taken in manufacturing and finance). In reality, the value created by an internal decision also depends on decisions made by external parties to which the company is linked, which could include customers, competitors, suppliers and/or regulators. For example, adding new functions to a product (e.g., web browsing, email, and a digital camera to a mobile phone) will only create

additional shareholder value if customers' needs have also changed in this direction, and if competitors' have not yet introduced superior offerings, and if suppliers' prices for critical components aren't higher than expected.

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

Finally, most companies face significant organizational limitations on their ability to navigate this landscape. As individuals, corporate leaders are subject to the same cognitive limitations as the rest of us. It takes more information to change their views than it does to initially form them. And once formed, leaders' existing views affect the information they pay attention to, and the weight they give to it, particularly when those views have generated above average results. Under these circumstances, information that conflicts with current views is often either overlooked or not taken as seriously as it should be. Moreover, these individual cognitive shortcomings are usually reinforced by group processes which promote conformity and discourage dissent. In many cases, all these factors cause companies to become overconfident, miss important changes in their environment, make poor investments, and eventually to end up on the receiving end of nasty business surprises that lead to substantial performance declines.

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

Consequently, very, very few companies consistently deliver above average shareholder returns over long periods of time. For example, in their paper “Industry, Corporate, and Segment Effects and Business Performance”, Ruefli and Wiggins found that between 1984 and 1996 only 18.8% of a sample of almost 5,000 companies were able to deliver superior performance (as measured by return on assets) in any given year, while only 3.5% were able to deliver superior performance in every year. Similarly, in their paper “Is Performance Driven by Industry or Firm Specific Factors?” Hawawini, Subramanian, and Verdin found that for most firms, industry factors had the biggest impact on performance. Only a few firms were either well above or well below the industry average. While most firms are able to institute the basic management practices needed to avoid poor performance, they find it exceptionally difficult to identify and successfully implement a set of strategic choices that will deliver sustained superior performance. A simulation based study by Philip Auerswald from Harvard (“The Complexity of Production, Technological Volatility, and Inter-Industry Differences in the Persistence of Profits Above the Norm”) covers the same issues, and reaches a similar conclusion. Finally, the existence of these dynamics also makes it very hard to predict a company’s future rate of growth (which, for most firms, is the major determinant of market value and shareholder returns). As Chan, Karceski, and Lakonishok document (in their paper “The Level and Persistence of Growth Rates”), company growth rates tend not to persist over time beyond what would be expected due to luck alone.

As you can see, the challenges facing a C.E.O. who is trying to consistently deliver superior shareholder returns are extremely daunting, and the probabilities of long term success are very low. By definition, the probability that an active investment manager will be able to consistently identify these few winning firms in advance therefore must be lower still – the probability of corporate success must be higher than the joint probability of corporate success and accurate foresight by our investment manager. Seen in this light, the heavy dependence on momentum (that is, on accurately forecasting other investors’ behavior) by successful active managers should come as no surprise. As Professor Russ Wermers concluded in his recent paper “Is Money Smart?” the “finding that [active fund] performance does not seem to persist after controlling for cash inflows [and the momentum buying they finance] casts doubt on previous studies that found that [active] managers have talents in choosing stocks that beat

their benchmarks.” So, once again, we reach a now familiar conclusion: it is extraordinarily difficult for most mortals to be consistently successful at active investment management over long periods of time, and the great majority of investors would be much better off investing in a diversified portfolio of index funds.

## **Retail Hedge Fund Index Products Finally Arrive**

In the past, we have written about hedge funds (also known as absolute value strategies), and how they might fit into an investor's portfolio (see our June, 2003 and April, 2002 issues). However, until now, we have not included them in our model portfolios because there have been no broadly diversified products available, at a reasonable cost, to most individual investors. The ones which have been on offer, known as "Funds of Funds", are quite expensive: in effect, you often pay quite high fees not only to the hedge fund managers themselves (typically one percent of the assets under management, plus twenty percent of the profits), but also to the Fund of Funds manager (typically a one to three percent management fee, plus a sales load and/or a share of the profits distributed by the underlying hedge funds). A second issue with the Fund of Fund concept is the question of fund selection: in effect, when investing in a Fund of Fund, you are hiring an active manager to hire other active managers. Not only is this an expensive proposition, but in effect, it raises the already daunting challenge of active management to the third power: you have to choose a winning fund of funds manager, who will pick winning hedge fund managers, who will make winning investments. Got it? Comfortable with it? Neither were we. Given the potential risks and returns presented by the Fund of Funds approach, we have not been of the opinion that most Funds of Funds represent a good value for investors.

This situation has now begun to change, with the introduction in Canada of hedge fund index products for individual investors, as well as new tracker funds in Ireland and Jersey. In this note, we will take a closer look at these, and explore their advantages and disadvantages. In January, we will publish updated model portfolios that include hedge fund index products in

their universe of possible investments (next month's end of year issue already has enough numbers in it!).

We'll begin with a critical point, from which other important ones flow. Hedge funds are not an asset class. Asset classes represent some type of claim on real productive assets that share common characteristics. As noted last month's article on active investment management, the return on an asset within an asset class has two components: compensation for the risk of the asset class itself (also known as "beta risk"), and compensation for risks unique to the specific asset under consideration (also known as "alpha risk"). When you hold a diversified portfolio of assets from within the same asset class, the alpha risks (and the returns associated with holding them) cancel each other out, and you are left with non-diversifiable risk (also known as beta or systemic risk), and the return for holding it. When you diversify your portfolio across asset classes, the beta risk is reduced, but not eliminated.

In contrast to a true asset class, the broad term "hedge funds" refers to a very diverse collection of actively managed investment strategies which aim to maximize the return for holding alpha risk in one or more asset classes. Investing in hedge funds is attractive for two reasons. First, because hedge fund managers seek to maximize alpha, their exposure to beta (or asset class) risk is often minimal (e.g., one could invest in specific shares, and use derivatives to hedge its exposure to overall market movements, leaving the fund with pure company specific alpha risk and return). As a result, hedge fund returns often have very low correlations with the returns earned on index funds which track the returns on different asset classes. Given this, including hedge funds in a portfolio should produce diversification benefits, in the form of a reduced level of overall risk.

Second, there is also an argument to be made that hedge funds should produce higher returns than actively managed mutual funds that invest in the same asset class. Three justifications are typically given to support this argument. First, hedge funds offer much better compensation packages to their managers than do actively managed mutual funds (e.g., twenty percent of the profits). Theoretically, this should attract the best active managers to hedge funds, and provide a strong incentive for them to make the best use of their abilities. Second, hedge

funds place far fewer restrictions on the ability of a manager to translate his or her forecasts into portfolio positions. For example, unlike mutual funds, many hedge funds can and do take short positions (i.e., sell a security they do not own in the expectation that its price will go down, and they can later buy it -- or "cover their short" -- at a price lower than the one at which they sold it), and employ leverage (that is, use debt to take on larger positions than they could using just the money investors have contributed to their fund). Finally, while actively managed funds have to keep a portion of their asset in low return cash deposits (to cover any imbalance between daily fund redemptions and sales), hedge funds do not, because their investors cannot withdraw their funds on short notice.

However, there is also an important theoretical disadvantages to investing in a diversified group of hedge funds (either through a Fund of Funds or through an index product). Hedge funds employ active management strategies, which, in aggregate, are a less than zero sum game. As previously noted, in a broadly diversified portfolio, alphas net out, and all that is left are the fees paid to the active managers (in short, all that is left is a very expensive de facto index fund!). This is true for people who invest in a broad range of actively managed mutual funds. And, in theory, it should also be true for people who invest in a wide range of hedge funds, unless hedge funds as a group earn a positive alpha (which, in turn, implies that some other group of active investors is earning a negative alpha). This is the question you must confront before you decide to invest in a broad hedge fund index product.

We'll tell you what we think. Given the extreme difficulty of being a consistently successful active manager, common sense suggests that the best of this group will seek out the opportunities that will offer them the greatest rewards for their skills. That means the best managers are more likely to be found at hedge funds than at actively managed mutual funds. On balance, this leads us to conclude that higher alphas are likely to be earned by hedge fund than by actively managed mutual fund managers within any given asset class. Moreover, the data show that, to date, the alphas on different hedge fund strategies generally have not cancelled each other out (although this still may happen, as more money flows into hedge fund strategies). On balance, we conclude that if a typical individual investor is going to use active management, the best way to do it is through a hedge fund index vehicle: the cost will

be lower than a Fund of Funds, the return will probably be higher than what you would earn on an equivalent allocation to actively managed mutual funds, and it also will probably have a lower correlation with the returns on the asset class index funds already in your portfolio.

Another issue which has concerned us is the specific funds that will be chosen for inclusion in the hedge fund index tracked by the new funds. Two considerations are involved. First, some strategies have more "capacity" than others. This means that some strategies can keep on delivering positive returns even when they receive large inflows of new funds (global macro funds are an example). In contrast, the returns on other strategies can be adversely affected by a large inflow of funds (e.g., distressed debt investing, where fund inflows would result in prices being bid up, which lowers expected returns). Second, even within broad classes of strategies, some funds may have more capacity than others, depending on the specific investment technique they are using.

To address these issues, the companies which produce the various broad hedge fund performance indexes have created subsets of them, which they usually call their "investable" index. These are the indexes which the new hedge fund based index products will attempt to track. What needs to be recognized clearly is that the decision on which hedge funds to include in these investable indexes is, at its core, an active management decision (which, to be accurate, could also be said about the decision to include or remove a company from an index like the S&P 500, which is only a subset of the broader market). Depending on the hedge funds selected, the returns on the investable index may track the overall hedge fund index returns quite closely, or they may diverge from them by a large amount. As these products are quite new, only time will tell.

Now, on to the new products themselves. The first is called the Tremont Hedge Fund Index Linked Trust, which will be traded on the Toronto Stock Exchange. It is designed to track the CSFB/Tremont Investable Hedge Index, which includes sixty different funds, representing the six largest funds in ten different broadly defined investment strategies. As we described in our April, 2002 article, we believe that CSFB/Tremont uses a very good index construction methodology. Moreover, we believe that using the largest hedge funds to construct the

investable index should ensure relatively close tracking with the overall CSFB/Tremont hedge fund index, because the latter is asset weighted. To make it more attractive to Canadian investors, this fund will employ a creative structure which will defer income and taxes until 2010. According to its preliminary prospectus, its management fee will be a very reasonable .55% per year, which is well below the fees usually charged by Funds of Funds.

The second offering is from One Financial. It takes the form of a capital guaranteed note, whose return is linked to the performance of the Morgan Stanley Capital International Hedge Invest Index. Like the CSFB/Tremont index, the latter is asset weighted, and utilizes a strong underlying methodology. The actual note (a debt instrument) matures in about ten years, and, as a return, offers the greater of ten percent or 150% of the return on the MSCI Hedge Invest Lyxor Tracker Fund (less a charge for the capital return guarantee). At 1.50%, the underlying expense ratio is higher than the one on the Tremont product. On balance, we prefer the Tremont offering to this one, because of its lower costs and cleaner structure.

Finally, at press time we were still trying to obtain more detailed information about the new MSCI Hedge Invest Lyxor Tracker Fund that has been registered in Ireland and in Jersey by a subsidiary of Societe Generale. However, the message seems to be clear: in 2004 we will no doubt see the launch of many similar products, which we will continue to track and evaluate. In the meantime, we will begin the new year with an alternative set of model portfolios which incorporate allocations to hedge fund index products.