

The Index Investor

Why Pay More for Less?

Global Asset Class Returns

YTD 27Feb04	In USD	In AUD	In CAD	In EURO	In JPY	In GBP
US Bonds	1.90%	-0.85%	4.86%	2.57%	3.37%	-2.65%
US Prop.	6.00%	3.25%	8.96%	6.67%	7.47%	1.45%
US Equity	3.60%	0.85%	6.56%	4.27%	5.07%	-0.95%
AUS Bonds	3.32%	0.57%	6.27%	3.99%	4.79%	-1.23%
AUS Prop.	7.14%	4.39%	10.10%	7.81%	8.61%	2.59%
AUS Equity	4.40%	1.65%	7.36%	5.07%	5.87%	-0.15%
CAN Bonds	-1.32%	-4.07%	1.64%	-0.65%	0.15%	-5.87%
CAN Prop.	3.26%	0.51%	6.21%	3.93%	4.73%	-1.29%
CAN Equity	3.10%	0.35%	6.06%	3.77%	4.57%	-1.45%
Euro Bonds	0.44%	-2.31%	3.40%	1.11%	1.91%	-4.11%
Euro Prop.	13.28%	10.53%	16.23%	13.95%	14.75%	8.73%
Euro Equity	3.80%	1.05%	6.76%	4.47%	5.27%	-0.75%
Japan Bonds	-1.13%	-3.88%	1.83%	-0.46%	0.34%	-5.68%
Japan Prop.	18.47%	15.72%	21.43%	19.14%	19.94%	13.92%
Japan Equity	0.70%	-2.05%	3.66%	1.37%	2.17%	-3.85%
UK Bonds	4.75%	2.00%	7.71%	5.42%	6.22%	0.20%
UK Prop.	19.15%	16.40%	22.11%	19.82%	20.62%	14.61%
UK Equity	4.20%	1.45%	7.16%	4.87%	5.67%	-0.35%
World Bonds	0.95%	-1.80%	3.91%	1.62%	2.42%	-3.60%
World Prop.	10.10%	7.35%	13.06%	10.77%	11.57%	5.55%
World Equity	3.85%	1.10%	6.81%	4.52%	5.32%	-0.70%
Commodities	11.70%	8.95%	14.66%	12.37%	13.17%	7.15%
Hedge Funds	2.01%	-0.74%	4.97%	2.68%	3.48%	-2.54%
A\$	-2.75%	0.00%	-5.71%	-3.42%	-4.22%	1.80%
C\$	2.96%	5.71%	0.00%	2.29%	1.48%	7.50%
Euro	0.67%	3.42%	-2.29%	0.00%	-0.80%	5.21%
Yen	1.47%	4.22%	-1.48%	0.80%	0.00%	6.02%
UK£	-4.55%	-1.80%	-7.50%	-5.21%	-6.02%	0.00%
US\$	0.00%	2.75%	-2.96%	-0.67%	-1.47%	4.55%

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 February, this benchmark had returned 3.3%, while our model portfolio had returned 4.8%. 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 3.3%.

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 2.9%, while our model portfolio had returned 4.2%, and the global benchmark had returned 2.7%.

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 2.2%, while our model portfolio had returned 3.1% and the global benchmark 1.5%.

The objective of our second set of model portfolios is to deliver less risk than their respective benchmarks, while delivering at least as much return. The benchmark for the first portfolio in this group is an aggressive mix of 80% domestic equities, and 20% domestic bonds. Through the end of last month, this benchmark had returned 3.3%, while our model portfolio had returned 4.6%. 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 3.3%.

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 2.9%, while our model portfolio had returned 3.5%, and the global benchmark had returned 2.7%.

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 2.2%, while our model portfolio had returned 3.2% and the global benchmark 1.5%.

The objective of our third set of model portfolios is not to outperform a benchmark index, but rather to deliver a minimum level of compound annual real return over a twenty-year period. Through last month, our 7% target real return portfolio had returned, in nominal terms, 4.4% year-to-date, our 5% target real return portfolio had returned, in nominal terms, 3.7%, and our 3% target real return portfolio had returned, in nominal terms, 3.6%.

Our fourth set of model portfolios are also target real return portfolios; however, they include the possibility of investing in a hedge fund index. Through last month, our 7% target real return HF portfolio had returned, in nominal terms, 4.0% year-to-date, our 5% target real return HF portfolio had returned, in nominal terms, 3.5%, and our 3% target real return HF portfolio had returned, in nominal terms, 3.7%.

Equity Market Valuation Update

Our equity market valuation analysis rests on two fundamental assumptions. The first is that the long term real equity risk premium is 4.0% per year. The second is the average rate of productivity growth an economy will achieve in the future. As described in our June, 2003 issue, we use both high and a low productivity growth assumptions. Given these assumptions, here is our updated market valuation analysis at the end of last month:

Country	Real Risk Free Rate Plus	Equity Risk Premium Equals	Required Real Return on Equities	Expected Real Growth Rate* plus	Dividend Yield Equals	Expected Real Equity Return**
Australia	3.21%	4.00%	7.21%	4.90%	3.73%	8.63%
Canada	2.48%	4.00%	6.48%	2.10%	1.78%	3.88%
Eurozone	1.47%	4.00%	5.47%	2.50%	2.43%	4.93%
Japan	1.60%	4.00%	5.60%	2.70%	0.89%	3.69%
U.K.	1.85%	4.00%	5.85%	2.50%	3.22%	5.72%
U.S.A.	2.05%	4.00%	6.05%	4.50%	1.59%	6.09%

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

** When required real equity return is greater than expected real equity return, theoretical index value will be less than actual index value – i.e., the market will appear to be overvalued.

Country	Implied Index Value*	Current Index Value	(Under) or Overvaluation in High Growth Scenario	(Under) or Overvaluation in LowGrowth Scenario
Australia	161.47	100.00	-61%	-13%
Canada	40.64	100.00	59%	67%
Eurozone	81.82	100.00	18%	46%
Japan	31.79	100.00	68%	77%
U.K.	96.12	100.00	4%	34%
U.S.A.	102.58	100.00	-3%	38%

* High productivity growth scenario.

This Month's Letter to the Editor

Does it make sense to change a portfolio's asset allocation all at once, or to do it over time?

Unfortunately, there is no clear-cut answer to your question about the best way to move from one's current portfolio to a new asset allocation. The main reason for this is the wide range of tax situations faced by different investors.

For example, quite a few of our readers have substantial amounts of unrealized gains and losses that could trigger significant tax consequences were they to substantially change either their asset allocation and/or holdings within different asset classes. Under these circumstances, deciding whether and how much change to undertake is a uniquely personal decision where a tax advisor has much more to offer than we do. At best, we provide general guidelines on asset location and tax efficiency (see the green button with this title on our home page), as well as some information on the trade-offs between index mutual funds and exchange traded index funds (ETFs). For example, the former are usually much better for people who are gradually changing their position, because the commission costs charged on ETF purchases are avoided. On the other hand, index mutual funds offer less control over the realization of future taxable capital gains distributions than do ETFs (though there are tax-managed index funds that also do an impressive job of holding down the realization of capital gains). Again, however, we note that because of the complexity of many people's tax situations, these general views must be integrated with (and occasionally traded off against) the specifics of an individual's situation.

On the other hand, when potential tax consequences aren't an important issue, our basic belief that it is usually impossible to accurately forecast future market returns leads us to prefer changing portfolio allocations slowly over time. In most cases, dollar cost averaging (i.e., sales and/or purchases of the same dollar amount each period) works to one's advantage, because over time it balances out the normal ups and downs in any market. That being said, experience has also taught us the folly of being completely inflexible ideologues on the subject of dollar cost averaging. We recognize that sometimes there are situations where an

asset class is so clearly overvalued that it behooves one to wait before adding to one's position in it (undervalued situations seem much more difficult to spot -- or at least feel confident about having spotted). The two biggest examples of this we cite are the overvaluation of the British pound in 1992 (when it was about to pull out of the EMU), and U.S. equities early in 2000 (but even in the case of the latter, there were plenty of voices suggesting they were still undervalued – see our March, 2000 issue). In our experience, these extremely overvalued situations seem to be relatively rare, but they do occasionally occur.

To help people identify extremely overvalued situations, each month's issue of *The Index Investor* contains our "Equity Markets Valuation Update", which is based on our application of the dividend discount model to different markets.

When it comes to whether or not an equity market is fairly valued, our methodology compares the future returns an equity market is expected to supply with those demanded by investors.

The returns the equity market is expected to supply are equal to the current dividend yield, plus the expected rate of dividend growth (we leave P/E changes out of our model, since they are essentially behaviorally driven).

The rate of return that a rational investor should demand in order to hold equity instead of risk free government bonds is equal to the current real rate of return on government bonds plus an appropriate equity market risk premium.

If the equity return the market is expected to supply is greater than the return a rational investor should demand, then our model will show that market to be undervalued. Under these circumstances, one would expect equity prices to increase, which would reduce the market dividend yield (dividends/market value) and bring the supply of returns down to the level of the returns that investors demand.

On the other hand, if the rate of return the market is expected to supply is less than the rate of return a rational investor would demand, then our model will show the market to be

overvalued. Under these circumstances, we would expect equity prices to fall, which would increase the dividend yield and raise the supply of returns up to the level that investors demand.

With respect to the valuation of fixed income markets, over and undervaluation substantially depend on the future changes in interest and exchange rates that one expects. We cover potential scenarios for how these may develop (and their implications for current fixed income market valuation levels) in our semi-annual economic outlooks.

This Month's Feature Articles: Key Points

This month's issue contains our annual review of recent academic research. Our first feature article focuses on two key questions: is classical efficient markets theory dead? And if it is, does index investing still make sense? We find that, indeed, the overwhelming weight of evidence leads one to conclude that the three pillars of classical efficient markets theory do not describe the way real world financial markets behave. First, rather than being purely rational, and with unlimited cognitive processing capacity, investors' decision making process is actually based on a combination of rational, intuitive and emotional factors. Second, rather than having access to perfect information, investors differ in their private information about fundamental asset values, and in their beliefs about how other investors will behave. Third, in the real world there are significant limitations on the ability of arbitrageurs to quickly correct market mispricings. However, while classical efficient markets theory is dead, the markets remain efficient, in the sense that it is still extremely difficult for anyone to consistently succeed (through skill) at active management (and next to impossible to identify these managers in advance). Why? Because a more realistic view of financial markets shows them to be a complex adaptive system whose behavior is beyond the power of man or machine to consistently predict. As a result, the case for indexing is stronger than ever.

Our second feature article continues our research review, and looks at findings of particular relevance to different strategies and products. Among other interesting findings, we show that most investors are still woefully under-diversified, that the term structure of real interest

rates is essentially flat, that U.S. real return bonds aren't as tax-inefficient as some have asserted, that on average private equity funds have a tough time beating the S&P 500 after their fees are taken into account, and that a distressing percentage of 401K plan trustees do a much less than impressive job of looking out for the best interests of their plan participants.

Annual Research Review Part One: Has The Death of Efficient Markets Theory Killed Indexing Too?

This month's feature articles are a two part summary of recent academic research findings in the area of finance and investments. In this one we will consider important recent developments in asset pricing theory -- that is, the factors and processes which determine the returns on different investments. In particular, we will focus on two key questions: is classical efficient markets theory dead? And if it is, does index investing still make sense?

In our second feature article this month, we will look at research findings that are related to a number of product and strategy issues. However, before we get into these issues, let's first step back and review why we take the time each year to review how these research findings are produced, and why it is important for us to understand them.

Let's start with the basics. Where does a theory come from? It usually starts with a rule, which is used to predict that a given outcome will occur in a given case. For example, nominal return bonds will have low returns when inflation is rising. And where do these rules come from? Sometimes they are simply asserted to be true: these are called axioms, or first principles. Other times they are arrived at inductively: that is, data are observed, a tentative cause is identified, the process is repeated and if the cause continues to be present it is asserted as a general rule. And what is the source of this data? Among the most important are time series (which study the change in a phenomena, like the returns on a value index, over time), cross-sectional data (which studies the differences in outcomes which all occur at the same time, such as the returns on different stocks in 2001), panel data (which looks at the evolution of cross-sectional data over time), the results of laboratory experiments involving real people (e.g., a simulated stock market with students playing the role of investors), and the

results of agent based simulations, in which software programs take on the roles of different players.

Finally, how do we choose between competing theories? Our natural tendency is to prefer the theory which has the most evidence supporting it. However, because the same evidence is often consistent with more than one theory, the advancement of knowledge via the scientific method is based on disproving theories rather than proving them. For example, a set of panel data that shows that companies with high book to market ratios (often known as "value" companies) delivered higher risk adjusted returns over a given period of time than companies with low book-to-market ratios ("growth" companies) can be used to support both a theory asserting that this represents compensation for a risk factor, or one that asserts it reflects the cognitive failings and irrational behavior of a substantial number of investors.

Unfortunately, unlike chemistry or physics experiments, it is usually impossible to perfectly replicate the conditions that led to the assertion of a social science theory. Therefore, the fact that a given theory doesn't produce the expected result when it is applied to a different case isn't always considered grounds for rejecting it. Proponents of the theory can always assert that the case itself was sufficiently different that the theory didn't apply. While agent based models are exceptions to this (because they can be perfectly replicated), they suffer from the criticism that the rules their agents use do not capture the full range of human behavior, and the complex mix of perception, memory, cognition and emotion upon which it is based. Given these limitations, the process of discrediting one social science theory and replacing it with another one is necessarily long, uncertain and more than a little contentious at times. In the meantime, we are left with the task of weighing the evidence in support of different theories. In practice, this requires us to answer two questions: (1) How relevant is the evidence? (i.e., assuming it is true, to what extent does it allow us to revise our belief in the probability that one or more theories is true?) (2) How credible is it? (i.e., what is the probability that it is true?)

While these concepts might seem quite abstract, when it comes to investing we need to understand them. More specifically, people who advocate the use of index products are

increasingly likely to be confronted (sometimes loudly, and in public) by someone who says, "Efficient markets theory is dead. Doesn't this invalidate the case for indexing?" In this article, we'll start by examining the first statement, and end with our answer to the question.

First, just what is "efficient markets theory?" In practice, it is a group of sub-theories that, taken together, enable investors to form a mental model of how financial asset returns are determined and use it to predict the results of different investment decisions. Essentially, efficient markets theory rests on three legs. The first is the assumption that investors are rational, and have unlimited mental processing capacity. The second is the assumption that these investors all have instant and costless access to perfect information about the assets in the market. And the third is the assumption that any departures from the first two assumptions will be instantly identified and arbitrated away. Given these assumptions, efficient markets theory -- in its strongest form -- concludes that trading volume should be relatively low, and driven by either the arrival of new information or investors' needs for liquidity (e.g., the need to sell shares to raise the cash to buy a car). A further conclusion of efficient markets theory is that active management is doomed to failure, since future asset prices (and returns) cannot be predicted in advance. As such, the best advice for investors is to hold a portfolio of assets that is large enough to ensure that company specific risks cancel each other out, and all that is left is undiversifiable (also known as market, systematic, or beta) risk. As a result, the return on any individual asset should be driven solely by the extent to which it varies with the return on the market. This is determined via a regression equation, in which an individual asset's relationship with the market is measured denoted by a Greek letter: beta. In short, efficient markets theory provides the original justification for index investing.

But wait, you say: does the evidence really support those three assumptions? Over the past few years, a growing body of evidence has been accumulated that suggests that each of these assumptions is probably wrong. In fact, multiple studies have shown that the return on a given asset is driven by many more factors than its relationship with market return. These findings have led some to say that beta, and the efficient markets theory, is dead. In past issues of *The Index Investor*, we have reviewed key research findings that have led to this

conclusion. To sum them up, it is reasonably safe to say that many researchers would agree that (a) investors can better be described as "boundedly rational", with limited cognitive capacity (e.g., attention, memory and the like) that sometimes leads to behavior that appears to be irrational (including overconfidence, excessive optimism, the repeated use of thinking short cuts that lead to systematic errors, and a tendency to give greater weight to either losses or gains, depending on the circumstances); (b) because of differences in access to private information, and/or the uneven speed with which public information diffuses, investors are heterogenous in the information they possess when making their decisions; and (c) the departures from rational asset prices that are caused by (a) and (b) can persist over long periods of time because of the existence of obstacles that prevent arbitrageurs from eliminating them. Broadly speaking, this is the set of assumptions that underlie the field of study called "behavioral finance". Judging by the number of leading behavioral finance academics who have set up firms to actively managing client investments, one can also infer that they believe that assumptions (a), (b), and (c) undermine the case for indexing, and potentially make active management a profitable activity. However, as we pointed out in our recent article on active management, given the short track records of many active managers, it is impossible to tell whether they result from skill or luck. And, more importantly, there is still no proven way (apart from luck) to identify in advance an active manager who will be successful over the long term. Given this, the case for index investing remains very much intact, even if one accepts the behavioral finance school's arguments on their own terms.

However, proponents of efficient markets have not been willing to go that far. In defense of their mental model, they start by conceding that the return on individual assets depends not just on the return on the market, but also on the extent of the asset's exposure to other risk factors. Three of these are typically cited: a "value" factor (which results in high book to market stocks having higher returns than low book to market stocks); a "size" factor (returns on smaller capitalization stocks are higher than returns on large capitalization stocks), and a "momentum" factor (short term returns on stocks which have done well in the recent past tend to be higher than returns on stocks which have not performed as well). Moreover, while recognizing that some investors may be irrational and ill-informed, proponents of efficient markets theory state that due to the interaction of multiple investors the market as a whole still

behaves as one would expect under the original efficient markets theory. To be sure, some investors will be able to earn either more return with more risk than the market portfolio, or less return with less risk, due to their choosing different exposures to value, size and momentum. However, successful long term active management -- defined as earning more return than the market portfolio with less risk -- remains impossible (except by luck). Under this "weak form" of the efficient markets theory, the case for index investing remains strong. As we have noted in previous articles, this view has also attracted criticism, particularly around the ill-defined risks for which the value, growth, and momentum factors are proxies. In response, the efficient market school has tried to link the value and size factors to an increased risk of financial distress (which they claim is higher at smaller and high book to market firms), and the momentum factor to risk related to the timing of the business cycle. This was pretty much the state of play at the time of last year's financial research review. However, over the past year a great deal more research has been produced which suggests the need to further revise our assumptions about investors, information, and arbitrage.

Let's start with investors. Two interesting recent papers (note: the papers we cite can generally be found on the internet by using a search engine like google, yahoo, ssrn or msn) provide further evidence that investors appear to behave irrationally from time to time and that the impact of their actions is not always immediately offset by the actions of more rational investors. In their paper "Did Pension Plan Accounting Contribute to the Stock Market Bubble?" Coronado and Sharpe assessed the extent to which equity investors were, in effect, fooled by pension plan accounting. To do this, they compare the value of different companies' net pension assets as disclosed in the Form 10-K filed with the U.S. Securities and Exchange Commission with the value implied by these companies' market capitalizations. They find that "overvaluation of pension earnings reached five percent of total market value in 2001." In a paper near and dear to our hearts, in their paper "Are Investors Rational? Choices Among Index Funds" Elton, Gruber, and Busse found that despite an average difference in return of 2.09% between the best and worst performing S&P 500 index fund between 1996 and 2001, the worst performing funds still received substantial cash inflows. The authors explain this result as follows: "Any market consists of a set of informed, rational investors, and a set of uninformed investors. Markets are made efficient by the arbitrage

activities of the informed investors. But the only thing an informed investor can do in the market for index funds is to buy the good-performing funds -- no arbitrage is possible. In such a market, all that is needed for inferior funds to exist and grow is a set of uninformed investors and a set of distributors who have an economic incentive to sell inferior products. In a market where arbitrage is not possible, we may be disappointed, but we should not be surprised when inferior products exist and even prosper."

One of the key assumptions made by the behavioral finance school has been that prospect theory can be used to explain investor behavior. This theory states that investors are about twice as sensitive to losses as they are to gains; as a result, they will tend to realize their profitable investments too soon, and hold onto their losers for too long. This is known as the "disposition effect." However, most of the research that gave rise to this theory has been based on the actions of individual investors. In "Prospect Theory and Institutional Investors" O'Connell and Teo tested its application to the group of investors who collectively account for the majority of trading volume in most financial markets. They found "no evidence whatsoever of disposition effects; rather the dominant characteristic [of the investors they studied] was aggressive risk reduction in the wake of losses." They also found that this phenomena was related to time (or, more accurately, the nearness to year-end and the final performance numbers that would determine bonuses). "Fund managers were conditionally more risk-tolerant in the first half of the year. Gains during this period lead to incremental risk taking, but there was no evidence of this during the second half of the year. Correspondingly, losses in the first half of the year produced very little risk reduction: it was only in the second half of the year that managers systematically cut risk following losses." Finally, they note that experience (learning) also plays an important role: "older, wiser funds did not take on more risk in the wake of gains, but cut risk more aggressively in the wake of losses." The authors conclude that the modified version of prospect theory first proposed by Barberis, Huang, and Santos (in "Prospect Theory and Asset Prices") best explains the behavior they observed. In this theory, rather than being a constant, an investor's degree of risk aversion changes in response to the evolution of gains and losses relative to some starting anchor value (reference point). As gains grow larger, the investor becomes less risk averse (i.e., they reduce their equity risk premium), which lowers their required rate of return and

drives asset prices still higher. However, as losses grow, so too does risk aversion and the required rate of return, which further accelerates the decline in asset prices. In short, the model proposed by Barberis, et al, whose presence was tentatively confirmed by O'Connell and Teo, implies much more volatile asset prices (and returns) than would be the case where all investors were rational and only changed their valuation of an asset in response to new information about its future cash flows or a change in interest rates.

These findings suggest that there is something more than purely cognitive forces at work in investors' decision making process. A number of other papers have explored just this point. In his paper "Risk As Analysis and Risk As Feelings" Paul Slovic reports his findings that "there are two fundamental ways in which human beings comprehend risk. The 'analytic system' uses algorithms and normative rules. It is relatively slow, effortful, and requires conscious control." In contrast, "the experiential system is intuitive, fast, mostly automatic and not very accessible to conscious awareness...It relies on images and associations, linked by experience to emotion and affect (a feeling that something is good or bad). This system represents risk as a feeling...Proponents of formal risk analysis tend to view affective responses to risk as irrational. Current wisdom disputes this view. The rational and the experiential systems operate in parallel and each seems to depend on the other for guidance...Both systems have their advantages, biases, and limitations...Rational decision making requires integration of both modes of thought." In their paper "How Do Investors Judge the Risk of Financial Items?", Koonce, McAnally, and Mercer finds that that both rational factors (probabilities and outcome values) as well as two emotionally based factors called "dread" and "unknown" drive perceptions of financial risk. As the authors describe it, "dread captures a risky item's perceived controllability and voluntariness, as well as the amount of worry and catastrophic loss potential associated with the item, while unknown captures a risky item's observability, its newness, the amount of knowledge one has about the item, and the immediacy of the item's effects." In terms of the relationship between these factors, they find that "higher potential loss outcomes lead to greater dread, and greater dread in turn leads to greater perceived risk."

Feelings, however, have an impact not only on risk, but also on expected returns. For example, in their paper "Optimal Expectations", Brunnermeier and Parker develop a theoretical model in which an investor's current satisfaction is based not only on his or her purely rational calculation of the present value of a portfolio's expected future returns, but also on his or her subjective assessment of the future feeling that outcome will produce. For this reason, an investor is made happier today by holding overly optimistic expectations about future returns. However, this tendency toward excessive optimism has to be balanced against the expected cost of the incorrect decisions that it produces (e.g., the difference between the returns on the optimistic portfolio allocation and the higher returns a purely rational, and more diversified portfolio allocation would have produced). The authors' model shows how the trade-off between current satisfaction and future costs (e.g., being exploited by more rational investors) produces an "optimal" level of investor over-optimism about future asset returns. It also shows how it leads the investor to systematically underestimate the variability of future returns (that is, to be overconfident as well as over-optimistic). Finally, the authors speculate about these investors' long term survival. On the one hand, their over-optimism and about returns and overconfidence about risks should, at the margin, cause them to lose wealth to more rational investors. On the other hand, because they hold a riskier portfolio than the rational investors, under favorable economic conditions the over-optimistic investors also realize higher long-term returns. On balance, the authors conclude that there is no a-priori reason to believe that both rational and over-optimistic investors cannot both exist in a market over long periods of time.

In another paper, "Investor Sentiment and the Cross Section of Stock Returns", Baker and Wurgler provide further evidence of the potentially powerful impact of emotional factors on asset prices and returns. After studying the 1962 to 2001 period, they find that "the cross-section of future stock returns is conditional upon beginning-of-period proxies for investor sentiment...When sentiment appears to be high, stocks that are likely to be relatively attractive to optimists and speculators and at the same time unattractive candidates for arbitrage -- young stocks, small stocks, unprofitable stocks, non-dividend paying stocks, high-volatility stocks, extreme growth stocks, and distressed stocks -- experience low future returns relative to other stocks. On the other hand, conditional on low beginning sentiment, these

cross-sectional patterns attenuate or reverse." They go on to note that "several aspects of these results cast doubt on the hypothesis that they reflect rational compensation for bearing systematic risk. Rather, they match simple theoretical predictions and line up well with historical accounts of bubbles and crashes."

Summing up this changing view of investors, Daniel Kahneman writes in his paper "Maps of Bounded Rationality" that the emerging "model of an [economic] agent has a different architecture, which may be more difficult to translate into the theoretical language of economics. [Its] core ideas include the two system approach, the large role played by System 1 [the experiential system described above], and the extreme importance of context dependence [e.g., as exemplified by the Barberis, Huang and Santos model]. The central characteristic of agents is not that they reason poorly, but that they often act intuitively [to conserve scarce cognitive resources]. The behavior of these agents is not guided by what they are able to compute, but by what they happen to see at the moment." In his paper, "Psychology and the Market", Edward Glaeser calls this "situationalism", which he describes as "the view that people isolate decisions and overweight immediate aspects of the situation relative to longer term concerns." However, with a strong dose of common sense, he goes on to note that "outside the laboratory emotionally powerful situational factors are almost always the result of actions by [others]... who are responding to a set of incentives." He notes that "the important question for economists is not whether consumers are rational as independent actors, but rather how heterogeneously rational consumers aggregate [supply and demand], especially when rational, self-interested [parties] are trying to exploit less than perfectly rational parties [by manipulating their situational context]."

In their paper "Behavior Based Manipulation", Zhou and Mei provide an example of just such a case. Their model explicitly investigates "how a smart manipulator interacts with irrational traders and the profit the manipulator makes from exploiting other investors' behavioral biases." As they put it, "the manipulator [seeks to use situational factors] to create more chances for irrational investors to make a mistake." In essence, their paper describes the theory behind a phenomenon -- "pump and dump" -- long known to market players but heretofore ignored by economists.

There is, however, another important question to answer: do the actions of irrational investors simply cancel each other out (providing support for the strong form of the efficient markets hypothesis), or do they systematically affect asset prices? Two recent papers have directly addressed this question. In their paper "Systematic Noise", Barber, Odean and Zhu note that in order for "the biases and sentiment of individual investors to affect asset prices...their preference for buying some stocks while selling others would have to be shared." After analyzing the trading records of 66,465 households, they find that this is in fact the case, and that "the trading of individuals is more coordinated than one would expect by mere chance." They go on to reach three other conclusions: (a) "individuals buy stocks with strong past returns"; (b) "their buying is more concentrated in fewer stocks than their selling", and (c) "they are net buyers of stocks with unusually high trading volume." In another paper, "Diversification Decisions of Individual Investors and Asset Prices", Goetzmann and Kumar "show that [the portfolios of] a vast majority of individual investors [in the data sample they studied] are under-diversified" and that "the least-diversified group of investors earned 2.4% lower return annually than the most diversified on a risk-adjusted basis." They go on to show that this systematic underdiversification by individual investors influences asset prices and returns.

Let's move on now to look at the information environment in which these investors operate. This issue has been analyzed on three levels of increasing complexity. The first is a situation in which different investors hold different information about the fundamental value of an asset. These differences could be due to (a) investors using different models to interpret the same set of publicly available information; (b) different investors receiving the same public information signals at different times; or (c) some investors having access to private information (which is not necessarily illegal insider information), while others rely solely on publicly available information. The presence of these "information asymmetries" has important consequences for asset pricing.

One of the earlier papers to address this was "Differences of Opinion, Rational Arbitrage and Market Crashes" by Hong and Stein. These authors put forth a model based on two key

assumptions: (1) there exist differences of opinion between investors about the fundamental value of an asset, and (2) some of these investors face short-sale constraints (e.g., mutual fund managers who are prohibited from selling short). When differences of opinion are large (which would be evidenced by high trading volume in the asset), some of the investors with the most negative views of the asset's value will not be able to trade. As a result, the market price for the asset will be above its fundamental value. Moreover, when the opinion of the more optimistic investors reverses (e.g., due to the arrival of a piece of unexpected bad news) and the asset's price begins to fall, the investors with the most negative opinion will come back into the market, further reinforcing the decline in the asset's price. In time series of asset returns, this model predicts the negative skewness we often observe. In their more recent paper, "Forecasting Crashes: Trading Volume, Past Returns, and Conditional Skewness in Stock Prices", Chen, Hong and Stein test to see if the predictions of their model are confirmed by historical market data. They find that it does: "negative skewness is most pronounced in stocks that have experienced (a) an increase in trading volume relative to trend over the past six months, and (b) positive returns over the past thirty-six months." Another paper "Differences of Opinion and the Cross Section of Stock Returns" by Diether, Malloy, and Scherbina reaches a very similar conclusion. They find that "stocks with higher dispersion in analysts' earnings forecasts earn lower future returns than otherwise similar stocks. This effect is most pronounced in small stocks and stocks that have performed poorly over the past year." They interpret the dispersion in securities' analysts' forecasts "as a proxy for differences of opinion about a stock", and conclude that their evidence "is consistent with the hypothesis that prices will reflect the optimistic view whenever investors with the lowest valuations do not trade."

The second level of complexity revolves around the specific interaction between more and less informed investors, rather than the more general case of interaction between investors who hold differing beliefs about the correct value of an asset. In particular, a number of researchers have asked whether better informed investors will drive less informed investors from the market. In "Asymmetric Information and Survival in Financial Markets" Emanuela Scubba from the University of Cambridge describes a model in which "uninformed investors trade on the basis of the information revealed by [i.e., inferred from] market prices. They can

approximate, but never attain the information possessed by informed investors." However, better-informed investors have to pay a cost for the information that gives them an advantage. As a result, information triggers a trade-off, between its cost and its expected benefit (which is reduced through imitative trading by uninformed investors). As a result of this trade-off, uninformed investors always account for some portion of the market (rather than being eliminated by better-informed investors), and in fact have an influence on market prices, which therefore tend to vary around, rather than settle at their correct value.

In his paper "Information Asymmetry, Price Momentum, and the Disposition Effect" Gunter Strobl from Wharton shows how the combination of investors with differing amounts of private and public information who interact with other investors who trade solely for liquidity related reasons (e.g., selling shares to buy a house, not because they have any particular information about their future value) can rationally produce the momentum and disposition effects that advocates of behavioral finance believe arise from irrational sources. He notes that "there are two reasons we might expect informed investors to exhibit disposition effects [that is, a tendency to sell winners and hold onto losers]. First, selling stocks at high prices and buying stocks at low prices always seems like a good idea when some investors trade for non-informational [liquidity] reasons." In other words, to accommodate the liquidity traders, informed investors would "be willing to buy stocks at prices below fundamental value, and sell stocks at prices above expected fundamental value." In addition, informed investors "may exhibit disposition effects if their informational advantage over uninformed investors sufficiently changes." He gives two examples. In the first, the arrival of a new positive public signal causes uninformed investors to push the value of an asset above the level justified by the private information. In this case, informed investors would logically sell their winners. In contrast, the arrival of a negative public signal might push the price of a stock below the level justified by the private information. In this case, an informed investor might hold onto, or even buy more of, his apparent losers. In contrast to informed traders, who tend to behave as contrarians, uninformed traders rationally behave as trend followers. Not knowing whether an increase in prices is due to the arrival of new private information that has caused informed investors to buy, or whether this is due to the actions of liquidity traders,

they logically decide to buy when prices rise, thus creating the observed momentum effects in asset prices.

A different set of papers takes a different approach to the interaction of investors with and without access to superior information about the value of an asset. In "Is Information Risk A Determinant of Asset Returns?" Easley, Hvidkjaer, and O'Hara test the theory that while some investors may be uninformed, they are not stupid. Recognizing their disadvantage versus better informed investors, they demand higher returns on assets that have a high risk of creating information advantages (i.e., a higher probability of trading based on private information). Analyzing 1983 to 1998 data on New York Stock Exchange traded stocks, they find that "a difference of ten percentage points in the probability of information based-trading between two stocks leads to a difference in their expected returns of 2.5% per year." In short, the authors conclude that, in an efficient markets context, the probability of private information-based trading is a priced risk factor, just like value, size, momentum, and the market. In a subsequent paper "Information and the Cost of Capital" Easley and O'Hara put forth a more detail model that predicts that investors require a higher rate of return when they perceive that a greater proportion of value-relevant information about a company is private (this is termed the "composition effect"), but a lower rate of return when a greater fraction of investors receive this private information (termed the "dissemination effect") and when it is more accurate (termed the "precision effect"). In their paper "Are Information Attributes Priced?" Botosan and Plumlee test Easley and O'Hara's predictions against a set of historical data and find that they are supported.

The third level of information complexity is based on the observation that in real life, traders care not only about the fundamental value of an asset, but also about other investors' perception of its fundamental value. The importance of this dynamic has long been recognized. For example, in his 1935 book The General Theory of Employment, Interest and Money, the great economist John Maynard Keynes wrote: "It might have been supposed that competition between expert professionals, possessing judgement and knowledge beyond that of the average private investor, would correct the vagaries of the ignorant individual left to himself. It happens, however, that the energies and skill of the professional investor and

speculator are mainly occupied otherwise. For most of these persons are, in fact, largely concerned, not with making superior long-term forecasts of the probable yield [return] of an investment over its whole life, but with foreseeing changes in the conventional basis of valuation a short time ahead of the general public. They are concerned, not with what an investment is really worth to a man who buys it for keeps, but with what the market will value it at, under the influence of mass psychology three months or a year hence...To change the metaphor slightly, professional investment may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole; so that each competitor has to pick, not the faces which he himself finds the prettiest, but those which he thinks likeliest to catch the fancy of the other competitors, all of whom are looking at the problem from the same point of view. It is not a case of choosing those which, to the best of one's judgement, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest."

"We have reached the third degree where we devote our intelligence to anticipating what average opinion expects average opinion to be. And there are some, I believe, who practice the fourth, fifth, and higher degrees...If the reader interjects that there must surely be large profits to be gained from the other players in the long run by a skilled individual who, undeterred by the prevailing pastime, continues to purchase investments on the best genuine long-term expectations he can frame, he must be answered, first of all, that there are indeed such serious-minded individuals and that it makes a vast difference to an investment market whether or not they predominate in their influence over the game-players. But we must also add that there are several factors which jeopardise the predominance of such individuals in modern investment markets. Investment based on genuine long-term expectation is so difficult today as to be scarcely practicable."

It is only recently, however, that finance researchers have tried to incorporate Keynes' insight into asset pricing models, and to better understand its implications. In their paper "Beauty Contests and Iterated Expectations in Asset Markets", Allen, Morris and Shin begin noting how, in the strong efficient markets model, the price of an asset today is equal to the

discounted present value of its expected future payoffs, and the present value of an asset tomorrow is equal to the present value at that point in time of the same expected future payoffs. However, when investors have differential private information about the value of the asset, its valuation becomes more difficult: its value today is equal to the discounted present value of tomorrow's average value of expected future payoffs, which itself will reflect the same process carried out further into the future. Under these circumstances, you cannot simply fold back expected future outcomes to obtain the present value of the asset, because the average expectation tomorrow may not equal the average expectation the day after tomorrow. The authors explore the implication of this key fact for asset pricing in a world with both private and public information: "Now suppose that the individual is asked to guess what the average expectation of the asset's payoff is. Since he knows that others have observed the same public signal, the public signal is a better predictor of average opinion, and he will put more weight on the public signal than the private signal...We will then tend to have asset prices overweighting public information relative to private information...If public information suggests that future payoffs will be high, then this can lead to high asset prices even if all traders have private information or judgement that the true value will be low."

In another paper, "Imperfect Common Knowledge and the Information Value of Prices", Amato and Shin note that "there are compelling reasons to believe that the impact of higher-order beliefs on the information efficiency of prices has become more detrimental over the past decade...Public information arguably wields greater influence in private decision making than ever before due to the emergence of widely and readily accessible mass media." This is very similar to the argument made by Robert Shiller in his book Irrational Exuberance that widespread positive media stories contributed to the development of the technology stock bubble. Amato and Shin also note how the weight given to public information signals becomes relatively more important as the competition between agents (e.g., investors) increases in intensity (a situation which well describes the competition between active investment managers).

In their paper "Higher Order Expectations in Asset Pricing", Bacchetta and van Wincoop attempt to estimate the extent to which the interaction of first order public and private

information (about asset prices) and second order information (about other investors' likely beliefs and actions) can lead to a "disconnect" between market prices and fundamental asset values, which they call "the higher order wedge." They conclude that it is perfectly rational for investors to expect that next period's market expectation of future payoffs is too high or too low. Investors expect the market to make expectational errors to the extent that they expect average private signals to differ from their own. Public information plays a key role here. When an investor receives private information that is less favorable than public information, he concludes that his private signal is weak, and expects others to have more favorable private information. Consequently, when the average private signal is weaker than public information, the majority of investors expect others to have more favorable private signals than their own. If private signals today are still informative tomorrow, the majority of investors then expect the outlook of the market to be too favorable tomorrow. Investors buy the asset in anticipation of this, further pushing up its price. Thus it is the combination of both noisy [i.e., imprecise] public and private information that makes it rational for investors to expect the market to make expectational errors." This creates the conditions for the existence of "a higher order wedge that disconnects an asset's market price from the present value of its future payoffs...Overly favorable public signals generally lead to both an overestimation of the present value of future payoffs and an overestimate of average private signals about those future payoffs. The latter leads to a positive higher order wedge, which therefore amplifies the impact on the asset price of expectational errors about future payoffs." Finally, the authors show that "the size of the wedge is largest for intermediate degrees at the quality of private information." When private information signals are very strong, public information is given low weight; similarly, when private information is very weak, public information receives most of the weight. In both cases, the potential size of the wedge is reduced, and asset prices are more reflective of their fundamental value.

Bacchetta and van Wincoop have also written another paper ("Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?") which compares the predictions of their model with actual market data. They begin by noting that "empirical evidence shows that macroeconomic fundamentals have little explanatory power for nominal exchange rates" in the short to medium run. They then describe a model that contains two types of

heterogeneity: differential information about the macroeconomic fundamentals, and two different types of traders, who either trade based on information or solely for liquidity needs. They show that "information dispersion leads to the magnification and persistence" of the impact of liquidity traders' actions, and that "higher order expectations...partly contribute to these results."

A final paper on this subject is "Stock Market Manipulations" by Aggarwal and Wu. It presents "a theory about and empirical evidence for the existence of stock market manipulations in the United States." In particular, the authors "consider what happens when a manipulator can trade in the presence of other traders who seek out information about a stock's true value." They conclude "in a market with manipulators, information seekers play an ambiguous role...making it easier for a manipulator to enter the market and potentially worsening market efficiency." In short, the presence of information seekers enables the manipulator "to distort a stock price away from its true value and profitably trade on this distortion."

Closely related to information-based issues are those related to liquidity. Financial markets provide both liquidity (to facilitate trades) and price discovery (to provide information). Hence, the increased focus on the role of information in financial markets is very closely related to a growing amount of research on the role of liquidity. As in the case of information, it is hard to develop a simple measure of liquidity, because at minimum it includes three concepts: (1) the breadth of the spread between prevailing bid and ask prices; (2) the quantity of an asset that is offered at the bid and ask price, which is also known as "depth", and (3) the extent to which trades of different sizes force changes in prevailing breadth and depth. This aspect of liquidity is also known as the "resilience" of a market. Three different liquidity related issues have received a lot of research attention. First, over very short (trade by trade) time frames, how does liquidity behave? Second, what causes liquidity crises? And third, is liquidity in fact a systematic risk factor that, like value, size, momentum, and perhaps information is also priced by rational investors in an efficient market? Let's look at each of these in turn.

The study of trade-by-trade data is also known as "market microstructure." To simplify matters, most of the studies in this area assume the existence of two kinds of trade. The first is a market order, which is executed immediately at the prevailing bid or ask price. The second is a limit order, which specifies the amount and price at which it can be executed. When a limit order is received, it is compared to the prevailing market price. If it matches, it is immediately executed. If it does not, it is stored in the limit order book, and held until it is either cancelled or executed as a result of changes in the market price. For example, consider a market where the current limit order book contained three sell orders: 50 shares at \$10, 25 shares at \$11, and 50 shares at \$12. What will happen when a market buy order for 75 shares arrives? Assuming no other orders arrive, 50 shares will be sold at \$10 and 25 will be sold at \$11, which in turn will force the prevailing ask or offer price to rise to \$12. An interesting question to ask is how realistic is this example? In other words, what causes large price changes in a market: is it the arrival of market orders, or is it changes in the limit order book (which, in essence, represents the current state of market liquidity)?

In their paper "Large Stock Price Changes: Volume or Liquidity?" Weber and Rosenow study just this question, and find that extreme price changes are caused by "a low density of orders stored in the limit order book... that is, by time changing liquidity rather than by large fluctuations in trading volume." Using a different data set, Farmer et al reach the same conclusion in their paper "What Really Causes Large Price Changes?" They also find that "the tail exponent of large price changes displays [which is closely related to the extent of the kurtosis in a stock's return series] appears to depend on parameters of the market: more lightly traded markets tend to display fatter tails, with more extreme risks. This has important practical implications, because it gives some understanding of what determines financial risks."

In a subsequent paper, "The Long Memory of the Efficient Market", Lillo and Farmer use historical data from the London Stock Exchange, and note the interesting finding that the signs of orders received in a market (be they market, limit, or cancellations) are positively correlated. "However, this predictability does not apply to price movements because liquidity [in the limit order book] and market order size compensate in an anti-correlated manner" [that

is, they react in the opposite direction, rising in response to orders with rising prices, and falling in response to orders with falling prices]. "As a result, the volume of orders...is also a long memory process. Thus, despite the striking predictability of every feature of the market except prices, the market nonetheless appears to remain roughly linearly efficient." Similar findings are reported by Chordia, Roll and Subrahmanyam in an analysis of New York Stock Exchange transactions published in their paper "Evidence on the Speed of Convergence to Market Efficiency." Their explanation of the process at work is as follows: "We interpret these results to reveal the actions of three distinct groups. Order imbalances in the first instance arise from traders who demand immediacy for liquidity or informational needs. Order imbalances are positively autocorrelated, which suggests either that traders are herding or spreading their orders out over time. Second, NYSE specialists react to initial order imbalances by altering their price quotes away from fundamental value in an effort to control inventory [and their risk exposure]. Finally, outside arbitrageurs (by way of market or limit orders) intervene to add market-making capacity by conducting countervailing trades in the direction opposite the initial order imbalances. This arbitrage activity takes at least a few minutes to start."

Apart from reinforcing an impression of short-term market efficiency, these findings raise an important question: what causes liquidity providers to respond in the opposite direction from the movement of prices? This quickly gets us back to the role of information. Bloomfield, O'Hara and Saar examine this issue in their paper "The 'Make or Take' Decision in An Electronic Market: Evidence on the Evolution of Liquidity." They report the results of an experimental market which "contains both informed traders who have superior information and liquidity traders who face both large and small liquidity needs." They find that "liquidity provision in a market changes dramatically over time, and the key to this evolution is found in the behavior of the informed traders. When trading begins, informed traders are much more likely to take liquidity, hitting existing limit orders so as to profit from their private information. As prices move toward true values, the informed traders shift to submitting limit orders...[In other words] when the value of their [private] information is low, they move very quickly to assume the role of dealers and trade primarily by supplying limit orders to the market... In an asymmetric information setting, it is the informed traders who ultimately have

the risk advantage because they know more about where the price should be. Thus, a market-making role arises endogenously in our electronic markets, adopted by traders for whom the risk of entering a limit order is lower than it is for other traders [i.e., uninformed traders whose limit orders can be exploited by traders with private information]."

Another approach to exploring this issue is to examine liquidity crises. In their paper "Liquidity Black Holes", Morris and Shin develop a model in which "traders with short horizons and privately known loss limits interact in a market for a risky asset...When the price of the asset falls close to the loss limit of the short horizon traders, selling of the risky asset by any trader increases the incentives for others to sell." The resulting "liquidity black hole is analogous to a run on the bank." Clearly, the increasing loss aversion model posited by Barberis, Huang and Santos would only accentuate this process. In "Predatory Trading", Brunnermeier and Pedersen make a similar point. They study "trading that induces and/or exploits other investors' need to reduce their positions." They "show that if one trader needs to sell, others also sell and subsequently buy back the asset at a lower price. This leads to price overshooting and a reduced liquidation value for the distressed trader. Hence the market becomes illiquid when liquidity is most needed." Finally, in "Model Uncertainty and Liquidity" Routledge and Zin start with the distinction between risk (which can be quantified) and uncertainty (which cannot be quantified). Given the importance of models in the valuation of assets, they suggest that an increase in uncertainty in effect reduces an uncertainty-averse trader's confidence in his model's outputs, and forces him to widen bid/ask spreads (i.e., to reduce the amount of liquidity he provides) to compensate.

All of these analyses of liquidity issues beg the question of whether or not it is a systematically priced risk factor in an efficient market. One of the earlier papers on this was "Liquidity Risk and Expected Stock Returns" by Pastor and Stambaugh, who concluded that stock returns were indeed "systematically related to fluctuations in aggregate liquidity." Moreover, based on their study of 34 years of data, they concluded "the average return on stocks with high sensitivities to liquidity exceeds that for stocks with low sensitivities by 7.5% annually, after adjusting for exposure to the market return, as well as the size, value and momentum factors." More recently, in his paper "Is There a Global Liquidity Factor?"

Christof Stahel concludes that "global and country-specific commonalities dominate industry effects as the source of common variation in liquidity" and that "global liquidity is a priced risk factor on [both] the portfolio and individual stock level." In "Flight to Quality, Flight to Liquidity, and the Pricing of Risk" Dimitri Vayanos presents a model in which the "preference for liquidity is time-varying and increasing with asset volatility...During volatile times, assets' illiquidity premia increase, investors become more risk averse, [and] asset [prices] become more negatively correlated with volatility." In short, illiquid assets are more sensitive to volatility, which as he notes, "has implications for evaluating the performance of strategies [such as those pursued by many hedge funds] which include investing in illiquid assets." He also finds that "an important factor driving the variation in liquidity premia seems to be the extent of uncertainty in the market."

Going slightly further, in "Asset Pricing with Liquidity Risk" Acharya and Pedersen attempt to integrate liquidity into a new version of the Capital Asset Pricing Model, while Lo, Petrov, and Wierzbicki integrate it into an asset allocation model in their paper "It's 11pm -- Do You Know Where Your Liquidity Is?"

The research findings we have reviewed on information and liquidity also have implications for the extent to which arbitrage takes place in a market to maintain its efficiency. In their paper "Synchronization Risk and Delayed Arbitrage" Abreu and Bruunermeier argue that "arbitrage is limited if rational traders face uncertainty about when their peers will exploit a common arbitrage opportunity. This synchronization risk arises because arbitrageurs become sequentially aware of the mispricing and incur holding costs [when they trade against it]." Given this, "rational arbitrageurs try to "time the market" rather than correct the mispricing right away, which leads to delayed arbitrage." They also show how the arrival of a strong public signal can spur arbitrageurs into collective action. Drawing from history, Temin and Voth make the same point in their paper "Riding the South Sea Bubble." They show how in 1720 Hoare's Bank in London "knew that a bubble was in progress, and that it invested knowingly in the bubble" because "it was profitable to ride the bubble" at least until a public signal prompts coordinated action by arbitrageurs. Finally, in "Aggregate Short Interest and Market Valuations" Lamont and Stein present further evidence on the limited impact of

arbitrage. They conclude that "arbitrageurs are reluctant to bet against aggregate market mispricings" and that "short selling does not play a particularly helpful role in stabilizing the overall stock market."

Taken together, all of these research findings paint a rich picture of financial markets populated by investors who possess different mixes of public, private, and higher order information, who sometimes act in ways meant to manipulate and exploit each other, and whose decision making results from a complicated and shifting mix of cognitive and emotional factors. This gives rise to periods of over and undervaluation which arbitrageurs do not immediately correct, as well as fluctuations in uncertainty and occasional liquidity crises. In short, they paint a picture of a complex adaptive system, whose dynamics are constantly evolving. Two further papers take this as their starting point, and explore its consequences. In their paper "Financial Markets Can Be At Sub-Optimal Equilibria", Joshi, Parker, and Bedau use an agent based simulation model of an evolving stock market that "consists of boundedly rational traders who learn through their market experience, continually adapting their behavior to changing market conditions."

These traders "lack perfect foresight about what other traders will think and do. At the same time, since a trader's profits depend on the behavior of other traders, each trader makes investment decisions based on the basis of her best guess about what other traders will be thinking and doing." They find that "a market operating under these conditions is a complex adaptive system consisting of a co-evolving ecology of heterogenous traders. A central factor governing the behavior of such markets is the rate at which traders revise and adapt their market forecasting methods. This revision rate determines the market's behavior, since different forecast revision rates promote the use of different kinds of forecasting methods in the population." In their model, a slow revision rate corresponds with a greater focus on the fundamental value of the asset, while faster revision rates imply greater use of technical trading rules that aim to take advantage of the actions of other traders (e.g., momentum, etc.). As the authors point out, "as more traders adopt technical trading rules, the incentives for their use can reinforce themselves, with positive feedback making them self-fulfilling prophecies." However, their widespread use "worsens everybody's forecasts by generating positive

feedback in prices, driving them away from their fundamental value and increasing noise [that is, the presence of signals with low information value]...When all traders engage in significant technical trading, they worsen each others' forecasts, causing a loss of efficiency, and a reduction in the traders' average earnings."

However, after multiple simulations of their market, the authors find that "a rapid forecast revision rate is the strategic equilibrium in which all traders revise their forecasts at the same rate." However, they also find that "in this strategic equilibrium the market is noisy and risky, with a high variance in asset prices and high levels of technical trading." Traders collectively would have higher average earnings if they slowed down their rate of forecast revision, and in effect focused more on the asset pricing fundamentals. However, this fails to happen in the model because of the potential advantages accruing to a self-interested trader who, under these circumstances, can increase his individual profits by defecting and using a technical trading rule.

In "Evolutionary Stable Stock Markets", Evstigneev, Hens and Hoppe employ a similar model, but, in addition to allowing for the mutation (adaptation) of trader strategies, also introduces selection pressure based on trader's results (i.e., traders who don't make a minimum amount of return lose their capital and exit the game). They encounter many of the same dynamics as Joshi et al, but with a critical difference. Over the long term the introduction of selection pressure favors those investors whose strategies are focused on fundamental value rather than technical trading rules. Their findings are consistent with one of the most famous quotes from Ben Graham, the father of value investing: "In the short term the market is a voting machine. In the long term it is a weighing machine."

This still leaves us with a final question: in light of these findings, can we still say that financial markets are efficient, and that the case for index investing remains intact? The advocates of active management would clearly like people to believe that the answer is "no and no." Unfortunately (for them) there is no evidence that they are correct. To be sure, accumulated research findings lead to the conclusion that the market is not strongly efficient. However, just as a Scottish "not proven" verdict is not the same as "innocent", the fact that

efficient markets theory has been undermined does not similarly invalidate the case for indexing. The original case for indexing was based on the conclusion that long term active management success was impossible because of the markets' efficiency. In recent years, as the strong version of the efficient markets theory declined, there arose an argument that asset prices were slightly predictable, and that successful active management was therefore possible. More recently, this view has been widely called into question. For example, in his paper "Dynamic Strategies and Asset Pricing Models", Cesare Robotti of the Federal Reserve Bank of Atlanta finds that "active investment strategies based on conditioning information strongly dominate passive [index] strategies in-sample [that is, over the period covered by the data used to construct the asset pricing model]." However, he also finds that active strategies "do not provide any convincing pattern of improved out-of-sample performance." That is, when you go beyond "back-testing" and actually try to use these forecasting models in the real world, you still can't beat your index fund competitors. But if efficient markets theory is dead, how can this be? From our perspective, the complex adaptive system view of financial markets provides the answer. One of the essential facts about complex adaptive systems is that changes in their behavior are non-linear; causes and effects can be widely separated in time and scale. They are also dynamic, experiencing periods of both stability and something akin to chaotic changes; usually they self-organize themselves to the border between these two regions. This makes it next to impossible to consistently forecast their behavior (although this may be possible to do, often in hindsight, over limited periods of time). To be sure, some people (and software programs) are better than others at developing an "intuitive feel" for the dynamics of a complex dynamic system and recognizing some of the patterns they create; however, there is no guarantee that this knowledge will remain relevant as the system itself continues to evolve.

For these reasons, successful long-term active management (based on skill) is still extremely rare, and difficult if not impossible for most mere mortals to identify in advance. The market is still efficient, in the sense that it is still extremely difficult to consistently make money by exploiting its inefficiencies. In sum, the case for indexing remains solidly intact, not because financial markets are efficient in the classical sense, but because their complexity now exceeds the capacity of man and machine to predict their future path.

Annual Research Review Part 2: Product and Strategy Implications

In this section of we will review a number of research studies with direct bearing on various products and strategies.

In their paper "Diversification Decisions of Individual Investors and Asset Prices", Goetzmann and Kumar find that "a vast majority of individual investors in our sample are underdiversified...More than 25% of investor portfolios in our sample contain only one stock, more than 50% of them contain fewer than three stocks, and in any given month only 5% to 10% of the portfolios contain more than ten stocks. As a consequence, investor portfolios have extremely high volatility (more than 75% of investor portfolios have higher volatility than the market portfolio) and they exhibit worse risk-return trade-off than randomly constructed portfolios." They also find that "the least diversified group of investors earn 2.4% lower return annually over the 1991-1996 period than the most diversified group on a risk-adjusted basis."

In "Downside Risk and Asset Pricing", Post and van Vliet starts with a simple question: in light of various research findings on the existence of non-market related risk factors (e.g., the size, value, and momentum effects), why do so many investors still put their money in market index funds? Using classical mean/variance optimization, this portfolio is less efficient than portfolios that include the other factors. To answer this question, the authors employ a criteria called second order stochastic dominance, or SSD for short. Without going into the gory statistical details, in determining the optimality of a portfolio, SSD takes into account the impact of skewness and kurtosis in addition to mean and variance. Using this criteria, they solve the mystery: the broad market index turns out to be superior, using SSD, than portfolios using size, value, and momentum that are optimal under the mean/variance criteria.

Past research has shown that individual investors place a heavy emphasis on past performance when deciding how to allocate their savings across different investment products. In "Inferences Regarding Investment Allocation Decisions in the Institutional Plan Sponsor

Market" by Heisler, Knittel, Neumann and Stewart the authors find that a similar focus on past performance does not characterize the decisions made by institutional plan sponsors. They note that "the consistency with which investment managers deliver active returns over multiple horizon, without regard to the magnitude of those returns relative to the S&P 500 plays a key role in determining the flow of assets and accounts among investment products." They also find that "the magnitude of [any] one year loss, as well as 3 and 5 year total returns are incremental factors in plan sponsor's allocation decisions."

In the past, we have written about how the careful division of investments between taxable and tax exempt accounts can materially improve long term returns. In "Optimal Asset Location and Allocation with Taxable and Tax Deferred Investing", Dammon, Spatt, and Zhang reach the same conclusion (but note that they only use bonds and equities as asset classes). They conclude that it is better to hold taxable bonds in the tax-deferred account, and equity in the taxable account. However, "it may not be optimal to allocate the entire tax-deferred account to taxable bonds if doing so causes the overall portfolio to be over-weighted in bonds...In this case, investors may hold a mix of stocks and bonds in their tax-deferred account, but only if they hold an all-equity portfolio in their taxable account." Similarly, an investor with a higher allocation to bonds may want to hold some of them (ideally tax-exempt bonds) in their taxable account, assuming their tax-deferred account already holds only bonds.

While we're on the subject of bonds, Ang and Bekaert have produced a very interesting paper on "The Term Structure of Real Rates and Expected Inflation." They find that the unconditional term structure of real interest rates is quite flat, "starting at a rate of about 1.7% and increasing to just over 1.8% at one year, before declining again to 1.7% at the five year maturity." Also, in their paper "Are Treasury Inflation Protected Securities Really Tax Disadvantaged?" Hein and Mercer from the Federal Reserve Bank of Atlanta show that they are not, and in fact on an after-tax basis have outperformed matched maturity conventional (nominal return) Treasury securities.

Moving on to another asset class, Nijman and Swinkels have a very interesting paper titled "Strategic and Tactical Allocation to Commodities for Retirement Savings Schemes." They

"find substantial differences in optimal strategic [commodities] allocations for pension plans with nominal and inflation-indexed liabilities. In the latter, commodities reduce the risk on the funding ratio by more than 30 percent."

In "Private Equity Performance", Kaplan and Schoar provide more information on the risks and returns of investing in venture capital and buyout funds. They begin by noting the self-selection bias in the Venture Economics data series they use: roughly fifty percent of the funds reported raised do not provide performance data. However, using the performance data for the remaining funds that do report it, they find that "on average, LBO fund returns net of fees are lower than those on the S&P 500; VC fund returns are lower than the S&P 500 on an equal weighted basis, but higher than the S&P 500 on a capital weighted basis." They also "document substantial persistence in fund performance in the private equity industry, for both LBO and VC funds." They attribute this to these funds' proprietary access to a flow of new investment opportunities, as well as to differences in their ability to add value to their respective investments. The authors also caution that "funds raised in boom times (and partnerships that are started during booms) are less likely to raise follow-on funds, indicating that these funds likely perform poorly."

Two other papers look at the choice between different types of hedge funds: "Fund of Hedge Funds Portfolio Selection" by Davies, Kat and Lu, and "Portfolios With Hedge Funds" by Chen, Feldman, and Goda. Both papers reach the same conclusion we did in our analysis of this asset class: Equity Market Neutral and Global Macro funds seem to provide the most benefit to portfolios that include a wide range of other asset classes.

Moving on to the analysis of specific investment products, in "Predictable Investment Horizons and Wealth Transfers Among Mutual Fund Shareholders", Woodrow Johnson makes three important points. First, the liquidity costs associated with the presence of both short and long term investors in a mutual fund can be expensive to the latter. He provides an estimate for one fund in which it amounts to .51% (51 basis points) per year in foregone returns. Second, there are factors that can be used to identify in advance mutual fund shareholders' likely holding period. Third, this argues strongly for either different funds for

each group, or for the imposition of additional fees which will avoid the transfer of wealth caused by this liquidity effect from long term to short term holders.

Last but certainly not least, Elton, Gruber and Blake have written a very important paper on "The Adequacy of Investment Choices Offered by 401K Plans." After examining over 400 plans, they conclude that "for 62% of the plans, the types of choices offered to plan participants are inadequate, and that over a twenty year period this makes a difference in terminal wealth of over 300%." Stunning. The authors also make a very strong case for offering plan participants index funds that track a wider range of asset classes. As they note, "investors in 401K plans are sacrificing significant return [they estimate 3.16% per year] because plan administrators are offering an incomplete set of investment alternatives." The authors also find that "funds included in the plans are riskier than the general population of funds in the same category." Specifically, "plan administrators offer plan participants mutual funds with less variance than randomly selected funds, but funds that are more highly correlated." Moreover, when the authors "examine one category of investment choices, S&P 500 index funds, they find that the index funds chosen by 401K plan administrators are on average inferior to the S&P 500 index funds selected by the aggregate of all investors." In sum, this paper further reinforces our long-held opinion that in the world of defined contribution pension plans there is a great need for more "prudent experts."