

The Index Investor

Why Pay More for Less?

Global Asset Class Returns

<i>Year to Date</i>	<u>In USD</u>	<u>In AUD</u>	<u>In CAD</u>	<u>In EURO</u>	<u>In JPY</u>	<u>In GBP</u>
US Equity	4.30%	(7.35%)	(5.39%)	(2.19%)	4.41%	4.88%
US Bonds	2.20%	(9.45%)	(7.49%)	(4.29%)	2.31%	2.78%
AUS Equity	12.70%	1.05%	3.01%	6.21%	12.81%	13.28%
AUS Bonds	10.68%	(0.97%)	0.99%	4.19%	10.79%	11.26%
CAN Equity	11.20%	(0.45%)	1.51%	4.71%	11.31%	11.78%
CAN Bonds	10.82%	(0.83%)	1.13%	4.33%	10.93%	11.40%
Euroland Equity	4.40%	(7.25%)	(5.29%)	(2.09%)	4.51%	4.98%
Euroland Bonds	8.20%	(3.45%)	(1.49%)	1.71%	8.31%	8.78%
Japan Equity	(7.80%)	(19.45%)	(17.49%)	(14.29%)	(7.69%)	(7.22%)
Japan Bonds	1.08%	(10.57%)	(8.61%)	(5.41%)	1.19%	1.66%
UK Equity	1.40%	(10.25%)	(8.29%)	(5.09%)	1.51%	1.98%
UK Bonds	0.64%	(11.01%)	(9.05%)	(5.85%)	0.75%	1.22%
World Equity	2.65%	(9.00%)	(7.04%)	(3.84%)	2.76%	3.23%
World Bonds	3.85%	(7.80%)	(5.84%)	(2.64%)	3.96%	4.43%
Commodities	(0.70%)	(12.35%)	(10.39%)	(7.19%)	(0.59%)	(0.12%)

Model Portfolio Update

The objective of our first set of model portfolios is to deliver higher returns than their respective benchmarks, while taking on no more risk. The benchmark for the first portfolio in this group is an aggressive mix of 80% domestic equities, and 20% domestic bonds. Through the end of April, this benchmark had returned 3.9%, while our model portfolio had returned 3.4%. We have also compared our model portfolios to a set of global benchmarks. In this

case, the global benchmark is a mix of 80% global equities, and 20% global bonds. Through the end of last month, it had returned 2.9%.

The benchmark for the second portfolio in this group is a mix of 60% domestic equities and 40% domestic bonds. Through the end of last month, it had returned 3.5%, while our model portfolio had returned 3.6%, and the global benchmark had returned 3.1%.

The benchmark for the third portfolio in this group is a conservative mix of 20% domestic equities and 80% domestic bonds. Through the end of last month, it had returned 2.6%, while our model portfolio had returned 2.8% and the global benchmark 3.6%.

The objective of our second set of model portfolios is to deliver less risk than their respective benchmarks, while delivering at least as much return. The benchmark for the first portfolio in this group is an aggressive mix of 80% domestic equities, and 20% domestic bonds. Through the end of last month, this benchmark had returned 3.9%, while our model portfolio had returned 3.6%. We have also compared our model portfolios to a set of global benchmarks. In this case, the global benchmark is a mix of 80% global equities, and 20% global bonds. Through the end of last month, it had returned 2.9%.

The benchmark for the second portfolio in this group is a mix of 60% domestic equities and 40% domestic bonds. Through the end of last month, it had returned 3.5%, while our model portfolio had returned 3.5%, and the global benchmark had returned 3.1%.

The benchmark for the third portfolio in this group is a conservative mix of 20% domestic equities and 80% domestic bonds. Through the end of last month, it had returned 2.6%, while our model portfolio had returned 2.9% and the global benchmark 3.6%.

The objective of our third set of model portfolios is not to outperform a benchmark index, but rather to deliver a minimum level of compound annual return over a ten-year period. Through last month, our 12% target return portfolio has returned 3.5% year-to-date, our 10% target

return portfolio has returned 3.2% our 8% target return portfolio has returned 3.4%, and our 6% target return portfolio has returned 2.7%.

Last month, the active portfolio was allocated as follows: 60% to the Vanguard Inflation Protected Securities Fund, 15% each to the Oppenheimer Real Assets Fund and the T. Rowe Price International Bond Fund, and 10% to the U.K. Equity Market iShare. These will not change next month. Year-to-date, our actively managed portfolio has returned 0.27%.

Equity Market Valuation Update

As we have previously noted, our valuation analysis rests on two fundamental assumptions: that over the long term, labor productivity growth in our six major regions will converge to between 2.5% and 3.5% per year, and that the long term real equity risk premium is 4.0% per year. Given those assumptions, here is our updated market valuation analysis at 30 April, 2003:

Country	Real Risk Free Rate	Equity Risk Premium	Required Real Return on Equities	Expected Real Growth Rate*	Div Yield	Expected Real Equity Return
Australia	2.97%	4.0%	6.97%	4.3%	3.80%	8.10%
Canada	3.08%	4.0%	7.08%	4.1%	2.10%	6.20%
Eurozone	1.97%	4.0%	5.97%	3.5%	3.30%	6.80%
Japan	1.62%	4.0%	5.62%	3.2%	1.20%	4.40%
U.K.	2.09%	4.0%	6.09%	3.5%	3.60%	7.10%
U.S.A.	2.78%	4.0%	6.78%	4.4%	1.80%	6.20%

**This reflects not only 3.5% productivity growth, but also expected labor force growth.*

Country	Implied Index Value*	Current Index Value	Current/Implied (productivity growth @3.5%)	Current/Implied (productivity growth at 2.5%)
Australia	324.91	228.29	70%	97%
Canada	158.34	224.69	142%	190%
Eurozone	158.63	118.73	75%	105%
Japan	34.59	69.75	202%	285%
U.K.	351.19	252.66	72%	100%
U.S.A.	283.36	374.67	132%	188%

* Assuming 3.5% future productivity growth

UK Recommends Sharp Cutbacks in Bundled and Soft Commissions

Last month, we wrote about the potential abuses that could occur under the current soft and bundled brokerage commission system in the United States. We showed how these practices could easily lead to excessive trading costs and lower returns for mutual fund investors. In a particularly well-timed follow up to this article, the Financial Services Authority in the United Kingdom has just released a major new report that recommends these practices be sharply curtailed.

Titled “Consultation Paper 176, “Bundled Brokerage and Soft Commission Arrangements, issued April, 2003”, the FSA report begins by noting that “Commissions paid by fund managers for the execution of trades on behalf of their customers [i.e., their investors] are a significant cost of investment fund management. In 2000, for example, UK fund managers paid about £2.3 billion in commissions from their customers’ funds to UK brokers. However, commissions may pay for more than the cost of trade execution. Estimates vary, but as much as 40% goes on additional services such as investment research and market information technology [eg., data feeds, Reuters and Bloomberg terminals], through “bundled” or “soft commission” arrangements...”

“In the institutional fund management market, the traditional model for commission payments is the full service broking arrangement between a broker and a fund manager. The broker receives ‘hard commission’ (i.e., an actual cash amount) from the fund manager at an agreed rate, on a transaction by transaction basis. In return, the broker supplies the fund manager with a package that includes both trade execution and other services, such as provision of research materials, or access to the broker’s own investment analysts. These other services are produced in-house, either by the broker or by another company in the same group, and cannot usually be purchased separately. Combining these additional services with trade execution in a single package and charging a single price for them is an example of the practice known as ‘bundling’.”

“Where a fund manager receives additional services from the broker, he is effectively getting a partial rebate on the commission he paid for the broker’s transaction service. Since, in most cases, no individual price is attributed to the constituent parts of the full broking service, their cost to the fund manager’s customers [that is, investors] is not transparent. The record of the transaction will show only the overall amount of the commission paid at the agreed upon rate.”

“Under the soft commission arrangements that are typical in the UK, a broker agrees to pay for certain goods and services that are supplied directly to the fund manager, usually by a third party. Between 50% to 57% of the funds generated from soft commission arrangements are spent on market information technology, including pricing and information services. The next largest categories of spending is research into companies and markets (25% - 30%), and computer hardware and software (14% - 17%)...”

“Bundling and softing...are a source of significant conflicts of interest, because they create incentives for fund managers to make trading decisions that do not necessarily serve the best interests of their customers...Because the costs paid for under bundling and softing arrangements are subject to less scrutiny by customers, fund managers may spend more than they need to on services...and have less incentive to ensure that they get best value for the money they spend...Bundling and softing [also] create an incentive for the fund manager to

undertake volumes of trading that may be motivated by a desire to obtain particular quantities of additional services, rather than to improve the performance of their customers' funds. [In addition], a fund manager's approach to broker selection may be unduly influenced by the nature of commission arrangements. As a result, the fund manager may select a brokers who offer generous bundling or softing terms but inferior execution quality, such as wider trading spreads or higher commission rates..."

"The control over these [potential conflicts of interest] exerted by normal market discipline is weak and uneven...Although fund managers are judged on fund performance, isolating the effect on performance of bundling and soft commission arrangements from any other factor is well-nigh impossible. More disclosure can help...However, we are not convinced that more disclosure [alone] will be enough to address the issues. In the retail funds market, we think it would have minimal impact..."

"Our general conclusion is that the economic benefits that are claimed for bundling and softing are unlikely to be realized, or may not be significant."

"We do not believe that our current regulatory regime addresses these issues satisfactorily, and conclude that further action is necessary. We are [proposing] two main measures. [First], we propose to limit the goods and services, beyond trade execution, that can be bought with commission or order flow. Specifically, we propose excluding market pricing and information services, such as dealing screens, which account for between 50% and 57% of soft commission credits. [Second], we propose that the cost of acquiring other services in a package along with trade execution should not be passed through automatically by a fund manager to customers' funds. This would apply in particular to the use of commissions to buy investment research. This does not mandate unbundling of [trade execution and research] by service providers, but focuses on greater transparency and accountability to fund management customers on the use of their funds...We propose that where a fund manager buys any other services in addition to trade execution with his customers' commissions, he should determine the cost of those services and rebate and equivalent amount to his customers' funds... The effect of this proposal is that the automatic pass-through of

commission to customers' funds would be limited to the cost of trade execution. [These] proposals will benefit consumers by ensuring that fund managers acting on their behalf have stronger incentives to obtain value for money, and to put in place trading arrangements that clearly operate in the best interests of retail fund owners.”

We believe the UK regulatory authorities are right on the money with their conclusions, and hope that their recommendations will be copied by regulators in other major markets.

SEC Settles Conflict of Interest Charges

On April 28, 2003, ten of the world's top investment firms agreed to a \$1.4 billion dollar settlement with securities regulators of charges involving conflicts of interest between their research and investment banking departments. As is typical of such settlements, the firms involved neither admitted nor denied their guilt; nevertheless, the nature of the settlement they agreed to makes it relatively easy to infer what the outcome would have been had they decided to let the case be decided in a court of law. Despite the settlement, it may still come to that, as the firms involved now face the prospect of thousands of suits by investors who believe they were misled by the investment research issued by these firms.

The firms involved in the settlement are a “who's who” of the world's major investment banking companies, and included Bear, Stearns, Credit Suisse First Boston, Goldman, Sachs, Lehman Brothers, J.P. Morgan, Merrill Lynch, Morgan Stanley, Citigroup/Salomon Smith Barney, UBS Warburg, and U.S. Bancorp Piper Jaffray.

In addition to the monetary payments, the firms are also required to comply with significant requirements that dramatically reform their business practices, including strengthening the separation between their research and investment banking departments.

As the SEC noted in announcing the settlement, “the enforcement actions allege that ... all of the firms engaged in acts and practices that created or maintained inappropriate influence by

investment banking over research analysts, thereby imposing conflicts of interest on research analysts that the firms failed to manage in an adequate and appropriate manner...The enforcement actions also allege that firms had issued fraudulent research reports, or research reports that did not provide a sound basis for evaluating facts, contained exaggerated or unwarranted claims about the covered companies, and/or contained opinions for which there were no reasonable bases.”

When the settlement was announced, SEC Chairman William Donaldson commented that “These cases reflect a sad chapter in the history of American business – a chapter in which those who reaped enormous benefits from the trust of investors profoundly betrayed that trust.” We couldn’t agree more. We think active investment management is already hard enough without having the deck stacked against you.

Three More Studies Highlight Indexing’s Advantage

This month we’d like to highlight three new studies that demonstrate the advantages of index investing over active management.

The Bogle Study

The first was presented by John Bogle (yes, that John Bogle) at the March 12, 2003 U.S. House of Representatives hearing on mutual fund fees. It had previously appeared in the *Journal of Portfolio Management*. It is based on data on U.S. mutual funds (unit trusts) for the ten years ended in June, 2001.

The fund results shown are adjusted for operating expenses, but not sales loads, taxes, or survivorship bias. As we will see, the latter is often overlooked in these studies, but is of considerable importance. It refers to the fact that many historical data series covering mutual fund performance do not include the results for funds that have gone out of business, or merged into funds that are more successful. As this happens far more frequently to actively

managed fund than to index funds, its effect is to impart an upward bias to the reported results for actively managed funds. Various studies have estimated the size of this survivorship bias to be between 1.5% and 3.1% per year, with the higher amounts applying to small cap funds.

Bogle’s analysis compared the impact of operating expense ratios on average fund returns over the ten year period. He compared the average returns earned by funds in the lowest quartile of operating expenses (whose average operating expense ratio was 0.6%) with those in the highest quartile (average ratio 1.8%). The results are shown in the following table:

Fund Type	Low-Cost Quartile	High-Cost Quartile	Low-Cost Advantage
Large Cap Value	14.8%	12.8%	2.0%
Large Cap Growth	14.2%	11.2%	3.0%
Mid-Cap Value	15.3%	12.5%	2.8%
Mid-Cap Growth	14.7%	12.5%	2.2%
Small Cap Value	16.8%	12.0%	4.8%
Small Cap Growth	15.4%	14.5%	0.9%

Virtually all index funds fall into this low cost group, and, as the table shows, the advantages produced by low costs are significant. Bogle’s analysis also looked at returns per unit of risk (also known as the “Sharpe Ratio”) for index versus actively managed funds. The results of this analysis are shown in the following table:

Fund Type	Index Funds	Actively Managed Funds	Index Fund Advantage
Large Cap Value	0.88	0.81	.07
Large Cap Growth	0.68	0.55	.13
Mid-Cap Value	1.00	0.82	.18
Mid-Cap Growth	0.48	0.45	.03
Small Cap Value	1.06	0.84	.22
Small Cap Growth	0.38	0.48	(.10)
All Funds	0.79	0.67	.12

As you can see, index funds appear to have a strong advantage in all categories except small cap growth. However, as Bogle points out, it must be remembered that the data for active fund returns have not been adjusted for survivorship bias. Over the ten years covered by Bogle's data, the average annual return for actively managed small cap growth funds was 14.38%, versus 12.62% for small cap growth index funds. The difference between these two amounts, 1.76%, is well within the 1.5% to 3.1% range estimated for the impact of survivorship bias (and remember, that the impact of the bias on small cap fund results was at the higher end of this range). In sum, it looks like more strong evidence in favor of index funds.

The WM Study

A more recent study, based on an entirely different data set of data from the UK, was conducted by the WM Company and sponsored by Virgin Direct Financial Services. It analyzed the performance of UK unit trusts (mutual funds) over the twenty years to the end of 2000, and compared their performance to the FTSE All-Share Index (a broad market index comparable to the Wilshire 5000 in the US).

The WM study found that, "over the twenty year period, of the 55 actively managed trusts with 20-year performance records, only eleven (20%) outperformed the index. Over the twelve years to the end of 2000 for which comparable data was available, the average tracker [index] fund return was greater than 70% of the active trusts."

The study also compared the performance of active versus index (tracker) funds over rolling five year periods. In this case, "it is clear that in general, over all [rolling] five year periods, the top quartile of the active trusts outperform the passive trusts. For an investor taking a five year time horizon, the chances of doing better using active rather than passive management was roughly on the order of one in four based on return evidence since 1990."

Much to our delight, the WM study also addressed the issue of performance persistence, and the ability to predict it in advance. It noted that "[this] outperformance of some active trusts

[over rolling five year periods] is to be expected but is not in itself an argument for active management. Evidence that such outperformance was consistent and predictable would be. Unfortunately, our previous reports have highlighted that outperformance tends not to repeat. Two quotes [from these previous reports] encapsulate our findings:

[First], “consistence of performance for active trusts (from one five year period to the next) would appear to be constrained by manager performance cycles. Few managers have been able to sustain outperformance (or underperformance) over the longer term. If this view is correct, then buying into a trust after several years of relative outperformance or selling after a period of relative underperformance may be highly detrimental to investor wealth.”

[Second], “this analysis would therefore suggest that ‘consistent’ performance [by actively managed funds] is on average apparent over a limited period only, certainly shorter than five years. Very few trusts exhibit consistent outperformance as they fail to position portfolios correctly to benefit from different phases of the economic cycle or they follow an investment style that by definition moves in or out of fashion in performance terms...Other reasons for lack of consistency could include the difficulties investment groups have in retaining investment staff who have exhibited short-term outperformance, and the difficulty in maintaining a consistent investment process as ‘hot’ money is directed at an outperforming trust.”

Regular readers will recognize that these conclusions echo points we have repeatedly made, and shown the evidence for, over the years in *The Index Investor*. These two additional studies only add to what has become a very, very strong case in favor of index funds.

Our Study

Clearly, much evidence based on comparable historical returns suggests that, due to their lower operating expenses and trading related costs (which include both trade execution at the fund level, and additional taxes at the investor level), index funds generally outperform actively managed mutual funds. In response to this, active management advocates sometimes

claim that index funds' superior performance is limited to certain periods of time, and that during other periods, the active approach performs better. This argument suffers from three shortcomings. The first is a lack of clarity about what constitutes a better performing active management approach to investing. While all active managers purchase assets in the belief that their price will rise in the future, value investors believe this will happen because the current market price is below the asset's theoretical (or "fundamental") worth. In contrast, momentum investors believe the price will rise because many other investors also will be buying the asset in question. In short, saying that active management sometimes outperforms indexing isn't enough: people making this argument should also be specific about the active management approach they are talking about.

The second shortcoming in the argument is that it fails to specify whether the occasional periods of superior performance by active managers are due to luck or skill (although the latter is clearly the impression the advocates are trying to give). Assuming it is the latter, the third shortcoming of the argument by active management advocates is that they usually fail to provide clear indicators one could use to identify (in advance) a period during which the value or momentum active management approach will outperform indexing.

Despite these shortcomings, many people have apparently found the active managers' argument to be quite persuasive (like the Siren's song, perhaps). The challenge we have been struggling with is to develop an effective way to evaluate it. We have worked on this for quite a while now, and have finally come up with some interesting results we'd like to share with the readers of *The Index Investor*. Not that you have to be convinced about the virtues of indexing, but rather to provide you with more ammunition to use the next time you encounter an outspoken active management advocate.

The problem with arguments based on different historical periods is that they rapidly take on the tone of a theological discussion: for every period you cite during which indexing prevailed over active management, I'll cite one when it did not. How do you resolve this? By comparing indexing to active management under a wide range of actual and theoretical historical scenarios. Fortunately, improved computer processing power and simulation

software makes it much easier to do this now that it was even five years ago. We'll begin with a description of the approach we took, and then move on to what we found in our analysis.

As noted, our goal was to simulate the performance of indexing versus active management under a wide variety of conditions. More specifically, we wanted to use an approach that was directly relevant to the situation of an investor trying to decide between a range of different types of mutual fund (e.g., index, value, etc.). Our model therefore needed to simulate conditions at three different levels. First, it needed to simulate the performance of the individual shares traded in the market. Second, it needed to simulate the performance of the different funds that bought and sold these shares based on the investment management approach they were using. Third, it needed to simulate the actual returns experienced by investors who had purchased interests in these different investment funds.

The Simulation Model

At the level of the individual shares traded in the market, we needed to provide information about their average returns, the standard deviation of these returns, and the correlations between the returns on different shares. One approach would have been to simply make up these numbers; however, this would have left us exposed to the criticism that we somehow had skewed this data to produce a better performance for indexing. On the other hand, we could have simply picked some actual shares at random from the thousands that are traded, and used actual historical data on them in our model. However, in this case we would be exposed to the criticism that the shares selected were somehow not representative of the equity market as a whole. In light of these criticisms, we took a third approach, and used as the shares in our simulated market the ten Dow Jones sector iShares that collectively make up the Dow Jones Total U.S. Equity Market Index. We were attracted to the fact that using only ten shares in our simulated market would make our computations easier while still accurately covering the entire U.S. equity market. The average returns, standard deviations, and correlations of returns for these ten sector iShares were based on monthly data covering 1992-2002. To make the model computations easier, we assumed that the shares paid no dividends, and that they all started out at the same price, and that subsequent price changes were based

on the shares' return (e.g., a share with a nine percent return in year two would also see its price increase by nine percent).

The next level of the simulation model – the investment funds – presented more challenges. To be fair, we modeled three funds that used different investment management approaches. The first is an index fund, which simply buys a portfolio of sector shares (based on their weight in the total market capitalization), and holds them. Over time, the weight of each sector in the index fund varies in line with the returns on the sector shares.

The second fund we modeled uses a value approach – that is, a belief that underperforming shares will eventually reverse course, and deliver attractive returns (i.e., “buy low, sell high”). Our challenge in this case was how to model the fund manager's decision process. For example, should the fund manager define “value” as a share with a low price relative to the others in the market, or as a share whose past year return was low relative to the other shares? And how many shares should the manager hold? The three lowest? The five (the latter would be the bottom half of all ten shares traded in our simulated market). In the end, we tested all of these approaches, and used the one that produced the highest returns for our value manager.

The third fund we modeled was a momentum fund (sometimes euphemistically called a “growth” fund), whose manager purchases shares that have performed well in the past in the belief that they will continue to do so in the future (i.e., “buy high, sell higher”). In this case, we confronted the same set of issues as we did with the value manager. How to define attractive shares for this fund? How many shares to hold? Again, we tested a variety of approaches, and chose the one that produced the best returns for our growth manager.

At the fund level of the model, we also had to specify the level of operating expenses (that is, the reported expense ratio) for each fund, as well as its trading costs. For the former, we used an estimate derived from large U.S. mutual funds, which have been in existence for all of the five years ending in 2002. This resulted in an expense ratio of 1.04% for our value fund, and 1.18% for our growth fund. This is a conservative estimate, because it does not include small cap funds. For example, a 2002 study by Lippper, Inc. estimated that the average operating

expenses charged by actively managed U.S. equity funds equaled 1.61% of their assets under management (we don't want to be accused of unfairly prejudicing our analysis against actively managed funds!). The expense ratio for our index fund is .25% (which is on the high side for an index fund focused on U.S. equities).

Trading costs presented an even thornier problem, because they aren't disclosed by U.S. mutual funds. They include not only explicit brokerage commissions, but also the impact of trades on the market price for a share, as well as changes in that share price that occur before the trade has completed. In our model, our estimates were based on data from the Plexus Group, which specializes in helping investment funds reduce their trading costs. Plexus estimates that, on average, the all in cost of a trade amounts to .53% (that is 53 one hundredths of one percent, or 53 "basis points") of the value of the trade for value funds, and 1.23% of the value of the trade for growth/momentum funds (for index funds we split the difference and used .88%, which again is probably on the high side). The different costs for value and growth funds reflect the different conditions under which each fund makes its trades. Momentum funds are generally trying to buy relatively large blocks of shares for which there is much demand (by other momentum funds) or sell them when there is very little demand (e.g., when the momentum funds are dumping them). In contrast, value funds tend to buy shares when demand for them is low, and sell them when demand for them is rising. Consequently, value funds have lower trading costs than momentum funds.

The final level of our simulation models includes four different investors. The first is a committed index investor, who buys and holds the index fund. The second is a committed momentum investor, who buys and holds the momentum fund. The third is a committed value investor, who buys and holds that fund. Our final investor is committed to active management in general, rather than to any specific approach. As a result, this investor chases performance, switching each year between the momentum and value funds, depending on which one had the best performance the previous year.

A very important part of our investor model is our treatment of taxes. To simplify our calculations, we assume that our investment funds trade only once, at the end of each year.

The capital gains they realize on these trades are distributed to investors (the capital losses are carried forward at the fund level, and reduce future capital gains). When they are received by investors, these capital gains are taxed at only 20% (that is, at a rate below the ordinary income tax rate). Finally, we ran the model more than once to simulate the impact of having our active investors pay front end loads (of 5.75%) when they purchased their respective funds, and the impact of using no-load funds (that is, funds that did not charge a front end sales load).

When our model is run, it simulates five years of activity. A single model run consists of 5,000 simulations. Each simulation begins with the random selection of an annual return for each of the ten shares from their respective distributions of possible returns. This is done independently for each of the five years. In other words, for a given sector share (e.g., Energy or Healthcare), a different return is selected for each year in the model from five different annual return distributions. Once these returns have been determined, they are used to calculate the end of year value for the shares owned by each fund (note that the beginning values for the momentum and value funds are reduced by sales loads when those are used). Each fund then applies its trading rule to determine which shares it will buy or sell before the start of the next year. These portfolio adjustments generate the fund's trading volume (and trading expense), as well as the capital gains that are passed on to its investors. Finally, at the end of each year, the fund's "reported return" is calculated. This is based on the percentage change in the value of its portfolio over the course of the year (net of operating and trading costs), but before the reduction in the value of the portfolio due to the payout of realized capital gains and losses.

We should clearly note that in our model, the index fund has no turnover each year, and does not realize any capital gains or losses. This is not what happens at a real index fund, where a low level (in comparison with actively managed funds) of turnover (21% on average) is caused each year by investors' liquidity needs (that is, to facilitate investors buying and selling shares in the fund) and by changes in the companies that make up the index the fund tracks. In our model, we have assumed no changes in the index (which would only affect the index fund, but happens irregularly, depending on the index), and no investor liquidity needs

(which would affect all funds). In other words, we don't think the lack of capital gains in our index fund is significant, because we have not included turnover caused by investor liquidity needs in the actively managed funds either.

At the investor level, we calculate the effective annual return net of any sales loads charged and any capital gains taxes paid. We also calculate an ending portfolio value after five years (based on an initial investment of \$100,000) and a compound annual rate of return (that is, a geometric return) over this period.

The Results

After much experimentation, we found that fund trading rules based on share prices (e.g., buy the three stocks that have the highest prices) produced higher returns (for investors) than trading rules based on share returns (e.g., buy the three stocks that had the highest returns last year). Price based trading rules enabled funds to hold onto a greater percentage of previous winners, whereas return based rules sometimes forced winners to be sold too soon.

We also found that, for the value fund, a trading rule that held five shares (that is, that allocated 20 percent of the fund's the portfolio to the five shares with the lowest end of the year price) produced higher returns (for investors) than trading rules that held fewer shares, because it limited capital gains realizations. In contrast, we found that holding only three shares produced the highest (investor) returns at the momentum fund, as the potential for excessive trading was offset (via the price rule) by holding on to winners for longer periods of time.

Let's now move on to the specific results from our model, given the trading rules we used. We should also note that the results reported below assume our active investors were using no-load funds. If load funds were used by active investors, the advantages of index investing only increase.

We'll start with some benchmark comparisons to real world data. In our model (where the funds trade only once per year), our value fund had an average annual turnover ratio (value of shares bought and sold as a percentage of the fund's total value) of 54%, while our momentum fund had an average turnover of 56%. In comparison, the previously mentioned sample of large cap U.S. equity funds over the 1998 –2002 period saw average annual turnover of 74% at value funds, and 98% at growth/momentum funds. This means that, if anything, our analysis underestimates the trading costs incurred by our respective actively managed funds, and overstates the returns they produce.

Now let's look at the compound annual after-tax returns that were realized by our four investors over five years. The table below shows the average return (over 5,000 model simulations), the median return (i.e., half the returns are above the median, and half the returns are below it), the standard deviation of the average return, and the ratio of the average return to the standard deviation of returns (i.e., the amount of return per unit of risk taken on).

Investor	Average After-Tax Return	Median Return	Std. Deviation	Avg/Std. Dev.
Index Investor	11.35%	11.38%	6.08%	1.87%
Momentum	9.73%	9.49%	7.13%	1.36%
Value	5.09%	5.19%	5.28%	.96
Perf. Chaser	6.90%	6.76%	6.57%	1.05%

As you can see in this table, index investing proved to be the superior strategy, on both an absolute return (Average Return) and a risk-adjusted return (Avg/Std. Dev) basis. Somewhat surprisingly, the second best approach turned out to be momentum, rather than value. Before doing this analysis, we would have guessed value. While this result may reflect the combination of the historical period from which our share price data was drawn (1992-2002 includes the internet boom, the golden age for momentum investing) and the momentum trading rule we used, it still surprised us, and led us to ask three further questions.

First, how many times did each strategy produce the highest return for an investor? The following table shows the percentage of times (out of 5,000 simulations) the returns on each strategy ranked first, second, third or fourth.

	First	Second	Third	Fourth	Total
Index	67%	26%	7%	0%	100%
Momentum	30%	48%	12%	10%	100%
Value	2%	16%	24%	58%	100%
Chaser	1%	10%	57%	32%	100%
Total	100%	100%	100%	100%	

Again, this analysis showed that the index approach was the one most likely to produce the top after tax return. It also showed that this was not guaranteed to be the case under all circumstances. Frankly, we quite like this table because it very clearly brings home a central point about investing: There are no guarantees -- the best we can do is make well informed decisions that have a high probability of producing the returns we seek.

The second subsequent analysis we performed was to run our model over eleven years (not five) using the actual returns for the Dow Jones Sector iShares over the 1992-2002 period.

In this case, the index investor realized a compound annual after-tax return (over the eleven year period) of 9.04%, the momentum investor 6.34%, the value investor 6.90%, and the performance chaser 4.81%. Once again, the indexing approach turned out to be superior, by a large margin. Another interesting question raised by this analysis is why momentum fared worse in this case than in the simulations. The simple answer is that the “real world” data represents only one possible set of outcomes, while our simulation results are based on 5,000 sets of possible outcomes.

A more complicated answer (which we have no analytical basis for verifying) may be that value strategies (as defined in our model) work better over longer periods of time. A variant of this answer is our original guess that value would outperform momentum was based on data that reflected a more simplistic view of these two strategies than the view we have taken

in our simulation model. For example, until recently, most of the companies that compile indexes created “growth” and “value” subsets of their data simply by dividing the underlying companies into two equal groups, based on a single criterion such as the ratio of their market to book value ratio. In contrast, in our model the active funds do not invest in all the shares in the market; rather, each year the momentum fund invests in the three shares (out of ten) with the highest price at the end of the preceding year, while the value fund invests in the five shares with the lowest price. In other words, our original guess that value would outperform momentum was inferred from data that was based an inaccurate view of how these two strategies are actually implemented in practice.

The third question we asked was whether the results we obtained were driven by tax considerations. To answer this question, we re-ran our analysis, and looked at the gross returns reported by the funds, rather than the after-tax returns to investors. This is equivalent to a situation in which our investors all held their funds in tax-deferred retirement accounts.

Here are the results of this analysis:

Investor	Average Pre-Tax Return	Median Return	Std. Deviation	Avg/Std. Dev.
Index Investor	11.35%	11.35%	6.13%	1.85%
Momentum	10.22%	10.15%	7.30%	1.40%
Value	5.64%	5.75%	5.36%	1.05%
Perf. Chaser	7.56%	7.45%	6.64%	1.14%

Once again, our index investor comes out ahead on average. However, how often can we expect this to be the case? The following table shows the percentage of times (out of 5,000 simulations) the returns on each strategy ranked first, second, third or fourth, on a pre-tax basis:

	First	Second	Third	Fourth	Total
Index	58%	32%	10%	0%	100%
Momentum	37%	41%	11%	11%	100%
Value	3%	15%	22%	60%	100%
Chaser	2%	12%	57%	29%	100%
Total	100%	100%	100%	100%	

In short, indexing's performance advantage over its actively managed competitors (on both an absolute and a risk-adjusted basis) was not wholly driven by its tax advantage, and remained even after this was removed.

One Final Test

In light of the results of our initial analysis, we decided to re-run our simulation model using an entirely different set of input data. This time we used data on returns, standard deviations, and correlations between 1978 and 1988 for the following ten asset classes: U.S. equities (using the Wilshire 5000 Index); European equities (MSCI Europe); Pacific Equities (MSCI Pacific); U.S. High Yield Debt; U.S. Investment Grade Debt, U.S. Real Estate Investment Trusts; Commodities (Goldman Sachs Commodities Index); and long term government bond returns (converted to U.S. dollars) for Australia, Germany, and the United Kingdom. The index fund in our model began the simulation with an equal 10 percent allocation to each of these asset classes. We used the same assumptions for operating expenses and trading costs as in our previous simulation, and assumed the use of no-load funds.

Here are the results of this analysis, based on pre-tax returns to our four investors after five years, based on 5,000 simulations:

Investor	Average Pre-Tax Return	Median Return	Std. Deviation	Avg/Std. Dev.
Index Investor	14.36%	14.27%	4.68%	3.07%

Investor	Average Pre-Tax Return	Median Return	Std. Deviation	Avg/Std. Dev.
Momentum	13.49%	13.40%	7.86%	1.72%
Value	10.98%	11.01%	4.31%	2.55%
Perf. Chaser	13.25%	12.97%	5.82%	2.28%

Once again, our index investor comes out ahead on average, especially on a risk adjusted basis. But how often can we expect our index investor to earn the highest pre-tax return? The following table shows the percentage of times (out of 5,000 simulations) the returns on each strategy ranked first, second, third or fourth, on a pre-tax basis:

	First	Second	Third	Fourth	Total
Index	35%	36%	29%	0%	100%
Momentum	39%	17%	14%	30%	100%
Value	12%	18%	19%	51%	100%
Chaser	14%	29%	39%	19%	100%
Total	100%	100%	100%	100%	

Again, indexing still held its own against competing strategies, even without its tax advantage. We found particularly interesting the distribution of outcomes for the momentum strategy: it was very much feast or famine in this case, with first and fourth place finishes much more common than seconds or thirds. Finally, while the value approach performed relatively better in this simulation than it did in the previous one, it was still well below what we would have guessed ahead of time.

Conclusions

The results of this analysis are both exciting and troubling. They are exciting because they addresses some of the key criticisms of previous analyses raised by active management advocates, while still reaching the same conclusion: over the long run, index investing delivers superior risk-adjusted results. On the other hand, the results are also troubling. First, we were frankly shocked by the extent to which the value strategy underperformed both momentum and indexing. Given that most of the underlying dynamics in our model were

random (apart from our choice of a simple trading rule for each actively managed fund), it is very clear that any argument in favor of value investing must rely very heavily on fund-specific advantages (e.g., superior information, or a superior model for making sense of it) rather than systemic ones inherent to the strategy in general. This in turn implies that managers at successful value funds are clearly adding unique value in exchange for the fees they charge their investors. Our analysis, however, suggests that it may be very difficult for these managers to achieve superior long-term results. Moreover, our analysis does nothing to alleviate our doubts about the ability of an average investor to identify superior value managers in advance.

A different, but related set of arguments applies to momentum (or more euphemistically, “growth”) funds. We were equally troubled by the strong performance of the momentum strategy in our model (on an absolute, if not always on a risk-adjusted basis), regardless of the underlying set of asset classes and historical returns data we used. Our results suggest that there is a systemic basis for the relative success of the momentum approach that has nothing to do with the skills of these funds’ managers. This leaves us wondering about just what it is that justifies the high fees charged by these funds.

Finally, the dismal returns earned by our performance chasing investor also need to be highlighted. Multiple research studies have shown that new investments into mutual funds are strongly related to their recent performance. Given this, it is no surprise that much of the advertising by actively managed funds focuses on past performance. Despite this, and in line with previous studies, our analysis finds that an investment strategy based on performance chasing is a recipe for inferior performance, and that patient index investing is a wholly superior approach. In other words, while the momentum strategy seems to work at the fund level (where it is applied to shares), it fails to do so when applied by individuals to different mutual funds. This is a very important distinction that is rarely if ever made in fund marketing materials.

In sum, after completing this analysis, our belief in the advantages of indexing is stronger than ever.

