

The Index Investor

Invest Wisely...Get an Impartial Second Opinion.

Contents

<i>This Month's Issue: Key Points</i>	1
<i>This Month's Letters to the Editor</i>	3
<i>Global Asset Class Returns</i>	5
<i>Asset Class Valuation Update</i>	6
<i>The Exchange Traded Fund Market</i>	17
<i>Active Management and Risk Budgeting</i>	28
<i>Product and Strategy Notes</i>	43
<i>2006-2007 Model Portfolios Update</i>	54

This Month's Issue: Key Points

Our first article this month reviews recent developments in the exchange traded funds market, and the direction they may take in the future. We start with the basic division the return on an asset into its two component parts: market return and unique return. All investors in an asset class earn the former, regardless of their skill. This basic asset class return is also known as “beta.” In contrast, across a whole asset class, positive and negative returns that are unique to a security (or a group of securities) cancel each other out, leaving only market return. Hence, earning a positive unique return (also known as “alpha”) is a reward for forecasting skill. We describe how alpha can be further sub-divided into different parts. Some alpha represents the reward for tilting a portfolio towards certain recognized risk factors that represent a sub-segment of the overall asset class (e.g., a tilt toward small cap value stocks) during periods when those tilts earn positive returns. The alpha that remains after subtracting the return on these “factor bets” comes from up to four sources: market timing, security selection, and providing insurance and/or liquidity to other investors. We note the confusion caused by the fact that returns on factor bets are tracked by indexes, as are market returns on broad asset classes. We note that index investing is not synonymous with passive investing. The latter

seeks to earn only market returns on broad asset classes that are not conditional upon skill. In contrast, positive returns on some index products (e.g., those that track factor-bets) are conditional upon forecasting skill.

We note how the exchange traded funds market now offers both pure market return products, factor-bet products, and products that are, for all intents and purposes, security selection products dressed up in index clothing. On the positive side, the expanding range of ETF products creates more opportunities to earn alpha at a low cost; on the negative side, it may cause overconfident investors to reduce their returns by overtrading.

Our second feature article expands on the first, and takes a deeper look at the arguments that support the use of active management. We do not doubt that successful, skill based active management is possible, and present evidence that supports this conclusion. However, we also explain why we believe that long-term success in this area (defined as outperformance of an asset class index fund after expenses and taxes) is beyond the reach of most individual investors. Despite this, there is abundant evidence that many investors will still choose to take both market and active risk in their portfolio. To help them better understand and manage the issues this raises, we present a short overview of risk budgeting. We show how it can be used to control the level of active risk exposure, allocate it efficiently, and highlight potential inconsistencies in a portfolio's structure.

In this month's product and strategy notes, we review four recent working papers from the IMF and World Bank that suggest that we are not alone in our belief that the unprecedented imbalances that now exist in the world economy are more likely to unwind with a bang than a whimper. We also show statistically how diversification across asset classes provides exposure to different underlying return generating processes. Our third note summarizes a fascinating new paper that compares mutual fund fees around the world. We conclude that it is time for a Canadian revolt. Our last note reviews a surprisingly difficult question: how do pension plan sponsors and product developers decide which indexes to track?

This Month's Letters to the Editor

What is your opinion of products that protect your principal while tying returns to an equity index?

These are known as “structured products” and are technically a type of debt instrument. The goal of most structured products is to provide a return that is greater than that available on straightforward debt but with similar principal protection. An alternative way to view them is as a straightforward equity investment with a series of rolling insurance contracts (say, put options) that protect the investor's initial capital against the possibility of a big fall in the market. Structured products have historically been quite popular in Europe, and, given investors' increasing concern with U.S. Equity market valuations, are now appearing in greater numbers in the United States.

Their clear advantage is ease of use. However, they also have some not-so-clear disadvantages. The most important is usually the pricing of the underlying insurance contracts. When you decompose a structured product, you often find that the implicit price being charged for the insurance (i.e., the put option) is higher than the market price; this difference presumably compensates the sponsor of the structured product for the cost of its creation and marketing, as well as providing a profit to them. Depending on the product, the underlying profit to the issuer can be quite rich.

The second issue is one that is often true of insurance in general: does a structured product “over-insure” an investor? For example, it may be the case that what the investor really seeks is not protection against all falls in the equity market (which is what the “guaranteed return of principal” approach provides), but rather “disaster insurance,” which only pays out in the case of a severe equity market decline. Because the probability associated with the latter is much lower than the former, the cost of this “disaster insurance” (i.e., a deep out of the money put option) is usually much lower than the cost of the insurance embedded into the structured product.

The third issue is whether an investor really needs a structured product at all. As we have stressed in our writing, it is the returns over time on the portfolio as a whole, rather than the ups and downs of any one asset class, that determines one's purchasing power. When the

equity market declines, other asset classes should increase in value, moderating the overall impact on the portfolio. Think of diversification across asset classes as a form of “self-insurance.” However, it seems to be in our human nature to focus on the individual elements of the portfolio as well as its overall return. Hence, the temptation to invest in structured products, and the appearance of many new ones as people become more nervous about the future state of the economy.

Global Asset Class Returns

<i>YTD 31 Jul 06</i>	<i>In USD</i>	<i>In AUD</i>	<i>In CAD</i>	<i>In EURO</i>	<i>In JPY</i>	<i>In GBP</i>	<i>In CHF</i>	<i>In INR</i>
Asset Held								
US Bonds	0.40%	-4.10%	-2.36%	-7.49%	-2.47%	-8.17%	-6.23%	3.51%
US Prop.	17.20%	12.70%	14.44%	9.31%	14.33%	8.63%	10.57%	20.31%
US Equity	3.10%	-1.40%	0.34%	-4.79%	0.23%	-5.47%	-3.53%	6.21%
AUS Bonds	-1.38%	-5.88%	-4.14%	-9.27%	-4.25%	-9.96%	-8.01%	1.72%
AUS Prop.	14.48%	9.98%	11.72%	6.59%	11.61%	5.91%	7.85%	17.59%
AUS Equity	11.43%	6.93%	8.67%	3.54%	8.56%	2.86%	4.80%	14.54%
CAN Bonds	3.57%	-0.93%	0.81%	-4.32%	0.70%	-5.00%	-3.06%	6.68%
CAN Prop.	10.82%	6.32%	8.07%	2.94%	7.96%	2.25%	4.20%	13.93%
CAN Equity	8.58%	4.08%	5.83%	0.70%	5.72%	0.01%	1.96%	11.69%
Euro Bonds	6.64%	2.14%	3.88%	-1.25%	3.77%	-1.93%	0.01%	9.75%
Euro Prop.	27.75%	23.25%	24.99%	19.86%	24.88%	19.18%	21.12%	30.86%
Euro Equity	14.76%	10.26%	12.00%	6.87%	11.89%	6.19%	8.13%	17.87%
Japan Bonds	1.39%	-3.11%	-1.37%	-6.50%	-1.48%	-7.18%	-5.24%	4.50%
Japan Prop.	0.87%	-3.63%	-1.89%	-7.02%	-2.00%	-7.70%	-5.76%	3.98%
Japan Equity	0.37%	-4.13%	-2.39%	-7.52%	-2.50%	-8.20%	-6.26%	3.48%
UK Bonds	8.18%	3.68%	5.42%	0.29%	5.31%	-0.39%	1.55%	11.29%
UK Prop.	31.10%	26.60%	28.34%	23.21%	28.23%	22.53%	24.47%	34.21%
UK Equity	16.25%	11.75%	13.50%	8.37%	13.39%	7.68%	9.63%	19.36%
World Bonds	2.45%	-2.05%	-0.31%	-5.44%	-0.42%	-6.12%	-4.18%	5.56%
World Prop.	16.13%	11.63%	13.37%	8.24%	13.26%	7.56%	9.50%	19.24%
World Equity	6.95%	2.45%	4.19%	-0.94%	4.08%	-1.62%	0.32%	10.06%
Commodities	3.50%	-1.00%	0.74%	-4.39%	0.63%	-5.07%	-3.13%	6.61%
Timber	-1.63%	-6.13%	-4.39%	-9.52%	-4.50%	-10.21%	-8.26%	1.47%
EqMktNeutral	5.31%	0.81%	2.56%	-2.58%	2.45%	-3.26%	-1.32%	8.42%
Volatility	23.86%	19.36%	21.10%	15.97%	20.99%	15.29%	17.23%	26.97%
Currency								
AUD	4.50%	0.00%	1.74%	-3.39%	1.63%	-4.07%	-2.13%	7.61%
CAD	2.76%	-1.74%	0.00%	-5.13%	-0.11%	-5.82%	-3.87%	5.86%
EUR	7.89%	3.39%	5.13%	0.00%	5.02%	-0.68%	1.26%	11.00%
JPY	2.87%	-1.63%	0.11%	-5.02%	0.00%	-5.71%	-3.76%	5.97%
GBP	8.57%	4.07%	5.82%	0.68%	5.71%	0.00%	1.95%	11.68%
USD	0.00%	-4.50%	-2.76%	-7.89%	-2.87%	-8.57%	-6.63%	3.11%
CHF	6.63%	2.13%	3.87%	-1.26%	3.76%	-1.95%	0.00%	9.74%
INR	-3.11%	-7.61%	-5.86%	-11.00%	-5.97%	-11.68%	-9.74%	0.00%

Asset Class Valuation Update

Our market valuation analyses are based on the assumption that markets are not perfectly efficient and always in equilibrium. This means that it is possible for the supply of future returns a market is expected to provide to be higher or lower than the returns investors logically demand. In the case of an equity market, we define the future supply of returns to be equal to the current dividend yield plus the rate at which dividends are expected to grow in the future. We define the return investors demand as the current yield on real return government bonds plus an equity market risk premium. As described in our May, 2005 issue, people can and do disagree about the “right” values for these variables. Recognizing this, we present four valuation scenarios for an equity market, based on different values for three key variables. First, we use both the current dividend yield and the dividend yield adjusted upward by .50% to reflect share repurchases. Second, we define future dividend growth to be equal to the long-term rate of total (multifactor) productivity growth, which is equal to either 1% or 2%. Third, we use two different values for the equity risk premium required by investors: 2.5% and 4.0%. Different combinations of these variables yield high and low scenarios for both the future returns the market is expected to supply, and the future returns investors will demand. We then use the dividend discount model to combine these scenarios, to produce four different views of whether an equity market is over, under, or fairly valued today. The specific formula is $(\text{Current Dividend Yield} \times 100) \times (1 + \text{Forecast Productivity Growth})$ divided by $(\text{Current Yield on Real Return Bonds} + \text{Equity Risk Premium} - \text{Forecast Productivity Growth})$. Our valuation estimates are shown in the following tables, where a value greater than 100% implies overvaluation, and less than 100% implies undervaluation:

<i>Australia</i>	Low Demanded Return	High Demanded Return
High Supplied Return	67%	102%
Low Supplied Return	103%	143%

<i>Canada</i>	Low Demanded Return	High Demanded Return
High Supplied Return	97%	161%
Low Supplied Return	181%	264%

<i>Eurozone</i>	Low Demanded Return	High Demanded Return
High Supplied Return	66%	109%
Low Supplied Return	111%	161%

<i>Japan</i>	Low Demanded Return	High Demanded Return
High Supplied Return	98%	194%
Low Supplied Return	243%	388%

<i>United Kingdom</i>	Low Demanded Return	High Demanded Return
High Supplied Return	49%	90%
Low Supplied Return	89%	137%

<i>United States</i>	Low Demanded Return	High Demanded Return
High Supplied Return	122%	185%
Low Supplied Return	210%	291%

<i>Switzerland</i>	Low Demanded Return	High Demanded Return
High Supplied Return	89%	152%
Low Supplied Return	169%	248%

<i>India</i>	Low Demanded Return	High Demanded Return
High Supplied Return	157%	235%
Low Supplied Return	289%	398%

Our government bond market valuation update is based on the same supply and demand methodology we use for our equity market valuation update. In this case, the supply of future fixed income returns is equal to the current nominal yield on ten-year government bonds. The demand for future returns is equal to the current real bond yield plus the historical average inflation premium (the difference between nominal and real bond yields) between 1989 and 2003. To estimate of the degree of over or undervaluation for a bond market, we use the rate of return supplied and the rate of return demanded to calculate the present values of a ten year zero coupon government bond, and then compare them. If the rate supplied is higher than the rate demanded, the market will appear to be undervalued. This information is contained in the following table:

	Current Real Rate	Average Inflation Premium (89-03)	Required Nominal Return	Nominal Return Supplied (10 year Govt)	Return Gap	Asset Class Over or (Under) Valuation, based on 10 year zero
Australia	2.43%	2.96%	5.39%	5.85%	0.46%	-4.28%
Canada	1.75%	2.40%	4.15%	4.31%	0.16%	-1.56%
Eurozone	1.87%	2.37%	4.24%	3.92%	-0.32%	3.07%
Japan	1.03%	0.77%	1.80%	1.94%	0.14%	-1.35%
UK	1.33%	3.17%	4.50%	4.60%	0.10%	-0.95%
USA	2.41%	2.93%	5.34%	4.99%	-0.35%	3.35%
Switz.	1.60%	2.03%	3.63%	2.70%	-0.93%	9.43%
India	3.49%	7.57%	11.06%	8.49%	-2.57%	26.38%

*Derived from ten year yield and forecast inflation

It is important to note some important limitations of this analysis. First, it uses the current yield on real return government bonds (or, in the cases of Switzerland and India, the implied real yield if those bonds existed). Over the past forty years or so, this has averaged around 3.00%. Were we to use this rate, the required rate of return would generally increase. Theoretically, the “natural” or equilibrium real rate of interest is a function of three variables:

(1) the expected rate of multifactor productivity growth (as it increases, so to should the demand for investment, which will tend to raise the real rate); (2) risk aversion (as investors become more risk averse they save more, which should reduce the real rate of interest, all else being equal); and (3) the time discount rate, or the rate at which investors are willing to trade off consumption today against consumption in the future. A higher discount rate reflects a greater desire to consume today rather than waiting (as consumption today becomes relatively more important, savings decline, which should cause the real rate to increase). These variables are not unrelated; a negative correlation (of about .3) has been found between risk aversion and the time discount rate. This means that as people become more risk averse, they also tend to be more concerned about the future (i.e., as risk aversion rises, the time discount rate falls).

All three of these variables can only be estimated with uncertainty. For example, a time discount rate of 2.0% and risk aversion factor of 4 are considered to be average, but studies show that there is wide variation within the population and across the studies themselves. The analysis in the following table starts with current real return bond yields and the OECD's estimates of multifactor productivity growth between 1995 and 2002 (with France and Germany proxying for the Eurozone). We then try to back out estimates for risk aversion and the time discount rate that would bring theoretical rates into line with those that have been observed in the market. The real rate formula is $[(\text{Time Discount Rate} + (1/\text{Risk Aversion Factor}) \times \text{MFP Growth})]$.

Real Rate 31Jul06	AUD	CAD	EUR	JPY	GBP	USD
Risk Aversion Factor	4.0	5.0	5.0	6.0	5.0	4.0
Time Discount Rate	2.00%	1.50%	1.50%	1.00%	1.50%	2.00%
MFP Growth	1.60%	1.20%	1.40%	0.60%	1.40%	1.40%
Estimated Real Rate	2.40%	1.74%	1.78%	1.10%	1.78%	2.35%
Actual Real Rate	2.43%	1.75%	1.87%	1.03%	1.33%	2.41%

Our analysis also uses historical inflation as an estimate of expected future inflation. This may not produce an accurate valuation estimate, if the historical average level of inflation is not a good predictor of average future inflation levels. For example, if expected future inflation is lower than historical inflation, required returns will be lower. Also, if one were to assume a very different scenario, involving a prolonged recession, accompanied by deflation, then one could argue that government bond markets are actually undervalued today.

Let us now turn to the subject of the valuation of non-government bonds. Some have suggested that it is useful to decompose the bond yield spread into two parts. The first is the difference between the yield on AAA rated bonds and the yield on the ten year Treasury bond. Because default risk on AAA rated companies is very low, this spread may primarily reflect prevailing liquidity and jump (regime shift) risk conditions. The second is the difference between BBB and AAA rated bonds, which may tell us more about the level of compensation required by investors for bearing default risk. For example, between August and October, 1998 (around the time of the Russian debt default and Long Term Capital Management crises), the AAA-Treasury spread jumped from 1.18% to 1.84%, while the BBB-AAA spread increased by much less, from .62% to .81%.

The following table shows the average level of these spreads between January, 1970 and December, 2005 (based on monthly Federal Reserve data), along with their standard deviations and 67% (average plus or minus one standard deviation) and 95% (average plus or minus two standard deviations) confidence range.

	AAA – 10 Year Treasury	BBB-AAA
Average	.97%	1.08%
Standard Deviation	.47%	.42%
Avg. +/- 1 SD	1.44% - .50%	1.51% - .66%
Avg. +/- 2 SD	1.91% - .03%	1.93% - .23%

At 31 July, 2006 the AAA minus 10 year Treasury spread was 1.22. This was above the long-term average compensation for bearing liquidity and jump risk (assuming our model is correct), and represents a significant increase over the previous month. This is consistent with the significant increase in the VIX we have seen over the past month (the VIX measures the future level of equity volatility implied by the current price of options on the S&P 500 Index).

At the end of the month, the BBB minus AAA spread was .90%, basically unchanged since the end of May. This was below the long-term average compensation for bearing default risk. The stability of this spread in the face of other developments (e.g., the rising VIX) lead us to conclude that it is more likely that corporate bonds today are overvalued than undervalued.

Finally, for an investor contemplating the purchase of foreign bonds or equities, the expected future annual percentage change in the exchange rate is also important. Study after study has shown that there is no reliable way to forecast this. At best, you can make an estimate that is justified in theory, knowing that in practice it will not turn out to be accurate. That is what we have chosen to do here. Specifically, we have taken the difference between the yields on ten-year government bonds as our estimate of the likely future annual change in exchange rates between two regions. This information is summarized in the following table:

Annual Exchange Rate Changes Implied by Bond Market Yields

	To AUD	To CAD	To EUR	To JPY	To GBP	To USD	To CHF	To INR
From								
AUD	0.00%	-1.54%	-1.93%	-3.91%	-1.25%	-0.86%	-3.15%	2.64%
CAD	1.54%	0.00%	-0.39%	-2.37%	0.29%	0.68%	-1.61%	4.18%
EUR	1.93%	0.39%	0.00%	-1.98%	0.68%	1.07%	-1.22%	4.57%
JPY	3.91%	2.37%	1.98%	0.00%	2.66%	3.05%	0.76%	6.55%
GBP	1.25%	-0.29%	-0.68%	-2.66%	0.00%	0.39%	-1.90%	3.89%
USD	0.86%	-0.68%	-1.07%	-3.05%	-0.39%	0.00%	-2.29%	3.50%
CHF	3.15%	1.61%	1.22%	-0.76%	1.90%	2.29%	0.00%	5.79%
INR	-2.64%	-4.18%	-4.57%	-6.55%	-3.89%	-3.50%	-5.79%	0.00%

Our approach to valuing commercial property securities as an asset class is hindered by a lack of historical data about rates of dividend growth. To overcome this limitation, we have assumed that markets are fairly valued today (i.e., the expected supply of returns equals the expected returns demanded by investors), and “backed out” the implied growth rates to see if they are reasonable in light of other evidence about the state of the economy (see below). This analysis assumes that investors require a 2.5% risk premium above the yield on real return bonds to compensate them for the risk of securitized commercial property as an asset class. The following table shows the results of this analysis:

Country	Real Bond Yield	Plus Commercial Property Risk Premium	Less Dividend Yield on Commercial Property Securities	Equals Expected Rate of Future Real Dividend Growth
Australia	2.43%	2.50%	6.4%	-1.5%
Canada	1.75%	2.50%	4.9%	-0.6%
Eurozone	1.87%	2.50%	3.0%	1.3%
Japan	1.03%	2.50%	1.3%	2.2%
Switzerland	1.60%	2.50%	3.3%	0.8%
United Kingdom	1.33%	2.50%	2.3%	1.5%
United States	2.41%	2.50%	4.1%	0.8%

A very rough way to test the reasonableness of these expected growth assumptions is to compare them to the expected real annual change in commercial rental income over the next five years. If you think the real growth estimates are too high, that implies overvaluation. On the other hand, if you think they are too low, that implies undervaluation. Since we expect a significant slowdown in the global economy over the next few years, we are inclined to view most of these implied real growth assumptions as too optimistic.

Our commodities asset class valuation analysis is focused on two drivers of near term returns: the “roll yield” (sale of futures contracts at close to the spot price as they mature, and reinvestment of the proceeds in a new, longer-dated contract) and unexpected changes in the spot price. With respect to the roll yield, the DJ AIG commodities index futures contract traded on the Chicago Board Options Exchange (CBOT) is currently contangoed, with a (5.2%) difference between the near and far term contract prices. With respect to the spot price, over the 1991 to 2005 period, the DJ AIG had an average value of 107.6, with a standard deviation of 21.9. The July 31 closing price of 178.03 was more than three standard deviations above the average. In light of history, the probability of spot price declines is much higher than the probability of spot price increases. Hence, we conclude that the balance of evidence suggests that commodities are overvalued today.

Our approach to assessing the current value of volatility (as measured by the VIX index) is similar to our approach to commodities. Between January 2, 1990 and December 30,

2005, the average value of the VIX was 19.45, with a standard deviation of 6.40. The one standard deviation (67% confidence interval) range was 13.05 to 28.85, and the two standard deviations (95%) range was from 6.65 to 32.25. On July 31, 2006, the VIX closed at 14.95. As you can see, this is below its long-term average value, and seems low in light of conditions in the economy and financial markets. Hence, we conclude that volatility is undervalued today.

Sector and Style Rotation Watch

The following table shows a number of classic style and sector rotation strategies that attempt to generate above index returns by correctly forecasting turning points in the economy. This table assumes that active investors are trying to earn high returns by investing today in the styles and sectors that will perform best in the next stage of the economic cycle. The logic behind this is as follows: Theoretically, the fair price of an asset (also known as its fundamental value) is equal to the present value of the future cash flows it is expected to produce, discounted at a rate that reflects their relative riskiness.

Current economic conditions affect the current cash flow an asset produces. Future economic conditions affect future cash flows and discount rates. Because they are more numerous, expected future cash flows have a much bigger impact on the fundamental value of an asset than do current cash flows. Hence, if an investor is attempting to earn a positive return by purchasing today an asset whose value (and price) will increase in the future, he or she needs to accurately forecast the future value of that asset. To do this, he or she needs to forecast future economic conditions, and their impact on future cash flows and the future discount rate. Moreover, an investor also needs to do this before the majority of other investors reach the same conclusion about the asset's fair value, and through their buying and selling cause its price to adjust to that level (and eliminate the potential excess return).

We publish this table to make an important point: there is nothing unique about the various rotation strategies we describe, which are widely known by many investors. Rather, whatever active management returns (also known as "alpha") they are able to generate is directly related to how accurately (and consistently) one can forecast the turning points in the economic cycle. Regularly getting this right is beyond the skills of most investors. In other words, most of us are better off just getting our asset allocations right, and implementing them

via index funds rather than trying to earn extra returns by accurately forecasting the ups and downs of different sub-segments of the U.S. equity and debt markets. That being said, the highest rolling three month returns in the table give a rough indication of how investors expect the economy and interest rates to perform in the near future. *The highest returns in a given row indicate that most investors are anticipating the economic and interest rate conditions noted at the top of the next column* (e.g., if long maturity bonds have the highest year to date returns, a plurality of bond investor opinion expects rates to fall in the near future). Comparing returns across strategies provides a rough indication of the extent of agreement (or disagreement) investors about the most likely upcoming changes in the state of the economy.

Three Month Rolling Nominal Returns on Classic Rotation Strategies in the U.S. Markets

Rolling 3 Month Returns Through July, 2006

<i>Economy</i>	Bottoming	Strengthening	Peaking	Weakening
<i>Interest Rates</i>	Falling	Bottom	Rising	Peak
<i>Style and Size Rotation</i>	Small Growth (DSG) -10.85%	Small Value (DSV) -4.20%	Large Value (ELV) 0.71%	Large Growth (ELG) -5.48%
<i>Sector Rotation</i>	Cyclicals (IYC) -5.53% Technology (IYW) -11.85%	Basic Materials (IYM) -8.72% Industrials (IYJ) -8.55%	Energy (IYE) 3.46% Staples (IYK) 0.51%	Utilities (IDU) 8.65% Financials (IYF) -1.93%
<i>Bond Market Rotation</i>	Higher Risk (LQD) 1.16%	Short Maturity (SHY) 1.06%	Low Risk (TIP) 1.72%	Long Maturity (TLT) 3.14%

The next table describes the typical cycles in the markets for commercial property and commodities. However, rather than being leading indicators of future economic conditions, commercial property and commodity market returns tend to coincide with current economic and interest rate conditions (i.e., those at the top of the same column, rather than the next one to the right). When many investors share the same expectations about future economic conditions, one would expect to see alignment between bond and equity market year-to-date

returns, and conditions in commodity and commercial property markets. However, we also note that this is when markets are most fragile; large moves can occur if something happens to change these closely aligned expectations. In contrast, when investors do not share the same expectations for the future, you would expect to see misalignment between year-to-date returns in bond, equity, commodity and commercial property markets.

Economy	Bottoming	Strengthening	Peaking	Weakening
Interest Rates	Falling	Bottom	Rising	Peak
Commodities				
Commodity Inventories	Peaking	Falling	Bottoming	Rising
Spot Prices	Bottoming	Rising	Peaking	Falling
Futures Prices Relative to Spot Price	Contango (futures higher than spot)	Uncertain	Backwardation (futures lower than spot)	Uncertain
Profitability of long commodity futures position, before diversification and collateral yields	Negative (falling spot and negative roll yield)	Uncertain (rising spot, uncertain roll yield)	Positive (rising spot and positive roll yield)	Uncertain (falling spot, uncertain roll yield)
Comm'l Property				
Commercial Property Vacancy Rates	Peaking	Falling	Bottoming	Rising
Rents	Low	Rising	High	Falling
New Construction Completion (space coming onto the market)	Falling	Bottoming	Rising	Peaking
Property Valuation Ratios	Bottoming	Rising	Peaking	Falling
Expected Future Property Returns	Peaking	Falling	Bottoming	Rising

The following table sums up our subjective view of possible asset class under and overvaluations at the end of July 2006. The distinction between possible, likely and probable reflects a rising degree of confidence in our conclusion.

Probably Overvalued	Commodities, Corporate Bonds
Likely Overvalued	Commercial Property, Most Equity Markets
Possibly Overvalued	
Possibly Undervalued	UK Equity, Australian Bonds
Likely Undervalued	Real Return Bonds, Equity Volatility
Probably Undervalued	Non-U.S. Dollar Bonds

The Exchange Traded Fund Market

The number of exchange traded funds listed on markets around the world is growing rapidly. However, the launch of so many new ETFs has also made this part of the investment landscape a much more confusing place. With that in mind, we have prepared this short guide to understanding exchange traded funds.

We will start with a short review of investment theory, and then use it to classify the many ETFs that are now offered in markets around the world. Let's begin with the definition of an "asset class." In our view, asset classes should be broadly defined, as they are distinguished by significantly different underlying economic return generating processes. Different statistical techniques can be used to perform this analysis, including correlation (true asset classes should have returns that have low correlations with each other) and principal components analysis (true asset classes should have different loadings on different return generating factors). For example, consider the difference between domestic investment grade bonds and emerging market equity. The underlying economic processes that generate the returns on these two investments are quite different from each other.

In contrast, the processes generating returns on "large cap" emerging market equities and "small cap" emerging market equities are quite similar, as evidenced by the high correlation between their respective returns. Hence we regard these two categories not as distinct asset classes, but rather as an example of "tilts" or sub-segments within the emerging market equity asset class.

An asset class's return generating process can be broken down into two parts. The first is common to all the securities that make up the asset class. It is often called the "systematic" or "non-diversifiable" return on the asset class. The second return source is either unique to a specific company, or common only to a subset of companies within the overall asset class (e.g., companies in the energy sector). Here is a simple example. Consider an asset class made up of only two securities, which have equal weightings on all possible measures (e.g., their market capitalization, their book value, their sales, etc.). The return on security A is 7%; the return on security B is 3%. The average return is 5%, which represents the systematic return on the asset class, which would be received by an investor who owned both A and B. The unique return on Security A is 2%, and on B it is (2%).

This simple example illustrates a number of critical points. First, at the asset class level, the unique returns (also called “alpha” returns) cancel each other out, leaving only the systematic return. Traditionally, this has been referred to as the “market” or “beta” return. Investing with the objective of earning only this broad asset class return should, in our view, be called “asset class investing”, “market investing”, “beta investing” or “passive investing.” As you can see, the distinguishing characteristic of the market return for an asset class is that it requires no ability or attempt to forecast A and B’s future returns. It simply seeks the return that comes from owning all the securities in the asset class.

Second, return is compensation for bearing risk. At the asset class level, you receive only systematic market return, which compensates you for bearing systematic market risk. This systematic return is composed of two parts: the risk free rate (which compensates you for deferring consumption) and an asset class risk premium. Most important, earning this asset class risk premium does not depend on skill.

A third insight from our simple example is that, when you take on additional unique risks that can be avoided through diversification, you may receive additional compensation in the form of positive alpha. In the short term, this can be due to either luck or skill. However, as you keep taking on unique risks over longer and longer periods, the probability increases that your average return will be either zero or (after your higher costs) negative – unless you have better than average forecasting skill, and can better (than the average active investor) distinguish between those unique risk exposures that will earn positive and negative alphas. This last point is an important one, and is often overlooked. Successful active investment doesn’t depend just on having some forecasting skill. Rather, it requires that your forecasting skill be superior to the average level possessed by the other active managers against whom you are competing. And that is a far higher bar than most people like to admit to themselves.

A fourth point is that forecasting skill must be based on either access to superior information and/or use of a superior model to make sense of publicly available information. It also must be based on the existence of a financial market that is not perfectly efficient. In an efficient market, skilled forecasting is impossible because market prices already incorporate all the public and private information available about a security, and the pricing insights generated by many different models. As we have noted many times in our writing, we believe that

financial markets are not always efficient. Rather, we see them as a complex adaptive system that, while strongly attracted to efficiency, seldom achieves it. Hence, we believe that skilled forecasting and successful active management that generate alpha are possible, though quite rare in practice (especially after the additional costs incurred are taken into account).

The fifth point from our simple example is that asset class or passive investing, as we define it is not quite the same as “indexing” which is often used as a synonym for it. In our example, we measured the performance of our two-security asset class by constructing an index, which put equal weights on A and B. The index return was equal to 5%. But suppose that A and B differed from each other in ways besides their return. For example, suppose A had higher revenues and book assets, but lower market capitalization than B. What weights would we then use to construct our index to measure the systematic return on our asset class? Reasonable people can and do disagree on the right answer to the question of how best to measure market returns. This is a critical point, because the use of two different benchmark indexes to measure market return will lead to two different estimates of the size of an active manager’s alpha. If that manager’s compensation is in any way tied to the amount of alpha he or she generates (as is usually the case), this creates a potential conflict of interest. An active manager has an incentive to choose a market benchmark that maximizes reported alpha, while the investor who is hiring the active manager has an incentive to use the most accurate market benchmark possible. As a practical matter, this benchmark decision is often made by a third party. In the case of institutional investors (e.g., pension plans), it is often made by a consultant hired by the plan sponsor. In the case of individual investors, it is often made by either a financial adviser or a rating service (e.g., Morningstar or Lipper). However, we note that the potential conflict of interest will still be present if the party making the benchmark decision (i.e., the consultant or rating service) derives any economic benefit from the generation of more alpha rather than less by active managers. In our view, this is often the case, with too little recognition of the conflicts of interest that are present.

With that introduction, let us move on to the market for managed investment products (e.g., hedge, mutual and exchange traded funds). Broadly speaking, these fall into three categories. Some products offer only systematic (beta) returns. Their objective is to provide a return equal to the average return on a broad asset class, such as domestic equity (however it is

measured by the index provider). Because capturing these returns requires no forecasting skill and only minimal trading, these asset class index funds charge very low expenses.

At the other end of the spectrum, some funds offer only unique (alpha) returns. An example of this type of fund is an “equity market neutral” hedge fund. The manager of such a fund attempts to do two things: (a) utilize his or her superior forecasting skill to identify securities and transactions that will produce positive alpha; and (b) use other transactions to eliminate the fund’s exposure to systematic (beta) returns. Because of the additional operations involved in comparison with a “beta only” fund, this “pure alpha” fund must charge higher expenses. Also, based on the assumption that it is easier to run a “beta only” fund than one that earns “pure alpha”, the manager of the alpha fund also expects to receive a larger portion of the fund’s returns, as compensation for the use of his or her relatively scarcer skill. We also note that you can construct an index to measure the average performance of all equity market neutral fund managers, even though there is no systematic (beta) return involved. As we said, beta investing is not the same thing as indexing.

Other products seek to provide investors with a bundled mix of beta and alpha returns. Most actively managed mutual funds are in this category. They buy and sell securities, but don’t eliminate their exposure to beta returns. To maximize their forecasting advantage, many active managers restrict their investing activities to a sub-segment of the broad asset class (e.g., small companies’ shares, or shares of healthcare companies), which provides a convenient basis for classifying these funds into sub-segments of the broad asset class.

Finally, in between “pure beta” and “pure alpha” funds lie products that use a low-cost indexed approach to track the performance of different sub-segments of the broad asset class, that are identified using a clear, publicly disclosed set of rules. For example, sector and style (e.g., small cap value) exchange traded funds are examples of these products, as are bond ETFs that track different maturity indexes. Clearly, because they hold portfolios of securities that differ from the composition of the overall asset class, the returns they produce are a type of alpha. Yet, because these approaches to earning alpha have become well-known and embodied in a rules-based index, they are, confusingly, often called “beta” or “factor beta.” (Alpha in beta clothing, if you will). As a result, the meaning of “alpha” has shrunk, and is now sometimes taken to include only active manager returns net of the return not only on the broad market but also on one or more sub-segment indexes (i.e., “factor betas”) that an investor can

buy for a relatively low price. To put it differently, “alpha” is now often taken to mean only an active manager’s gross return, less the return on the market and the return on relevant “factor betas” to which the active manager has decided to be exposed via the investments his or her fund makes.

This alpha arguably comes four sources. The first is market timing, or the skill to profitably switch between asset class or factor exposures. The second is security selection, which reflects superior skill in forecasting the returns of individual assets. The third is skill in profitably providing insurance to other investors, for example by selling them put options that limit their downside risk (but increasing the manager’s). This is theoretically an attractive source of alpha, as it is not an inherently zero sum game. The fourth source of alpha involves earning a fee for providing liquidity to other investors. Again, this has the advantage of not being an inherently zero sum game. However, liquidity within some asset classes (with equities in the lead) is rapidly migrating to “factor-beta” status.

In our view, the creation of “factor beta” index products has been both a blessing and a curse. On the one hand, they have made it possible to implement a wider range of forecasts at lower cost. On the other hand, they have probably created a dangerous amount of confusion in many investors’ minds. Too many people appear to be under the illusion that they can earn alpha over a long-term holding period simply by using these “factor beta” index funds to permanently tilt their portfolios one way or another. In a reasonably efficient market, this should be impossible. Rather than alpha, a long-term one way factor beta tilt (e.g., toward small cap value stocks) should produce either lower returns but with lower risk than the overall market, or higher returns with higher risk. To believe that it will produce positive alpha requires acceptance of two additional premises.

The first is that some investors will systematically, over long periods of time, and for one or many reasons, make valuation mistakes. There is some evidence that this may happen. For example, immediate liquidity needs will always force some investors to sell securities they know are undervalued. And some investors will, because of overconfidence or their use of a momentum strategy, tend to buy securities that are overvalued. However, the second premise is that there are permanent barriers that prevent other investors from arbitraging away most of the alpha that these mistakes are expected to produce, by buying (and bidding up the price of) the undervalued securities, and selling short the overvalued securities. The evidence suggests

that this premise is much weaker than the first one (see, for example, “The Limits of the Limits to Arbitrage” by Brav and Heaton). Moreover, if both these premises are true, historical data should show significant positive risk adjusted returns from permanently tilting one’s portfolio towards a sub-segment. But this is not what we find.

One way to measure the effectiveness of an active management strategy is by using something called the “Information Ratio.” To calculate this, you start with the return on the sub-segment tilt (e.g., the return on a small cap value index) and subtract from it the return on the broad asset class index. Over many periods, the average of this result is the active return on this strategy or its “gross alpha.” If you subtract the expenses you pay to the active manager from this, it is the “net alpha.” The next logical step is to relate this to the amount of risk that was taken to earn the alpha. This is measured by the standard deviation of the alphas, which is known as “active risk” or “tracking error.” The Information Ratio therefore measures the risk adjusted return of the active strategy, by dividing the active return (alpha) by the active risk (tracking error) that was taken on to earn it. Information ratios of .50 or more are generally considered excellent performance by an active manager (although this varies by asset class, with higher IRs generally needed for top quartile performance in asset classes where returns are more volatile).

The following table shows the annualized net alphas, tracking errors, and information ratios for four common sub-segment tilts over two different ten year periods, covering 1979 to 1988, and 1989 to 1998. All the data are in nominal terms, and are based on the Wilshire Indexes.

	Net Alpha (assumes 25 bp expenses)	Active Risk	Information Ratio
1979-1988			
Large Value	1.60%	7.00%	0.23
Large Growth	-0.85%	3.54%	-0.24
Small Value	6.16%	7.77%	0.79
Small Growth	1.31%	9.29%	0.14
1989-1998			
Large Value	-2.01%	5.70%	-0.35
Large Growth	3.06%	4.51%	0.68
Small Value	-3.70%	7.44%	-0.50
Small Growth	-3.25%	10.97%	-0.30

This table makes a central point: over the twenty years covered by the data, there was no “free lunch.” As is true of all active management returns, alpha could only have been earned through the use of superior forecasting skill, and not simply by a permanent tilt toward one or more sub-segments of the U.S. equity market.

Philosophically (and practically, if you are an active manager), “factor beta” index products also create a “where will this all end?” issue with respect to the morphing of "alpha" into "factor beta". In theory, there are multiple criteria (factors) that could be used to automatically divide the securities in a broad asset class into smaller sub-segments, whose average returns can then be measured by an index and termed a "factor beta." Why just stop with industry sectors and sub-sectors, company market capitalization, and ratios like market/book and price/earnings that are used to define “value” and “growth” categories? Why not use some measure of economic profits, or non-market capitalization measures of size, or the absolute amount of dividends paid, and create indexes and ETFs that track each of them? In fact, as further described below, exchange traded funds based on all of these concepts have recently either been registered or launched.

The clear implication is that in a world of digital information and low cost computing power, the number of possible categorization/segmentation schemes, and thus indexes and sources of “factor beta” is very large indeed. This logically suggests that many more factor-beta index tracking ETFs should be introduced in the future. What then, will “alpha” signify when this ETF innovation process has run its course? Logically, it will refer to unique returns that, to the extent they can be forecasted, require the use of (a) superior non-public (though not illegal) fundamental information about company or industry (e.g., research into changing customer needs, and which companies are best positioned to satisfy them); (b) superior insight into the future behavior of other investors; and/or (c) the use of a model that generates asset price forecasts from public information (about either fundamental information or investor behavior), whose assumptions are not made public and/or are constantly updated.

So, to sum up this section, investment economics has not changed. An asset's return generating process still have two parts: one systematic and one unique (and diversifiable). However, these basic ingredients have been repackaged and combined into a confusing range of investment products. Some of these offer systematic “market returns” on broad asset classes

(call this “classic beta”); some other products offer sources of unique (alpha) returns at the sub-segment level at a relatively cheap price (“factor beta”); and yet another set of products offers relatively expensive sources of unique returns, which are thankfully still called “alpha.”

Now let’s move from investments in general, to the exchange traded funds (ETF) market in particular. Again, we will start with some definitions. The first is between an open end investment company (e.g., an OEIC in Europe, or a mutual fund in the United States) and a closed end investment company (e.g., a unit trust in Europe, or a closed end fund or ETF in the United States). When a person invests in an open end fund, the fund issues new shares to the investor and receives his or her money in return. That money is then invested in securities issued by other companies. If the person wants to switch out of this investment, he or she sells his shares back to the fund, which redeems them. It raises the cash needed to pay the investor by selling some of the securities it owns. In sum, at an open end fund, the number of fund shares outstanding goes up and down as investors purchase and redeem them, which also causes fluctuations in the value of the securities owned by the fund.

In contrast, a closed end fund issues shares only once to investors, at the time of its initial public offering. It then uses these funds to make investments in securities issued by other companies. If an investor wants to sell his or her shares in a closed end fund, he or she can only do so by finding another investor willing to buy them – the closed end fund does not continuously issue and redeem shares. To facilitate the matching of buyers and sellers, closed end funds are listed on stock exchanges. Exchange traded fund are quite similar to closed end funds, and are also traded on a stock exchange. Historically, the main difference between them was that exchange traded funds tracked an index, while closed end funds were actively managed vehicles. However, as we shall soon see, this distinction is no longer true.

At the highest level, it is now possible for investors in many markets around the world to use exchange traded funds to gain indexed exposure to many, if not all, broadly defined asset classes. For example, the following table shows the different ETF products from companies like Barclays (iShares), Rydex, State Street (SPDRs and StreetTRACKS), and Vanguard that could be used to accomplish this in the United States:

Broad Asset Class	Exchange Traded Funds
Real Return Bonds	TIP
Domestic Investment Grade Bonds	AGG
Foreign Currency Bonds	No ETF. However, could use actively managed Closed End Funds (IMF or JGG), or ETFs that track foreign currencies (and earn local money market rates of interest), such as FXE, FXB, FXC, and FXF.
Domestic Commercial Property	VNQ, RWR, ICF, IYR
Foreign Commercial Property	No ETF yet, but many are in registration.
Commodities	Only ETF is DWC, which tracks the Deutsche Bank Liquid Commodity Index. However, could also use Barclays exchange traded notes (ETNs) to gain exposure to either the Dow Jones AIG Commodity Index (our preferred index for this asset class) or the Goldman Sachs Commodities Index.
Timber	No ETF
Domestic Equity	TMW, VTI, IWV, IYY, ISI
Foreign Equity	EFA
Emerging Equity	EEM, VWO
Equity Volatility	No ETF

There are also a wide (and growing) range of ETFs available that enable investors to take “factor beta” active management tilts (potentially involving both long and short positions) within most of these asset classes. For example, domestic equity investors can take size (large, mid, small, and microcap), style (growth and value), sector (e.g., energy) and sub-sector (e.g., oil and gas exploration) tilts. Using new exchange traded funds from ProShares, they can also make leveraged investments (whose returns are 1.5 to 2 times greater than the return on the underlying index), and take short positions (these funds’ return is the inverse of the return on the underlying index). Foreign equity investors can also take size, style and sector tilts, as well as regional and country tilts. Emerging markets equity investors can take regional and country tilts. Investment grade bond investors can take maturity and credit tilts. Commodity investors

can tilt toward oil, gold and silver. Other tilts (e.g., toward energy, metals and agricultural sub-sectors) may be on the way. We also understand that domestic property investors will soon be able to take sector tilts (e.g., industrial, office, retail, lodging and housing), while foreign property investor will be able to take regional tilts.

So far, so good. However, it gets more confusing. A growing number of new exchange traded funds are based on an indexing methodology that, in our opinion, essentially makes them functionally indistinguishable from actively managed mutual funds. In our view, the critical distinction is who is making the active management forecast and resulting investment decisions. In the case of the low-cost alpha products noted above, which offer a wide variety of sub-segment tilts, it is the investor who is making the forecast, deciding to tilt his or her investment within a given asset class, and then implementing this decision via an ETF that tracks a corresponding index. The new “active EFT” products operate in a very different manner. They typically track an “index” that includes securities selected on the basis of a model (i.e., an algorithm) designed to identify investments that will outperform a given benchmark. As the model operates over time, the securities included in the “index” also change. In our view, this is nothing more or less than a relatively low cost quantitative active management strategy cleverly placed in an index “wrapper” to enhance its customer appeal.

Probably the best know examples of this approach are the exchange traded funds that use something other than market capitalization to determine index weights. In Research Affiliates’ fundamental indexing methodology, a company’s weight is determined by a mix of its revenue, cash flow, book value and dividends. In Wisdom Tree’s dividend weighting methodology, a company’s weight is determined by the absolute size of the dividend it pays relative to other companies. These algorithms are quite straightforward, but are still designed to produce index returns that are higher than those on a comparable market capitalization weighted index.

PowerShares uses a more complicated multifactor screening model to identify the securities and their weights in the indices that underlie its exchange traded funds. Another example of this approach is the First Trust IPOX-100 ETF. It tracks the IPOX Composite Index, which is described as “a rules-based value-weighted index measuring the average performance of U.S. IPOs during their first 1000 trading days. Index constituents are selected

based on quantitative initial screens and range from large, mature companies and fast growing IPOs to IPOs underperforming the market.”

Many more of these “active index” exchange traded fund are on the way. For example, Claymore Investments has registered ETFs that will be based on a sector rotation strategy, purchases and sales of shares by corporate insiders, and shares with low analyst coverage. Similarly, First Trust has registered a new exchange traded fund that will invest in the forty companies (from with a group of 250 large companies) with the highest economic profit (their screen is based on the Deutsche Bank Cash Return on Invested Capital methodology, which is a variant of other residual income valuation techniques such as Stern Stewart’s Economic Value Added or Boston Consulting Group’s Cash Flow Return on Investment). And in Europe, ABN Amro has created products based on an index that is a rules-based approach to seasonal timing of equity markets. An interesting question is whether the returns from all these products will be referred to as new sources of “beta.” On the other hand, something tells us that nobody will end up calling a new product Deutsche Asset Management has registered a new source of beta. The DB Currency Index Value Fund, will, if it is approved by the Securities and Exchange Commission, essentially be an ETF that gives investors access to a hedge fund-type “pure alpha” currency strategy.

As with all active management methodologies that attempt to generate alpha through the use of a superior model, these new “active index” exchange traded fund products run two risks. The first is that the models they use will be copied by competitors, who will bid up the prices of the securities it identifies as undervalued, and thus reduce their expected returns. This risk seems particularly acute for those “active index” ETFs whose model assumptions are publicly available (like fundamental and dividend weighting). The second risk is that changes in the structure of the economy will invalidate the assumptions of the models that underlie these new active ETFs. Just because a model has worked in the past does not mean that it will work in the future.

In sum, the exchange traded fund market has come a long way from its origins, and now contains a growing mix of “beta” and “alpha plus beta” products. Their existence has also made it possible to create your own equity market neutral “pure alpha” product, for example by investing \$1,000 in a long-only “active index” product and another \$1,000 in an ETF whose returns are equal to the inverse of the relevant asset class or sub-segment index. Granted, using

two ETFs is not a very efficient way of doing this; the return would be higher if you could offset the market risk by selling futures instead (which raises the question of how much longer it will be until such products appear). But still, today's ETF market represents a great leap forward from a few short years ago. And it helps point the way to what may lie ahead.

In so far as all these new exchange traded fund products have expanded the choices available to well-informed investors, and enabled them to implement strategies at a lower cost than before, they represent a positive development. On the other hand, not all investors are well-informed, and most do not possess forecasting skills that are consistently superior to those of the average active manager (particularly after expenses, trading costs and taxes are taken into account). This seems especially true of individual investors. For example, in their recent paper "Do Noise Traders Move Markets?" Barber, Odean and Zhu found that "stocks heavily bought by individuals in one year underperform stocks heavily sold by 4.4% in the following year." In addition, study after study has found that overconfidence is a hallmark of human nature (e.g., see "Sensation Seeking, Overconfidence, and Trading Activity" by Grinblatt and Keloharju). In light of these findings, the expanded range of exchange traded fund products may also have created more potential for investor disappointment. Exciting though all these new ETFs may be, they have not changed the mathematical fact that for every positive alpha earned there is a negative one somewhere else. As always, the best advice for most investors is to proceed with caution into the brave new world of exchange traded funds, with a healthy respect for the immense difficulty of being a consistently successful active manager, and a prudent awareness of the fact that low cost asset class index funds will almost certainly outperform a rapidly rising percentage of active managers as the investor's time horizon lengthens.

Active Management and Risk Budgeting

Innovation in the markets for exchange traded, mutual and hedge funds is creating many new opportunities for investors. Increasingly, it is possible for them to separate broad market returns and systematic risk from unique returns and active risk. This raises the question of whether an investor should either actively manage some part of his or her portfolio, and/or allocate some portion of it to active management by others.

The basic argument in favor of active management rests on two premises: (1) financial markets are not always efficient, and (2) some people have the talent to predict their future outcomes with a degree of accuracy that exceeds simple luck. We believe that there is strong evidence in favor of both propositions.

Over many years, we have repeatedly stressed our view that financial markets are best seen as a complex adaptive system (CAS). The defining characteristic of such systems is that, while strongly attracted to equilibrium, they are seldom in it. In addition, in such systems, cause and effect are often widely separated and non-linear, causing their behavior to alternate between regimes of higher and lower predictability. The first reason to view the financial markets as a CAS is because they trade claims on the underlying real economy, which itself is a CAS. The best-known analysis of the latter is [The Economy as an Evolving Complex System](#), a two volume series by Arthur, Durlauf and Lane. The second reason is that the interaction of different types of investment strategy (e.g., passive indexing, trend-chasing/momentum, and fundamental value) in the financial markets also produces non-linear price changes and, potentially, additional inefficiencies. Some of the best analyses of this include, “The Price Dynamics of Common Trading Strategies” by Farmer and Joshi, “Agent Based Financial Markets” by Blake LeBaron, “Interacting Agents in Finance” by Cars Hommes, “The Case for Market Inefficiency” by Bird, He, Thosar and Woolley, “Market Mood, Adaptive Beliefs and Asset Price Dynamics” by Dieci, Foroni, Gardini and He, and “The Adaptive Markets Hypothesis” by Andrew Lo.

Finally, three additional factors contribute to the inefficiency of financial markets. The first is the uneven rate at which new information diffuses among a population of investors (e.g., see “Industry Information Diffusion and the Lead-Lag Effect in Stock Returns” by Kewei Hou). The second is the fact that human beings are not perfectly rational information processing machines, and instead often use biased thinking shortcuts and let emotions affect our decisions (e.g., see “A Survey of Behavioral Finance” by Barberis and Thaler, “From Efficient Market Theory to Behavioral Finance” by Robert Shiller, and “Affect and Financial Decision Making” by Richard Peterson. And the third factor is that not every trade is made with the objective of maximizing expected value. For example, some people sell securities they believe are undervalued because they need to raise money; in other cases, central banks

may engage in losing trades in pursuit of higher priority monetary or exchange rate policy goals.

However, it is not sufficient that financial markets are not always in equilibrium, and therefore, to some extent, inefficient (i.e., that traded financial assets are both over and undervalued). A belief in the potential efficacy of active management also requires that at least some investors be able to predict future asset prices with a degree of accuracy that exceeds simple luck. Obvious examples like Warren Buffett and Bill Miller show that it is possible to consistently demonstrate superior forecasting skill over long periods of time. However, a number of studies have concluded that skilled managers like Miller and Buffett are quite rare. Most of these studies look at the probability that a set of active manager returns could be produced by something other than luck. Technically, the dividing line between skill and luck is usually taken to be a “t-statistic” value of 2 or more, which signifies 95% confidence that the conclusion reached is accurate. For example, in “Do Hedge Funds Deliver Alpha? A Bayesian and Bootstrap Analysis”, Kasowski, Naik and Teo find that “the performance of the top hedge funds (ranked by the t-statistic of alpha) cannot be attributed to sample variability or luck alone. This is true across all fund investment categories, and is robust to controlling for incubation and backfill bias, short-term serial correlation in returns, and structural breaks.” Another example is “False Discoveries in Mutual Fund Performance: Measuring Luck in Estimated Alphas” by Barras, Scaillet and Wermers. They find that while the impact of luck is substantial, there are some funds whose superior performance is almost certainly due to manager skill. Kosowski, Timmermann, Wemers and White reach a similar conclusion in their paper “Can Mutual Fund Stars Really Pick Stocks?” Finally, in “Mutual Fund Performance: An Empirical Decomposition”, Russ Wermers makes the important point that the presence of skill does not guarantee success as an active manager. Rather, it must produce alphas that are large enough to compensate for the higher expenses and transaction costs associated with actively managed funds. Wermers goes on to show that not all skilled managers can pass this test.

Many other studies show why most investors are not consistently skilled active managers. For example, in their recent paper, “Do Noise Traders Move Markets?”, Barber, Odean and Zhu find that stocks sold by individual investors outperform (over the next year) those they buy (for more evidence, see also “Trading is Hazardous to Your Wealth” by Barber and Odean; “Once Burned, Twice Shy” by Odean, Srahhilevitz and Baber, and “Who Loses

from Trade?” by Barber, Lee, Liu, and Odean) . Many other papers have explored possible reasons why most individual investors are poor active managers. Possible causes include an excessive focus on securities that are in the news (see, for example, “All That Glitters” by Barber and Odean, and “Profiting From Predictability” by Seasholes and Wu, which shows how other investors profit from attention-driven buying), late receipt of information compared to institutional investors, overconfidence, the use of biased thinking shortcuts (“heuristics”) that reflect the need to conserve our limited cognitive resources (e.g., see “Losers, Winners, and Biased Trades” by Johnson, Tellis and Macinnis), narrow (“myopic”) framing of investment issues (e.g., in terms of securities or asset classes, rather than an overall portfolio, or in terms of performance versus peers, instead of a liability that needs to be funded), and the affect of emotions on decision making (e.g., losing money seems to hurt twice as much as making money feels good).

If an investor does not believe that he or she is a skilled active investment manager, the decision about whether to add active risk to his or her portfolio rests on three further premises: (1) it is possible to identify other managers who are skilled; (2) It is possible to access their funds; and (3) the costs they charge for their services, plus any taxes attributable to their trading, are less than the alpha they are expected to produce. In our view, at least for retail investors, these three hurdles are significantly higher than the first two.

Let’s start with the challenge of identifying a skilled active investment manager. As is true in similar decision situations, the objective is to identify pieces of evidence that have a high power to discriminate between skilled and unskilled managers and come from reliable sources. A typical starting point is a manager’s historical “information ratio” or “IR.” In the simplest case, to calculate this, you start with the active manager’s return and subtract from it the return on a broad asset class index (i.e., subtract the market return from the manager’s return). Over many periods, the average of this result is the active return on this strategy or its “gross alpha.” If you subtract the expenses you pay to the active manager from this, it is the “net alpha.” The next logical step is to relate this to the amount of risk that was taken to earn the alpha. This is measured by the standard deviation of the alphas, which is known as “active risk” or “tracking error.” The Information Ratio therefore measures the risk adjusted return of the active strategy, by dividing the active return (alpha) by the active risk (tracking error) that was taken on to earn it. Information ratios of .50 or more are generally considered excellent

performance by an active manager (although this varies by asset class; as a general rule, the higher the volatility of asset class or sub-segment returns, the higher the IR needed to be a top quartile active manager). Unfortunately, identifying skilled active managers is not as easy as simply identifying those with the highest IRs.

To begin with, it is clear that the choice of asset class benchmark affects a manager's estimated alpha and IR. For example, choosing the S&P 500 (which is an index of large capitalization stocks) to represent the market return can clearly distort the estimated alpha for an active manager who specializes in small capitalization value stocks. In addition, to the extent that the returns on either the market index or the active manager's portfolio are not normally distributed, the reported alpha will also be distorted. However, even when they are accurately calculated, historical information ratios seem to have relatively little predictive value, especially if they are positive.

One problem is that the historical track records of most active manager are too short to make their IRs statistically significant. For example, an IR of .50 is often used as a benchmark for top quartile active manager performance. However, in order to be 95% sure that a manager achieved this due to skill and not luck, you would need 16 years of annual data (statistically, this would achieve a T-Ratio of 2.0). This is an important point, because multiple studies have found that the impact of luck is substantial, particularly among funds with high positive reported alphas (e.g., see "False Discoveries in Mutual Fund Performance: Measuring Luck in Estimated Alphas" by Barras, Scaillet and Wermers, "Can Mutual Fund Stars Really Pick Stocks?", by Kosowski, Timmerman, Wermers and Whilte, and "Do Hedge Funds Deliver Alpha? A Bayesian and Bootstrap Analysis" by Kosowski, Naik and Teo).

Given the limitations of the investment ratio, other quantitative indicators of manager skill are also used. For example, many information vendors decompose manager alpha (and its associated active risk) into its component parts. These usually include size (e.g., small cap), style (e.g., value), sector (e.g., healthcare) and other factor tilts, market timing (e.g., switching tilts over time, or moving into and out of cash), and security selection (i.e., the specific securities purchased within the given size/style/sector universe). It is generally thought that a high proportion of alpha from security selection is a sign of higher skill, as is stability over time in the proportion of alpha (and active risk) from different sources. However, these conclusions have to be balanced by a recognition that even a skilled active manager's investment process

(just like a superior company's value creation process) will probably not work all the time; rather, it will tend to perform better (i.e., produce significant positive alphas) under some economic and market conditions, and worse under others (see, for example, "Time Changing Alpha and Active Fund Performance Evaluation" by Heaney, Hallahan, Joseve and Mitchell).

Given these limitations with quantitative indicators, the process of identifying skilled active managers must also incorporate qualitative factors. As you would expect, there is no shortage of ideas about what these should include. Logically, you would expect a qualitative evaluation of an active manager to cover a number of areas, including, for example:

1. What are the true sources of their reported historical alpha?
 - a. Is any of it due to exposure to the market returns of other asset classes (e.g., beta masquerading as alpha?)
 - b. How much is due to providing insurance and liquidity to other investors?
 - c. How much comes from market timing?
 - d. How much from well chosen factor exposures (e.g., to different countries or sectors)? And do these reflect forecasting decisions, or are they simply permanent tilts?
 - e. And how much from security selection?
2. How do they generate investment ideas?
 - a. What source of disequilibrium are they attempting to exploit?
 - b. How durable are they?
 - c. How does the manager identify the opportunities disequilibrium creates? (e.g., using what type of superior information and/or superior model?)
 - d. How broadly do they search for ideas?
3. How to they decide which ideas to act on?
 - a. What does the real decision process look like (most human decisions are not made on the basis of elaborate processes)? What are the five to seven key variables (the normal number in human decision making) that drive it?
 - b. To what extent is their decision process subject to bias?
4. How to they combine their best ideas into a portfolio?
 - a. How to they estimate expected risk and return?
 - b. What constraints to they face that could reduce alpha (e.g., inability to take short positions, or limits on acceptable deviation from a benchmark index)?
5. How to they implement their portfolio decisions?
 - a. How do they control their risk exposures?
 - i. E.g., in the case of managers earning alpha from the provision of insurance and liquidity to others?
 - b. Do they have an effective information systems and an efficient back-office operation?
 - c. What do they do to minimize transaction costs (both explicit, like commissions, and bid/ask spreads, and implicit like the impact of their trading on prices and the opportunity cost incurred while executing their trades?)

6. How do they monitor their alpha generation process, and decide when to sell and/or to reduce risk?
7. How do they incorporate feedback and learning, to improve their investment process?
8. How strong is their team?
 - a. How different are they as individuals? Theoretically, different perspectives should help generate alpha from different sources.
 - b. How well do they work together? The potential value of different perspectives will be lost if they cannot work well together.
 - c. How closely are their incentives aligned with those of investors in their fund?

As you can see from this list, identifying skilled active managers is a challenging process. Some institutional investors (e.g., pension and endowment funds) are clearly good at this task; however, given the amount of time and qualitative information involved, it is easy to conclude that it is beyond the reach of all but the most diligent and talented individual investors.

Because they cannot conduct the type of time consuming manager research undertaken by institutional investors, most individuals seeking skilled active managers for their investments will likely turn to a fund rating service like Morningstar. Will this help them achieve their goal? To be sure, Morningstar's new rating system (implemented in June, 2002) appears to work better than its previous one. As documented in a recent study ("Morningstar Mutual Funds Redux" by Morey and Gottesman), funds with more stars in 2002 had a strong probability of outperforming those rated with fewer stars over the subsequent three-year period. However, this study suffers from two significant limitations. First, saying that a fund was likely to outperform one with fewer stars is not the same as generating positive alpha and a high information ratio after sales loads, expenses, and (depending on the investor's situation) taxes. In fact, most of the funds in the study generated negative alphas after adjusting for sales loads and expenses. In addition, while they acknowledged the issue, Morey and Gottesman did not try to adjust their findings for the tendency of strong performing mutual fund managers to extend their performance by a strong tilt to momentum in subsequent years, as described by Wermers (in "Is Money Really 'Smart'?") and Grinblatt, Titman and Wermers (in "Momentum Investment Strategies, Portfolio Performance, and Herding: A Study of Mutual Fund Behavior"). In sum, individual investors simply do not have available today a rating system that enables them to quickly and easily inspect mutual funds' previous alphas, active risk, and information ratios. This is not to say that such systems are not possible. Indeed, a number of recent papers describe approaches to identifying skilled mutual fund managers that appear quite

promising (e.g., “Using Hierarchical Models to Shrink Alphas” by Geoffrey Friesen, “Bayesian Alphas and Mutual Fund Performance” by Busse and Irvine, “The Ratings Game” by Michael Stutzer, and “The Right Answer to the Wrong Question” by Harlow and Brown). It remains to be seen whether these advanced methodologies will ever come into widespread commercial use.

However, let us assume that you are a sophisticated individual, adviser, family office or institution, and, after conducting an in-depth investigation, believe you have identified skilled managers who can generate alpha and top-quartile information ratios in the future. Before deciding to allocate a portion of your portfolio to active risk, you still have two significant obstacles to overcome.

The first is gaining access to these manager’s funds. Some of the best performing mutual and hedge funds are closed to new investors. The logic for doing this was clearly described by Berk and Green, in their paper “Mutual Fund Flows and Performance in Rational Markets.” Lacking perfect information, investors will view high returns as indicators of manager skill, and make new investments in a top-performing fund. To the extent that the fund manager’s compensation is linked to total assets under management, his or her pay will increase. In other words, the inflow of funds will provide him or her with compensation for his or her superior skill (assuming the returns weren’t due to luck alone). Yet this same inflow of funds will make the manager’s work far more difficult, because it is much harder to identify large investment opportunities that promise high returns than it is small ones. Moreover, it is harder to execute large trades (without moving the price against you) than it is smaller ones. Berk and Green argue that over time, this process logically will result in higher compensation for skilled managers, without providing significantly higher returns for investors. But what happens if, to some extent, the manager’s compensation is also tied to the size of the alpha he or she generates? In this case, there is a rational trade-off to make between two sources of compensation: adding assets to the fund or limiting its size to maximize its potential alpha. The result of this choice is often high performing funds that are closed to new investors.

Nevertheless, let us assume that you are able to access a mutual or hedge fund that you expect will generate alpha and a high information ratio in the years ahead. There is still one more hurdle to cross. You also have to decide whether the fees and expenses charged by the fund in question (and taxes related to its trading, if your investment is held in a taxable account)

will be less than the alpha the fund is expected to produce. A number of analysts have raised serious questions about this issue. For example, in the hedge fund world, many investors have chosen to use “Funds of Funds” either to diversify their exposure to different strategies and/or to gain access to top quality managers. Yet a number of analysts have questioned whether the fees charged by these Funds of Funds, which come on top of the fees charged by the underlying hedge funds, more than offset the alpha the latter are expected to generate (see, for example, “Fees of Fees on Funds of Funds” by Brown, Goetzmann, and Liang).

The picture may be arguably worse in the mutual fund world. For example, in “Measuring the True Cost of Active Management by Mutual Funds,” Ross Miller started with the observation that most actively managed mutual funds are “long only” products that, in effect, combine market and active returns. Thus, the market risk portion of an actively managed mutual fund’s return could have been obtained (at a much cheaper cost) by buying an index product. Given this, Miller made two adjustments to active funds’ reported results. First, he subtracted the index fund’s performance from the mutual fund’s performance to derive the latter’s alpha. Second, he subtracted the index fund’s fees from the mutual fund’s fees to derive the true cost charged to produce that alpha. Miller found that “at the end of 2004, the mean “active expense ratio” [i.e., the true cost of generating the alpha] for the large-cap equity mutual funds tracked by Morningstar was 7% “ [of the fund’s assets]. This was “over six times their published expense ratio of 1.15%.” It was also arguably higher than the pricing charged by many hedge funds (2% of the assets under management, and 20% of the profits generated).

On the other hand, a more recent paper, “Poor Man’s Hedge Funds? Performance and Risk-Taking of Hedged Mutual Funds” by Agarwal, Boyson, and Naik provided somewhat more encouraging news. Its authors found that mutual funds that employ “hedge-fund” like strategies (e.g., long/short and market neutral investing) significantly outperform traditional long-only actively managed mutual funds, delivering substantially higher alphas for not much more in additional fees (note that we use some of these products to implement our model portfolios’ allocations to equity market neutral). In addition, as described in another article in this issue, there are also a growing number of low cost, actively managed exchange traded funds coming to market, which may soon include market neutral, “pure alpha” products. So there is some cause for hope that in the future individual investors will be better able identify well-managed active risk products, and access them at a reasonable cost.

At this point, we will assume that you have not only identified top quartile active managers whom you believe will generate alpha and high information ratios in the future, but that you can access their funds at an acceptable cost. The next question is the right way to add them into your portfolio, alongside asset class index funds designed to capture broad market returns. This brings us to the subject of what is known as “risk budgeting.” This term refers to an integrated process that encompasses three decisions:

- The total amount of risk to take in a portfolio;
- Its division between market risk and active risk;
- The optimal allocation of active risk between active managers.

Let’s look at each of these issues. Assuming a portfolio’s returns are normally distributed, the probability of a loss of more than a certain amount over a given period of time can be determined using its expected annual average return and its standard deviation. In a normal distribution, 95% of the outcomes lie above the amount equal to the return less 1.65 standard deviations. And 99% lie above return less 2.33 standard deviations. For example, let us assume that a portfolio composed solely of investments in broad asset class index products has an expected annual real return of 5.9% and a standard deviation of 6.6%. The 95% confidence limit for the maximum loss that could be sustained over a one year holding period can be obtained by subtracting 1.65 times the standard deviation from the expected return, or $5.9\% - (1.65 \times 6.6\%)$. Thus, one can be 95% confident (assuming the assumptions are accurate), that any annual real loss experienced by the portfolio will be less than or equal to about (5.0%). In other words, given my assumptions about the portfolio’s future risk and return, and that its returns are normally distributed, I am 95% confident that over any one year my portfolio will not lose more than about (5.0%), in real terms. Similarly, I am 99% confident that over any one year period it will not lose more than (9.5%), which is equal to the expected return less 2.33 standard deviations [$5.9\% - (2.33 \times 6.6\%)$]. This type of analysis is widely used by institutions, where it is known as “value at risk” or “VAR.” However, it is not without its limitations, and is at best an imperfect tool (but still a big improvement over what was used before it arrived).

For example, VAR calculations will be inaccurate when one or more of a portfolio’s

assets have returns that are not normally distributed, and or returns that are not independent (since serial correlation results in underestimated volatility). There is also the problem that volatility and correlation change over time, which makes VAR estimates over longer periods more problematic than over the short periods (e.g., one day to one month) for which this methodology is most commonly used. Another problem with VAR (which it has in common with portfolio optimization) is that historical risks and returns may be a poor predictor of the future. Closely related to this is the problem that future returns are much more difficult to predict than future risk, as measured by the standard deviation and correlation of returns. Hence, another way to perform VAR analysis is without returns. In this case, the 95% risk limit would be expressed solely in terms of standard deviation. For example, the 95% downside limit would be 1.65 times the portfolio standard deviation of 6.6%, or about (10.9%), and the 99% limit would be $2.33 \times 6.6\%$, or about (15.4%). Remember this 15.4 number. We will keep using it as our example progresses.

Now, suppose I thought that, given my capacity (and willingness) to bear risk, the 15.4% limit was too low, and that, in pursuit of higher returns, I was willing to take on risk up to a 99% confidence limit of 17%. How could I add this risk to my portfolio? Basically, I have two options, assuming that my current portfolio already reflects an optimal allocation of market risk, given my target return.

The first alternative would be to add leverage to the market return portfolio (i.e., borrow money and use it to add to my asset class index fund investments). Adding leverage boosts both expected return and expected risk. Leverage equal to 9.4% of the portfolio's value would increase its 99% confidence limit standard deviation from 15.4% to 17.0% [= $15.4/(1-0.094)$].

A second alternative would be to use actively managed investment products (i.e., pursue alpha) that would add active risk to my existing market risk exposure. Let's assume I choose the latter approach. How would I calculate how much active risk to add? If we assume that all sources of active risk (across all asset classes) are uncorrelated with the any of the market risks already present in the portfolio (a computationally convenient and therefore typical assumption, but usually not an accurate one – e.g., see “The Implications of Blending Specialist Active Equity Managers” by Gallagher and Gardner), then the answer is straightforward, but the math is a bit complicated. You basically solve for the amount of active risk that, when added to your existing market risk exposure, produces your target 99% confidence level risk of 17.0%

The math challenge comes from the fact that when combining different sources of risk you have to convert standard deviations into variances by squaring them first. Converting back to a standard deviation (the usual way we express risk) requires that you take the square root of the combined variance. That sounds painful, but it's really not in practice. Let's go back to our example. Our target risk of 17% is the square root of some number. That number equals 17^2 , or 289. We know that market risk will provide a substantial part of that 289. In fact, it will provide 15.4^2 , or 237. That means that active risk must provide $289-237$, or 52. The square root of 52 is about 7.2 (which means we are back in the world of standard deviation). However, this 7.2 represents the 99% confidence limit. To get back to the actual active risk limit, we have to divide 7.2 by 2.33 to get 3.1, which is the amount of active risk we need to add to our portfolio to achieve our target total risk.

Hopefully, that math wasn't too painful. The next question is how best to add this active risk to our portfolio. Should we give it all to one active manager, or divide it up? And if the latter, how should we go about it? We can treat this the same as any other portfolio optimization problem. Our challenge is to divide the 3.1% active risk between active managers to maximize their (and our) expected alpha. To do this, we need estimates of the active risk each manager typically takes on as part of their investment process (e.g., active equity managers typically use much more than active bond managers), the correlation of their returns with other sources of active risk, their expected information ratios (since alpha equals the IR times the active risk taken), and estimates of our own uncertainty about their skill (i.e. a confidence interval for the manager's future IR). Describing the active management challenge in these terms makes three points quite clear.

The first is that because correlation between alpha sources increases their effective active risk, uncorrelated sources of alpha are highly valuable to a portfolio. Unfortunately, most hedge fund strategies today are "directional", and net long (i.e., net buyers of) securities in one or more asset classes. This often creates positive correlations between these active strategies and market returns on one or more asset classes. In other words (and this is a critical point), relatively few "hedge fund" strategies still adhere to the original meaning of the term, which signified, in essence, "pure alpha" strategies that were not exposed to (i.e., "hedged against") market risks. Today, again speaking in broad terms, only three types of hedge fund strategy still usually meet this original test: equity market neutral, convertible arbitrage, and

fixed income arbitrage. However, since the latter two typically use substantial amounts of leverage, we use only the first – equity market neutral – in our model portfolios.

Moreover, active strategies often turn out to have positive correlations with each other (as anyone who has seen the recent common trends in hedge fund returns can attest). The typical risk budgeting assumption that active strategy alphas are effectively uncorrelated is a dangerous illusion that could one day come back to haunt a lot of investors, just as it did in the case of Long Term Capital Management (which failed spectacularly in 1998 when its allegedly uncorrelated alpha strategies turned out to be anything but when global fixed income markets encountered a series of negative surprises).

The second point is that, because the future market returns on broad asset classes are relatively easier to predict than future returns (alphas) on active management strategies (which are, after all, a zero sum game), it is not surprising that most investors take on far less active risk than market risk in their portfolios. The third point is that the risk budgeting perspective tends to highlight the inefficiency of the traditional approach to portfolio management, which, after determining asset class weights, seeks to identify superior active managers (whose alphas are often correlated) to implement it. This can lead to larger than intended active risk exposures. Moreover, by forcing active managers into traditional asset class categories, it can limit investor access to many active strategies (e.g., portable alpha) that don't neatly fit into an asset class box.

Besides enhancing an investor's ability to manage his or her use of active risk, risk budgeting also helps to highlight potential inconsistencies within portfolios. Specifically, it allows an investor to compare the share of total portfolio return and portfolio risk that is coming from different asset classes and active risk exposures, given a set of portfolio weights. For example, the following table is based on an equal allocation to eleven asset classes plus the equity market neutral strategy, and historical U.S. dollar real returns data between 1994 and 2004. Note how different investments' share of total portfolio expected return and risk differ from their equal 8.33% asset weights.

	Share of Total Expected Portfolio Return	Share of Total Expected Portfolio Risk	Share of Return Less Share of Risk
Real Bonds	4.9%	1.3%	3.5%
Dom. Bonds	4.8%	0.8%	4.0%
Foreign Bonds	6.2%	5.5%	0.7%
Dom Property	14.2%	7.5%	6.7%
Foreign Property	6.4%	17.3%	-10.9%
Commodities	7.9%	3.1%	4.8%
Timber	7.4%	1.4%	6.1%
Dom. Equity	11.6%	8.7%	3.0%
Foreign Equity	5.8%	13.1%	-7.3%
Emg. Equity	4.0%	18.9%	-14.9%
Equity Volatility	18.0%	22.2%	-4.2%
Equity Market Neutral	8.7%	0.2%	8.6%
	100.0%	100.0%	0.0%

In equilibrium (rarely achieved, but a good starting point nonetheless), an asset class's expected returns should be proportional to its expected risk. Logically, in pursuit of a given target return, a portfolio's weights should therefore produce equal proportions of expected return and risk from the different positions taken (e.g., from different asset classes). As you can see, that is clearly not the case in our equally weighted portfolio. Because of the mismatch between their shares of total portfolio return and total portfolio risk, foreign property and emerging markets seem significantly overweight (with shares of return well below their shares of risk), while domestic property, timber, and equity market neutral seem significantly underweight. However, these conclusions critically depend on the historical risk, return and correlation assumptions used to reach them.

Logically, there are a number of possible explanations for the apparent portfolio mismatches, including assumptions about future levels of return, and changes in risk and correlation that differ from the historical data. For example, the allocation to foreign property could be justified in light of an expected increase in its future relative returns, and/or an expected decrease in its relative volatility and/or correlation with returns on other asset classes. Risk budgeting cannot tell you whether these assumptions are right or wrong. Rather, what it does quite well is highlight potential inconsistencies in a portfolio's structure, and challenge an investor to explain them. This cannot but help improve the investment management process,

and, one hopes, the results that are eventually realized.

So, to sum up, for some investors successful active management is clearly possible, despite that fact that on the whole it is a zero sum game (in fact, a negative sum game after active managers' expenses are taken into account). Complex adaptive systems like the real economy and financial markets are rarely in equilibrium, though they are strongly drawn toward it. This leads to incorrect valuation of financial assets, and the opportunity to earn alpha through the application of superior forecasting skills. It seems equally clear that at least some (but, relatively speaking not very many) investment managers possess this superior forecasting skill. Unfortunately, identifying these truly talented managers is anything but easy, and logically requires the expenditure of significant resources. In particular, individual investors do not have the benefit of the time and information that are available to institutional investors to help them identify skilled active managers.

Moreover, identifying skilled managers (with an acceptable degree of confidence in an inescapably uncertain conclusion) is only part of the battle. It is frequently difficult to access the funds they manage. In addition, there is also the risk that their fees will exceed the benefits they are expected to produce, either because their skills decline, and/or their success attracts so much money that their expected returns decline.

Still, there are bright spots, for those who (after adjusting for natural overconfidence) believe they have what it takes to succeed at active management. These include new techniques for identifying skilled managers, the encouraging risk adjusted performance of new mutual funds that use "hedge fund" type strategies, and the potential for more low cost market neutral ("uncorrelated alpha") exchange traded fund products. Finally, risk budgeting offers a powerful set of techniques that, while not yet widely available to individual investors (though Risk Metrics Inc. is making significant progress in this area), holds great promise for improving the discipline with which they (and their advisers) implement the investment management process.

Product and Strategy Notes

The IMF's Economic Worries

Based on the assumption that the subjects for their research papers are not simply drawn out of a hat, three recent working papers published by the IMF present an interesting picture of what may be on their mind these days. The first is “The Euro’s Challenge to the Dollar” by Ewe-Ghee Lim. It provides an excellent overview of the arguments in favor of and against the proposition that the Euro will account for an increasing percentage of the currencies in which financial assets are denominated. The argument in favor of a growing role for the Euro is based on the size of the region’s economy, and the obvious benefits of diversifying a portfolio’s currency exposures. The argument against a growing role for the Euro is based on the high costs associated with switching away from the dollar. For example, the price of many internationally traded goods are expressed in dollars today. The paper also notes another possible cause, which was clearly described in a 2005 paper by Chinn and Frankel (“Will the Euro Eventually Surpass the Dollar?”). This is the risk that “U.S. macroeconomic policy eventually undermines confidence in the dollar, in the form of inflation and depreciation.” In their paper, Frankel and Chinn emphasize that this change is likely to be non-linear, and occur more rapidly after a “tipping point” has been reached.

Perhaps not-so-coincidentally, the second IMF working paper is titled “Output Drops and the Shocks that Matter” by Becker and Mauro. They find that the largest drops are associated with external financing shocks (i.e., a sharp reduction in foreigners willingness to lend to a country) and terms of trade shocks (e.g., as are caused by a sharp rise in import or fall in export prices, like a spike in the price of oil). They also find that sharp “output drops are usually associated with major disruption for the residents of affected countries, both directly and often through ensuing, prolonged growth slowdowns.” Finally, rather than beating around the bush, the third working paper confronts the issue head on. In “How Might a Disorderly Resolution of Global Imbalances Affect Global Wealth?” Francis Warnock’s analysis finds “that were we to witness a simultaneous, unexpected 10 percent decrease in the U.S. dollar, U.S. equity markets, and dollar-denominated bonds, foreigners would, in sum, lose roughly \$1.2

trillion in foreign currency terms of financial wealth, an amount equivalent to almost 5 percentage points of non-U.S. GDP.” And that is based on just a 10 percent fall. The logical question to ask is whether that is a reasonable assumption.

A final working paper from the World Bank indirectly helps to answer that question. In “How Will China’s Savings-Investment Imbalance Evolve?”, Louis Kuijs makes an important distinction between the levels of savings and investment in China, and the balance between them. For example, surplus of savings over investment equal to 10 (which would cause a current account surplus of the same amount) could reflect savings of 80 and investment of 70, or savings of 35 and investment of 25. Kuijs provides an excellent analysis of why the levels of savings and investment are so high today in China. These include not only demographic factors (e.g., an aging population), but also policy decisions (e.g., high savings by public sector enterprises, owing to their reluctance to pay dividends). Kuijs shows how policy changes (e.g., a shift away from investment led growth, and more toward domestic consumption) could be used to reduce the levels of savings and investment. However, he also shows how these changes would not have a major impact on the savings/investment balance, and hence on China’s tendency to run substantial current account surpluses in the years ahead.

This brings us back to the IMF’s question of how we will eventually reduce the unprecedented savings/investment and current account imbalances that exist in the world economy today. The two largest counterparts to the United States yawning current account deficit are China’s and the oil exporting countries’ surpluses. Exchange rate changes alone (e.g., the depreciation of the U.S. dollar versus most other currencies) are incapable of reversing this situation, because the underlying trade flows are relatively insensitive to price. On the one hand, many companies have emphasized supply chain efficiency improvement, and extended significant parts of them into China. They oftent simply do not have alternative sources of supply. And on the other hand, this is also true of oil, the consumption of which has thus far also proven relatively insensitive to price increases. So, if price changes cannot, on their own, reduce global imbalances to a sustainable level, that leaves reductions in incomes to make up the difference. And this brings us to a critical uncertainty about which we have written over and over again: the political willingness and ability of the United States middle class to sustain a

prolonged real stagnation or decline in their incomes. If they resist this course of action, or, alternatively, if they accept it but it gets out of hand and triggers a classic debt contraction and deflationary depression, there is but one remaining choice: a defacto default by the United States on its external obligations, in the form of a prolonged period of very high inflation. Time will tell how this turns out. However, as we have repeatedly noted, given the highly leveraged state of American household balance sheets, and rising domestic and international political polarization, we would not bet against the high inflation scenario eventually occurring. And more and more, it seems the IMF is moving in this direction too.

Asset Classes and Return Generating Processes

The economic logic of asset class diversification is to reduce risk by gaining exposure to multiple return generating processes that have some (and hopefully a significant) degree of independence from each other. This raises a critical question: what evidence exists that this is what happens in practice?

The answer to this question requires that we accurately describe the return generating processes for different asset classes. In theory, each of these can be described as a model containing a number of variables (also known as "factors") and relationships between them and the return on the asset class in question. To the extent that the variables and/or relationships differ across asset classes, we can conclude that diversification results in exposure to different economic return generating processes.

So far, so good. However, as is true in so many cases in life, there is a catch. It turns out that accurately describing the return generating processes for different asset classes isn't as easy as you might think. The greatest amount of research may be found for the equity asset class. Three different approaches have been used.

The first uses macroeconomic variables to describe the return generating process. However, most of these models have ended up with less explanatory power than expected (in statistical terms, their regression models had low r-square coefficients). To be sure, a number of economic variables were found to be significant in terms of their impact on equity returns, such as inflation, real interest rates, exchange rates, and economic growth. However, these

studies have encountered two problems. First, in aggregate, their explanatory power isn't very strong (for recent reviews of this research, see "Economic Hedging Portfolios" by Van den Goorbergh, de Roon, and Werker, or "A Framework for Exploring the Macroeconomic Determinants of Systematic Risk" by Anderson, Bollerslev, Diebold and Wu). Second, because the economic variables themselves are not independent of each other (e.g., real interest rates interact with economic growth), it is hard to use them to distinguish one asset class's return generating process from another.

A second approach uses variables (factors) derived from the equity market itself to explain individual stock returns (and, in a few studies, bond default premia too). The best known example of this approach are the so-called "Fama French Carhart" factors, which, in addition to the return on the market itself, also use premiums for small sized firms (technically, the factor is equal to the return on small cap firms less large cap firms), high book/market firms (sometimes known as the value factor), and momentum (which measures the difference in returns between firms whose prices have recently risen and those whose prices have fallen). However, as many commentators point out, while this approach has greater statistical explanatory power (compared to macroeconomic factors), because it is based on factors derived from market returns, it doesn't get at the underlying drivers of the return generating process. Hence, many papers have been written about possible relationships between the Fama French Carhart factors and their possible economic drivers. For example, in "Macroeconomic Risks and the Fama and French/Carhart Model", Aretz, Bartram, and Pope find that these factors aggregate and summarize a larger number of macroeconomic variables. However, once again, we are left with the problem that (a) the FFC factors have not been applied to a wide range of asset classes, and (b) the underlying macroeconomic variables are not independent.

The shortcomings of the first two approaches have led to the use of a third one, derived from signal processing theory. This uses advanced quantitative methods to decompose a stream of information (be it noise or the returns on a group of asset classes) into underlying explanatory variables that are, by definition, statistically independent of each other. The two most common techniques in this area are Principal Components Analysis and Independent Components Analysis. The main difference between them is that PCA assumes that the signals it is analyzing (say, the return on an asset class) are normally distributed, while ICA does not make this assumption. Both usually assume that the signal (e.g., the return on an asset

class), can be described in terms of a linear combination of the independent statistical factors (e.g., .2 times factor A plus .4 times factor B). The main advantages of this approach is that it automatically decomposes the return generating process for different asset classes into statistically independent variables, and calculates the extent to which each variable can explain the variation in asset class returns. Hence, it can answer our initial question about whether investing in different asset classes exposes us to different underlying return generating processes. The main disadvantage is that it is not automatically clear what these independent variables mean in "real world" terms. Generating these interpretations is still as much art as science.

Given the strengths and weaknesses of the different approaches, we performed a Principal Components Analysis (PCA) on sixteen years of U.S. dollar real return data (1989 to 2004) for twelve asset classes to investigate the extent to which they provide exposure to different underlying return generating processes. The twelve asset classes included real return bonds, domestic investment grade bonds, high yield bonds, foreign currency bonds, commodities (as measured by the Dow Jones AIG Index), timber, domestic equity, foreign equity, emerging equity and expected equity market volatility, as measured by the VIX index.

Our PCA found that five variables explained 80% of the variation in the real returns on these twelve asset classes over the period we studied. The first variable explained 34%, the second 17%, the third 12%, the fourth 9% and the fifth 8%. Seven more independent variables explained the remaining 20% of the variation in asset class returns. Moreover, and this is critical, we found that different asset classes have different exposures to each of these independent variables. This is evidenced by the correlations between the returns on different asset classes and changes in the values of different variables, as shown in the following table:

	Variable 1	Variable 2	Variable 3	Variable 4	Variable 5
Pct. of Asset Class Returns Explained	34%	17%	12%	9%	8%
Real Bonds	.17	(.47)	(.26)	(.17)	.16
Dom. Bonds	.11	(.53)	.29	(.25)	(.23)
HY Bonds	(.33)	(.08)	.14	(.48)	(.36)
For. Bonds	.04	(.54)	.26	.36	(.02)
Dom. Prop.	(.27)	(.25)	(.24)	(.33)	(.13)

	Variable 1	Variable 2	Variable 3	Variable 4	Variable 5
For. Prop.	(.34)	(.27)	(.13)	.37	.07
Commod.	.00	(.18)	(.67)	(.13)	.32
Timber	(.04)	(.02)	.43	(.36)	.79
Dom. Eq	(.44)	.02	.19	.02	.01
For. Eq	(.41)	(.14)	.08	.32	.21
Emerg Eq.	(.41)	.04	(.10)	.09	.00
Volatility	.36	(.15)	.09	.23	(.02)

To be sure, in terms of their correlations with the independent variables, some asset classes are more similar than others. For example, real return bonds and domestic bonds; domestic and foreign property, commodities and timber, and domestic and foreign equity. But even in these cases, there are significant differences.

However, as noted above, interpreting the meaning of the independent variables identified by the Principal Components Analysis is something of an art. That being said, the data shown in this table is suggestive of possible explanations. The first variable is strongly positively related to the returns on volatility, real bonds, and domestic bonds, and negatively related to returns on other asset classes, especially the more risky ones. In our view, this variable seems to capture changing market psychology, which, as you can see, have a powerful impact on asset class returns. However, we cannot say for sure whether the variable reflects changes in the level of market uncertainty perceived by investors, or in their aversion to risk (for more on this, see “Risk, Uncertainty and Asset Prices” by Bekaert, Engstrom, and Xing). The second variable has relatively strong negative relationships with real return bonds, domestic bonds, and foreign bonds. It is somewhat less negatively related to returns on domestic and foreign property. From this we infer that this is variable related to variation in the real rate of interest.

The third independent variable is harder to interpret. Our inference is that it may have something to do with unexpected disinflation or deflation, which would logically boost returns on domestic and foreign nominal return bonds, and timber, while causing returns on commodities, domestic property and real return bonds to fall.

The fourth variable has a strong positive relationship with the returns on all three foreign asset classes, and a negative one with the return on high yield bonds, and to a lesser extent domestic property, timber and commodities (which are both priced in dollars) and domestic bonds. In our view, this variable is related to changes in the dollar exchange rate.

The fifth variable, with positive loadings on timber, commodities and real return bond returns, and a negative loading on domestic investment grade and high yield bonds, seems to be related to changes in inflation rates.

While they explain less of the overall variability in the returns on the twelve asset classes, the remaining seven independent variables (which aren't included in the table above) also have some interesting loadings on different asset classes. In particular, we find high loadings on emerging market equities, domestic property and foreign property. In sum, based on this analysis, we conclude that, as expected, diversification across broadly defined asset classes reduces an investor's risk by providing exposure to a number of different underlying return generating processes.

Finally, this analysis raises an interesting question about the future direction in which the investment management industry is likely to evolve. The expected and historical returns on any investment can be expressed as premiums for bearing different types of risk. Today, that risk is usually decomposed into market risk associated with different asset classes, and different types of unique risks within them (e.g., based on exposure to an industry sector, or a country, or large cap stocks or high risk bonds). The analysis in this article suggests that in the future, market risk exposure defined in terms of asset classes may give way to a definition that is based on variables similar to the ones identified in our principal components analysis, provided that investments can be identified whose returns mimic the changes in the independent variables. Clearly, we're not there yet. However, we strongly suspect that this is what lies ahead.

It's Time for a Canadian Revolt

At least against high mutual fund fees. That is the main conclusion that jumps from the pages of a fascinating new working paper by Khorana, Servaes and Tufano. In "Mutual Fund Fees Around the World," they report on a 2002 cost comparison across 47,000 funds offered in 18 different countries. The following tables show the average fund total expense ratio (as a percentage of assets under management) and the average total shareholder charges (which amortizes front and back end sales loads over a five year holding period).

Bond Funds

	Total Expense Ratio	Total Shareholder Charges
Australia	1.25%	1.64%
Canada	2.25%	4.10%
France	.86%	1.23%
Germany	.82%	1.34%
Switzerland	.92%	1.40%
United Kingdom	.97%	1.86%
United States	1.15%	1.41%
Global Avg.	1.28%	1.52%

Equity Funds

	Total Expense Ratio	Total Shareholder Charges
Australia	1.70%	2.16%
Canada	2.87%	4.93%
France	1.45%	2.12%
Germany	1.41% %	2.22%
Switzerland	1.48% %	2.15%
United Kingdom	1.39%	2.43%
United States	1.71%	1.99%
Global Avg.	1.87%	2.26%

One last statistic from the study absolutely bowled us over. Part of a fund's total expense ratio reflects the amounts paid to its investment manager (the remainder covers marketing, administration, and auditing charges etc.). Globally, these averaged 1.03%. Not much, right? However, applied to the total assets under management, they added up to U.S. \$63.6 billion. That is an absolutely stunning amount, especially when you consider how few of the funds in question will outperform a comparable index fund over a five-year holding period.

How Are Benchmark Indexes Chosen?

With so many indexes available today, how do investors choose between them? Depending on the purpose for which the index is being used, many different factors can come into play.

Perhaps the simplest thing to be said is that index selection is a complex process. To begin with, there are multiple “buyers” involved. Let’s consider two examples, one involving institutional investors, and the other retail investors. In the case of the former, three different parties are typically involved. The first is the party (e.g., a pension plan sponsor) that typically has a fiduciary obligation to see that funds are prudently invested. The second is the consultant whom the sponsor often hires to provide advice on the setting of investment objectives, the establishment of a risk budget and asset allocation policy, and how to select one or more passive and active managers to implement it. And the third party is/are these managers. In the case of a retail investor, the roles are quite similar, with an investor paying a financial adviser for asset allocation and fund selection advice, and various fund providers seeking that business.

Let’s look at the different roles indexes play in the institutional process. To begin with, the plan sponsor and his or her consultant will want to have good data available about the historical market return and risk on different asset classes, and on the relationships (e.g., correlations) between them. This argues for using as broad an index as possible, to accurately capture the true market return on an asset class. It also argues for the use of indexes that are based on fully transparent and easily understood methodologies. However, there may still be questions about how best to weight a broad index. For example, we have frequently noted that the most stable source of long-term commodity returns seems to be related to the benefits of holding futures contracts that are well diversified across energy, metals, and agricultural products. This argues in favor of the Dow Jones AIG Commodity Index, and against indexes like the Goldman Sachs Commodity Index in which energy products have a very high weight. Similarly, we have also noted (in our December, 2004 issue) that strong arguments can be made against the use of market capitalization weighting for bond indexes, and in favor of equally weighted indexes that cover a grid of maturity (duration) and risk parameters. Closely related to the weighting issue is what is known as “free float adjustment.” This refers to the fact that some shares that have been issued by companies are not freely traded (e.g., they are owned by an alliance partner). If these shares were included in an index, it would make it difficult for index fund managers to replicate it. Hence, most indexes remove these shares from their calculation, and include only the “free float” – that is, the proportion of a company’s shares that are frequently traded in the public market.

When it comes to establishing a risk budget (and evaluating active managers), it is also

important to have an index that allows a consistent approach to calculating the impact of different “factor betas” (e.g., tilts toward small cap, illiquid, and/or momentum stocks, long maturity bonds, or energy commodities) to better identify manager alpha that is created through skillful (or lucky) market timing and security selection. Hence, the availability of these sub-indexes is important, as is their alignment with the strategies used by active managers. For example, use of an overly narrow definition of “small cap value” stocks may lead to overstated alphas if a manager routinely holds stocks not included in the universe covered by the index (the differences in returns and risks between two U.S. small cap indexes, the S&P 600 and the Russell 2000, are often cited as a classic example of this problem). In addition, as hedge funds employ ever more complex active strategies, the need for consistent factor-betas across asset classes (e.g., liquidity) is becoming more important, to ensure accurate measure of exposures and calculation of manager alpha and risk. Looked at this way, it is easy to see how an index selection decision can raise conflict of interest issues where managers (and sometimes consultants) receive performance fees from the plan sponsor that are tied to the alpha they generate. And the same is true for a financial planner making actively managed fund recommendations to a retail investor. Clearly, the choice of benchmark index for the calculation of fund alphas will have a strong impact on the investor’s perception of the adviser. It is not hard to see the potential for conflict of interest here too.

Needless to say, the companies that manage index funds are not silent on the issue of index choice. However, they bring another set of issues to the table. One of them is the trade-off between coverage and the cost of creating a fund based on an index. While broad indexes create a more accurate picture of market risk, they can also be expensive to track, since they include many smaller and less liquid issues. To some extent, this can be overcome through the use of statistical sampling techniques that allow the broad index returns to be replicated without having to own all the constituent securities. However, the extent to which an index is “investable” at a reasonable cost will still come up. Closely related to this is the frequency at which the index is reviewed, and securities are added to and subtracted from it. If this happens quite frequently, and involves a large number of transactions, the costs will mount. As a result, more and more indexes have begun to include so-called “buffer zones” that include securities that have moved just outside an index’s official inclusion boundaries. Because many of these companies later move back within the index’s boundaries, letting them stay in the index helps

minimize transaction costs. Finally, at the margin, passive fund managers also prefer to create products that track indexes that are well known and trusted by their target customers. For some sophisticated institutional investors, this may be a minor consideration; however, for some retail investors, it may be much more important.

In addition to these considerations, both passive and active managers also prefer indexes which quickly provide accurate data about their constituent securities, including fundamental information such as their index weighting, price/earnings ratios, dividend yields and the like. This information is useful not only to investors in products that track the index, but also to active managers who are trying to outperform it. For this same reason, active managers also prefer indexes for which substantial historical data is available, and ones on which derivative instruments (e.g., futures, swaps and options) are traded with sufficient liquidity to support their strategies at a reasonable cost. In sum, choosing between indexes is often a complex process that involves many trade-offs, many of which are invisible to plan sponsors and individual investors.

2006-2007 Model Portfolios Update

Our model portfolios are constructed using a simulation optimization methodology. They assume that an investor understands the long-term compound real rate of return he or she needs to earn on his or her portfolio to achieve his or her long-term financial goals. We use SO to develop multi-period asset allocation solutions that are “robust”. They are intended to maximize the probability of achieving an investor’s compound annual return target under a wide range of possible future asset class return scenarios. More information about the SO methodology is available on our website. Using this approach, we produce model portfolios for six different compound annual real return targets: 7%, 6%, 5%, 4%, 3%, and 2%. We produce two sets of these portfolios: one assumes only investments in broad asset class index funds. These are our “all beta” portfolios. The second set of model portfolios includes equity market neutral (uncorrelated alpha) funds as a possible investment. These assume that an investor is primarily investing in index funds, but is willing to allocate up to ten percent of his or her portfolio to equity market neutral investments.

We use two benchmarks to measure the performance of our model portfolios. The first is cash, which we define as the yield on a one year government security purchased on the last trading day of the previous year. For 2006, our U.S. cash benchmark is 4.40% (in nominal terms). The second benchmark we use is a portfolio equally allocated between the ten asset classes we use (it does not include equity market neutral). This portfolio assumes that an investor believes it is not possible to forecast the risk or return of any asset class. While we disagree with that assumption, it is an intellectually honest benchmark for our model portfolios’ results.

The year-to-date nominal returns for all these model portfolios are shown in the tables on the following pages. Mutual and exchange traded funds that can be used to implement these model portfolios’ asset allocations are listed on our website.

Model Portfolios Year-to-Date Performance

<i>These portfolios seek to maximize the probability of achieving at least the target real return over twenty years, at the lowest possible risk.</i>			
	YTD 31Jul06	Weight	Weighted Return
	In US\$		In US\$
7% Target Real Return	<i>YTD Returns are Nominal</i>		
<i>Asset Classes</i>			
Real Return Bonds	-0.1%	0.0%	0.0%
U.S. Bonds	0.4%	0.0%	0.0%
Non-U.S. Bonds	4.5%	10.0%	0.5%
Domestic Commercial Property	17.2%	0.0%	0.0%
Foreign Commercial Property	15.7%	5.0%	0.8%
Commodities	3.5%	17.5%	0.6%
Timber	-1.6%	2.5%	0.0%
U.S. Equity	3.1%	55.0%	1.7%
Foreign Equity (EAFE)	11.5%	5.0%	0.6%
Emerging Mkt. Equity	7.4%	5.0%	0.4%
Equity Market Neutral	5.3%	0.0%	0.0%
		100.0%	4.5%

	YTD 31Jul06	Weight	Weighted Return
	In US\$		In US\$
6% Target Real Return	<i>YTD Returns are Nominal</i>		
<i>Asset Classes</i>			
Real Return Bonds	-0.1%	2.5%	0.0%
U.S. Bonds	0.4%	0.0%	0.0%
Non-U.S. Bonds	4.5%	10.0%	0.5%
Domestic Commercial Property	17.2%	0.0%	0.0%
Foreign Commercial Property	15.7%	0.0%	0.0%
Commodities	3.5%	15.0%	0.5%
Timber	-1.6%	5.0%	-0.1%
U.S. Equity	3.1%	45.0%	1.4%
Foreign Equity (EAFE)	11.5%	15.0%	1.7%
Emerging Mkt. Equity	7.4%	7.5%	0.6%
Equity Market Neutral	5.3%	0.0%	0.0%
		100.0%	4.6%

	YTD 31Jul06	Weight	Weighted Return
	In US\$		In US\$
5% Target Real Return	<i>YTD Returns are Nominal</i>		
<i>Asset Classes</i>			
Real Return Bonds	-0.1%	2.5%	0.0%
U.S. Bonds	0.4%	27.5%	0.1%
Non-U.S. Bonds	4.5%	12.5%	0.6%
Domestic Commercial Property	17.2%	0.0%	0.0%
Foreign Commercial Property	15.7%	0.0%	0.0%
Commodities	3.5%	15.0%	0.5%
Timber	-1.6%	5.0%	-0.1%
U.S. Equity	3.1%	17.5%	0.5%
Foreign Equity (EAFE)	11.5%	10.0%	1.2%
Emerging Mkt. Equity	7.4%	10.0%	0.7%
Equity Market Neutral	5.3%	0.0%	0.0%
		100.0%	3.5%

	YTD 31Jul06	Weight	Weighted Return
	In US\$		In US\$
4% Target Real Return	<i>YTD Returns are Nominal</i>		
<i>Asset Classes</i>			
Real Return Bonds	-0.1%	5.0%	0.0%
U.S. Bonds	0.4%	32.5%	0.1%
Non-U.S. Bonds	4.5%	15.0%	0.7%
Domestic Commercial Property	17.2%	5.0%	0.9%
Foreign Commercial Property	15.7%	0.0%	0.0%
Commodities	3.5%	10.0%	0.4%
Timber	-1.6%	10.0%	-0.2%
U.S. Equity	3.1%	0.0%	0.0%
Foreign Equity (EAFE)	11.5%	15.0%	1.7%
Emerging Mkt. Equity	7.4%	7.5%	0.6%
Equity Market Neutral	5.3%	0.0%	0.0%
		100.0%	4.1%

	YTD 31Jul06	Weight	Weighted Return
	In US\$		In US\$
3% Target Real Return	<i>YTD Returns are Nominal</i>		
<i>Asset Classes</i>			
Real Return Bonds	-0.1%	12.5%	0.0%
U.S. Bonds	0.4%	42.5%	0.2%
Non-U.S. Bonds	4.5%	10.0%	0.5%
Domestic Commercial Property	17.2%	0.0%	0.0%
Foreign Commercial Property	15.7%	0.0%	0.0%
Commodities	3.5%	12.5%	0.4%
Timber	-1.6%	0.0%	0.0%
U.S. Equity	3.1%	5.0%	0.2%
Foreign Equity (EAFE)	11.5%	10.0%	1.2%
Emerging Mkt. Equity	7.4%	7.5%	0.6%
Equity Market Neutral	5.3%	0.0%	0.0%
		100.0%	2.9%

	YTD 31Jul06	Weight	Weighted Return
	In US\$		In US\$
2% Target Real Return	<i>YTD Returns are Nominal</i>		
<i>Asset Classes</i>			
Real Return Bonds	-0.1%	22.5%	0.0%
U.S. Bonds	0.4%	32.5%	0.1%
Non-U.S. Bonds	4.5%	10.0%	0.5%
Domestic Commercial Property	17.2%	0.0%	0.0%
Foreign Commercial Property	15.7%	0.0%	0.0%
Commodities	3.5%	7.5%	0.3%
Timber	-1.6%	5.0%	-0.1%
U.S. Equity	3.1%	7.5%	0.2%
Foreign Equity (EAFE)	11.5%	10.0%	1.2%
Emerging Mkt. Equity	7.4%	5.0%	0.4%
Equity Market Neutral	5.3%	0.0%	0.0%
		100.0%	2.5%

	YTD 31Jul06	Weight	Weighted Return
	In US\$		In US\$
Equally Weighted Portfolio	<i>YTD Returns are Nominal</i>		
<i>Asset Classes</i>			
Real Return Bonds	-0.1%	10.0%	0.0%
U.S. Bonds	0.4%	10.0%	0.0%
Non-U.S. Bonds	4.5%	10.0%	0.5%
Domestic Commercial Property	17.2%	10.0%	1.7%
Foreign Commercial Property	15.7%	10.0%	1.6%
Commodities	3.5%	10.0%	0.4%
Timber	-1.6%	10.0%	-0.2%
U.S. Equity	3.1%	10.0%	0.3%
Foreign Equity (EAFE)	11.5%	10.0%	1.2%
Emerging Mkt. Equity	7.4%	10.0%	0.7%
		100.0%	6.2%

<i>These portfolios seek to maximize the probability of achieving at least the target real return over twenty years, at the lowest possible risk.</i>		<i>Unlike the other target return portfolios, these allow investment in uncorrelated alpha (equity market neutral) funds.</i>	
	YTD 31Jul06	Weight	Weighted Return
	In US\$		In US\$
7% Target Real Return	<i>YTD Returns are Nominal</i>		
<i>Asset Classes</i>			
Real Return Bonds	-0.1%	0.0%	0.0%
U.S. Bonds	0.4%	0.0%	0.0%
Non-U.S. Bonds	4.5%	0.0%	0.0%
Domestic Commercial Property	17.2%	0.0%	0.0%
Foreign Commercial Property	15.7%	7.5%	1.2%
Commodities	3.5%	12.5%	0.4%
Timber	-1.6%	7.5%	-0.1%
U.S. Equity	3.1%	52.5%	1.6%
Foreign Equity (EAFE)	11.5%	5.0%	0.6%
Emerging Mkt. Equity	7.4%	10.0%	0.7%
Equity Market Neutral	5.3%	5.0%	0.3%
		100.0%	4.7%

	YTD 31Jul06	Weight	Weighted Return
	In US\$		In US\$
6% Target Real Return	<i>YTD Returns are Nominal</i>		
<i>Asset Classes</i>			
Real Return Bonds	-0.1%	0.0%	0.0%
U.S. Bonds	0.4%	5.0%	0.0%
Non-U.S. Bonds	4.5%	7.5%	0.3%
Domestic Commercial Property	17.2%	0.0%	0.0%
Foreign Commercial Property	15.7%	0.0%	0.0%
Commodities	3.5%	10.0%	0.4%
Timber	-1.6%	10.0%	-0.2%
U.S. Equity	3.1%	37.5%	1.2%
Foreign Equity (EAFE)	11.5%	15.0%	1.7%
Emerging Mkt. Equity	7.4%	10.0%	0.7%
Equity Market Neutral	5.3%	5.0%	0.3%
		100.0%	4.4%

	YTD 31Jul06	Weight	Weighted Return
	In US\$		In US\$
5% Target Real Return	<i>YTD Returns are Nominal</i>		
<i>Asset Classes</i>			
Real Return Bonds	-0.1%	2.5%	0.0%
U.S. Bonds	0.4%	22.5%	0.1%
Non-U.S. Bonds	4.5%	10.0%	0.5%
Domestic Commercial Property	17.2%	2.5%	0.4%
Foreign Commercial Property	15.7%	0.0%	0.0%
Commodities	3.5%	15.0%	0.5%
Timber	-1.6%	5.0%	-0.1%
U.S. Equity	3.1%	12.5%	0.4%
Foreign Equity (EAFE)	11.5%	10.0%	1.2%
Emerging Mkt. Equity	7.4%	10.0%	0.7%
Equity Market Neutral	5.3%	10.0%	0.5%
		100.0%	4.2%

	YTD 31Jul06	Weight	Weighted Return
	In US\$		In US\$
4% Target Real Return	<i>YTD Returns are Nominal</i>		
<i>Asset Classes</i>			
Real Return Bonds	-0.1%	12.5%	0.0%
U.S. Bonds	0.4%	27.5%	0.1%
Non-U.S. Bonds	4.5%	0.0%	0.0%
Domestic Commercial Property	17.2%	0.0%	0.0%
Foreign Commercial Property	15.7%	2.5%	0.4%
Commodities	3.5%	12.5%	0.4%
Timber	-1.6%	7.5%	-0.1%
U.S. Equity	3.1%	5.0%	0.2%
Foreign Equity (EAFE)	11.5%	15.0%	1.7%
Emerging Mkt. Equity	7.4%	7.5%	0.6%
Equity Market Neutral	5.3%	10.0%	0.5%
		100.0%	3.8%

	YTD 31Jul06	Weight	Weighted Return
	In US\$		In US\$
3% Target Real Return	<i>YTD Returns are Nominal</i>		
<i>Asset Classes</i>			
Real Return Bonds	-0.1%	12.5%	0.0%
U.S. Bonds	0.4%	32.5%	0.1%
Non-U.S. Bonds	4.5%	12.5%	0.6%
Domestic Commercial Property	17.2%	0.0%	0.0%
Foreign Commercial Property	15.7%	0.0%	0.0%
Commodities	3.5%	10.0%	0.4%
Timber	-1.6%	5.0%	-0.1%
U.S. Equity	3.1%	5.0%	0.2%
Foreign Equity (EAFE)	11.5%	10.0%	1.2%
Emerging Mkt. Equity	7.4%	5.0%	0.4%
Equity Market Neutral	5.3%	7.5%	0.4%
		100.0%	3.0%

	YTD 31Jul06	Weight	Weighted Return
	In US\$		In US\$
2% Target Real Return	<i>YTD Returns are Nominal</i>		
<i>Asset Classes</i>			
Real Return Bonds	-0.1%	17.5%	0.0%
U.S. Bonds	0.4%	35.0%	0.1%
Non-U.S. Bonds	4.5%	12.5%	0.6%
Domestic Commercial Property	17.2%	0.0%	0.0%
Foreign Commercial Property	15.7%	0.0%	0.0%
Commodities	3.5%	7.5%	0.3%
Timber	-1.6%	5.0%	-0.1%
U.S. Equity	3.1%	5.0%	0.2%
Foreign Equity (EAFE)	11.5%	5.0%	0.6%
Emerging Mkt. Equity	7.4%	5.0%	0.4%
Equity Market Neutral	5.3%	7.5%	0.4%
		100.0%	2.4%