

Market Efficiency in an Irrational World*

Kent Daniel
Kellogg School of Management
and NBER
Northwestern University
2001 Sheridan Road
Evanston, IL 60208-2006
847-491-4113
kentd@nwu.edu

Sheridan Titman
University of Texas at Austin
and NBER
College of Business Administration
University of Texas
Austin, Texas 78712-1179
512-232-2787
titman@mail.utexas.edu

Abstract:

This paper explains why investors are likely to be overconfident and how this behavioral bias affects investment decisions. Our analysis suggests that investor overconfidence can potentially generate stock return momentum and that this momentum effect is likely to be the strongest in those stocks whose valuation requires the interpretation of ambiguous information. Consistent with this, we find that momentum effects are stronger for growth stocks than value stocks. A portfolio strategy based on this hypothesis generates strong abnormal returns that do not appear to be attributable to risk. Although these results violate the traditional efficient markets hypothesis, they do not necessarily imply that rational but uninformed investors, without the benefit of hindsight, could have actually achieved the returns. We argue that to examine whether unexploited profit opportunities exist, one must test for what we call “adaptive-efficiency,” which is a somewhat weaker form of market efficiency that allows for the appearance of profit opportunities in historical data, but requires these profit opportunities to dissipate when they become apparent. Our tests reject this notion of adaptive-efficiency.

*We thank David Chapman, Savjiv Das, David Hirshleifer, Ravi Jagannathan, Hong Yan and an anonymous referee for helpful comments.

1 Introduction

The traditional economic paradigm assumes that individuals are “rational,” meaning they make optimal decisions based on the information available to them. In the field of asset-pricing, the implications derived from this view, described in papers by Grossman and Stiglitz (1980) and others, are that prices reflect all available information, so that opportunities to earn extraordinary returns arise only from private information. According to this view, investors without special information, no matter how savvy they may be, cannot improve on the performance of index portfolios.¹

The recent behavioral finance literature offers an alternative paradigm in which individuals make systematic mistakes in the way that they process information. In this paper, we consider one of these biases, that of *overconfidence*. There are a number of reasons why we concentrate on this particular bias: first, it is perhaps the best established of these biases. Second, we view it as a bias that is likely to manifest itself in the sort of analysis that is necessary in security valuation. Finally, we argue that overconfidence is a trait that is likely to arise through evolutionary selection. We believe this last argument is extremely important. Behavioral biases that distort decisions with no offsetting benefits are likely to be eliminated by natural selection. Hence, we should be considerably more skeptical about the existence of biases that cannot be explained in this way.

Although many economists are sympathetic to the view that behavioral biases play a role in economic decisions, they generally believe that irrational investors have only a minor effect on prices. The standard argument is that, in competing to take advantage of the profit opportunities created by the trades of irrational investors, rational investors will push prices to a level where the profit opportunities virtually disappear. Thus, in the end, prices will be determined in the market “as if” all investors are rational.

However, there is now considerable evidence at odds with this “as if” hypothesis. Daniel, Hirshleifer, and Subrahmanyam (1998) summarize this evidence and argue that,

¹Roberts (1967) and Fama (1970) lay out the definitions of *weak*, *semi-strong*, and *strong*-form efficiency. There is considerable evidence that is at odds with the *strong-form* efficient markets hypothesis that investors with private information cannot make extraordinary profits. When we refer to traditional market-efficiency, we mean *semi-strong* form efficiency.

taken as a whole, the data is more consistent with an overconfidence model than a standard rational model. The empirical evidence presented in this paper provides still more evidence inconsistent with the “as if” proposition. In particular, we find that a portfolio strategy suggested by the overconfidence theory consistently realized excess returns in the 1963 to 1997 period. This long-short (zero cost) portfolio strategy, which is based on market capitalization, momentum, and book-to-market ratios, has a Sharpe ratio that is more than three times as high as that of the market and realizes positive returns in virtually every year.

The returns from this strategy cannot be explained by traditional measures of risk; indeed, the betas of the strategies are negative. Also the returns cannot be explained by the popular Fama and French (1993) three-factor model.² Moreover, they cannot be explained by less traditional risk factors that measure the degree to which these portfolios contribute to the tracking error of portfolio managers who are evaluated relative to a value-weighted index like the S&P 500. In addition, we argue that it is unlikely that the strong performance of these strategies were spuriously uncovered as a result of the data mining efforts of the hundreds of academics and practitioners who have analyzed this data.

Although this evidence indicates that overconfident investors systematically influence market prices, it does not necessarily imply that the past 30 years offered easy profit opportunities to quantitative investors who were not inflicted with behavioral biases. As we will later argue, anticipating and exploiting apparent profit opportunities would have required an implausible level of knowledge about the rationality and trading strategies of other investors. In reality, most investors will have incomplete knowledge of the tastes, strategies, and cognitive limitations of other investors and, like academics, can only imperfectly infer these from the past pattern of returns.

To account for the fact that investors must learn about pricing anomalies from historical prices we introduce a somewhat weaker concept of market efficiency that we call

²While we do not wish to rule out the possibility that some model with full rationality might explain the returns of this strategy, we note that, given the high Sharpe-ratio of the portfolio, such a model would require investors with very peculiar preferences, based on the arguments of Hansen and Jagannathan (1991).

adaptive-efficiency. In a market that is adaptive-efficient, pricing anomalies may be observed in historical data. However, if investors do indeed learn from the past price history, these anomalies will not persist for too long. To test for adaptive efficiency we examine the return pattern of a meta-strategy that mechanically picks strategies based on the performance in the previous 10 years. As we show, this meta-strategy has continued to earn consistent profits, rejecting our notion of adaptive-efficiency.

2 Behavioral Finance

The psychology literature describes a myriad of behavioral biases that can potentially explain almost any pricing anomaly one might imagine. However, we think the most prominent anomalies can be explained by overconfidence, which is consistent with anecdotal evidence as well as experimental evidence.

In this section, we summarize the experimental evidence on overconfidence and then discuss why individuals are likely to have evolved in ways that make them overconfident. In our view, presenting an evolutionary rationale for overconfidence is extremely important. We would expect a trait like overconfidence to be eliminated through natural selection if it simply resulted in irrational decisions. However, since overconfidence enhances genetic survival in other ways, it makes the experimental evidence demonstrating overconfidence that much more convincing.

This section also explains why the previously mentioned book-to-market and momentum effects might arise because of overconfidence and suggests a more refined trading strategy that might better exploit this overconfidence bias. This refinement is based on the idea that the momentum effect is likely to be more important for growth stocks than value stocks, which is the motivation for the empirical analysis in Section 3.

2.1 Overconfidence

Many economists would agree that their notion of rationality should not be taken too literally. First, this notion implicitly assumes that individuals have an unlimited ability

to both observe and process information. In reality, of course, individuals have limited processing ability, and hence use vague, *ad hoc* rules to translate the information they receive into estimates of cash flows and firm valuations.³ For example, investors can't really incorporate the news about the antitrust proceedings against Microsoft into concrete views about the future competitiveness of the industry and how this in turn will affect Microsoft's future cash flows. Instead, they must do much of their analysis based on "hunches" or "feelings," which can easily be influenced by behavioral biases.

One of the most strongly documented behavioral biases is overconfidence. In their summary of the microfoundations of behavioral finance, DeBondt and Thaler (1995) state that "perhaps the most robust finding in the psychology of judgment is that people are overconfident." Psychologists, physicians and nurses, engineers, attorneys, negotiators, entrepreneurs, managers, investment bankers, and market professionals such as security analysts and economic forecasters all exhibit overconfidence.⁴ Moreover, some evidence suggests that experts tend to be more overconfident than relatively inexperienced individuals (Griffin and Tversky (1992)).

Experimental evidence also suggests that the degree to which individuals are overconfident depends on the situation. Overconfidence is generally stronger for more diffuse tasks for which feedback is slow (*e.g.*, making diagnoses of illnesses), as opposed to more mechanical tasks (*e.g.*, solving arithmetic problems) which provide immediate and conclusive outcome feedback, such as weather forecasting or horse-racing handicapping, (see, *e.g.*, Einhorn (1980)).

While there are clearly disadvantages associated with overconfidence, there are offsetting benefits that suggest that, overall, overconfidence might increase the chances of individuals passing on their genes. The evolutionary theories suggest that those individuals who *appear* to be the strongest and the smartest, *i.e.* are more likely to attract mates and reproduce.⁵

³This is closely related to some of the ideas of Herbert A. Simon. See, for example, Simon (1982).

⁴See Odean (1998b) for a good summary of empirical research on overconfidence.

⁵In the biology literature, Trivers (1985, 1991) provides arguments along these lines. In the economics literature, Waldman (1994) and Hirshleifer (1999) shows how overconfidence among males can be an evolutionarily stable strategy.

For similar reasons, appearing to be confident might enhance shorter-term economic survival. Even in the money management business where results are easy to measure and reward, we expect that holding past investment performance equal, portfolio managers who appear more confident will more successfully attract clients. The ability to act as though one is smart and strong is therefore a survival tendency, which provides a comparative advantage to individuals with inflated opinions of themselves.^{6,7}

An important ingredient in this theory is that individuals can better fool others about their abilities if they can first fool themselves. In other words, self-confident individuals will appear to be more competent than individuals who are insecure about their own abilities. As a result, individuals who successfully filter information in ways that add to their self-confidence may, in theory, be more successful than individuals who always interpret information rationally.

2.2 Rational Behavior

Before we discuss how overconfidence biases decision making, it is useful to first describe what economists mean by rational decision making. When a rational investor values a stock, like Microsoft, he or she must combine information from a number of different sources. They must, for example, combine new information with existing information, and information they collect on their own with information provided by others, like a firm's management and accountants. Rational investors combine these different sources of information using what is known as Bayes' rule, which specifies that the weights placed on the different pieces of information should be proportional to their respective precisions.

⁶Of course, there is a cost associated with having an opinion of oneself that is too inflated, which suggests that there is an optimal level of overconfidence. Hirshleifer (1999) presents an equilibrium model that formalizes the tradeoff between the benefits and costs of overconfidence.

⁷Another argument as to why overconfident traders will not influence prices is that their irrationality will cause them to lose money to rationals, and without much wealth they will not have much influence on market prices. However, agents who are more confident may trade more aggressively and accumulate wealth more quickly (De Long, Shleifer, Summers, and Waldmann (1991)). Also, attribution bias might lead successful traders to be more overconfident (Gervais and Odean (1998)). Finally, feedback will be very noisy; informed agents make money because they have valuable information – they lose money because they trade too aggressively. As long as they are not too overconfident they will still make money (Daniel, Hirshleifer, and Subrahmanyam (1998)).

Of course, an investor's ability to process information is limited. As a result, investors will probably use *ad-hoc* rules to combine their different sources of information, and will therefore undoubtedly make "mistakes" in this process. However, these *ad-hoc* rules should on average provide a good approximation to Bayesian analysis if there is no systematic bias in the way that investors make mistakes. In contrast, as we discuss in the next section, an overconfident investor makes systematic mistakes that bias the values they place on stocks.

2.3 How Overconfidence Biases Decisions

Overconfidence has both a direct and an indirect effect on how individuals process information. The direct effect, discussed in Daniel, Hirshleifer, and Subrahmanyam (1998), is simply that individuals place too much weight on information they collect themselves since they tend to overestimate the precision of that information. This direct overconfidence effect can explain the book-to-market effect, since overconfident investors tend to underweight information provided by accountants.⁸ However, as shown in Daniel, Hirshleifer, and Subrahmanyam (1998), the direct effect of overconfidence cannot explain the momentum effect.

The indirect effect of overconfidence arises because individuals filter information and bias their behavior in ways that allow them to maintain their confidence. Psychologists have developed theories of what they call cognitive dissonance, attribution bias and conservatism bias to describe this type of behavior. In summary, they tend to ignore, or at least underweight, information that one way or another lowers their self-esteem. For example, investors may be reluctant to sell their losers since this requires that they in some sense admit to making a mistake (see Odean (1998a)), which could lead to a loss in confidence which would have other deleterious consequences. For similar reasons, investors may systematically overweight information that tends to support earlier decisions and to filter out information that suggests that earlier decisions were a mistake.

⁸Daniel, Hirshleifer, and Subrahmanyam (1998) show that this argument also provides an explanation for evidence that suggests that investors underreact to corporate disclosures and decisions

Daniel, Hirshleifer, and Subrahmanyam (1998) and Barberis, Shleifer, and Vishny (1998) discuss how these kinds of biases can explain the momentum effect. Interestingly, these models illustrate how overconfidence can generate both overreaction and underreaction and how both can be consistent with the momentum effect that we observe.⁹ Specifically, Barberis, Shleifer, and Vishny (1998) describe a conservatism bias that results in investors overweighting their prior beliefs and thereby underreacting to new information.¹⁰ In Daniel, Hirshleifer, and Subrahmanyam (1998) investors first estimate an investment's value and then estimate the precision of their valuation. Because of attribution bias, they overweight information that confirms their original estimate, and underweight information inconsistent with their views for some time. As a result, their estimate of the precision of their valuation increases over time, leading to momentum as a sort of "delayed overreaction."¹¹

Overconfidence is not likely to equally bias the pricing of all securities. As we mentioned in the last subsection, experimental evidence suggests that overconfidence is likely to influence the judgment of investors relatively more when they are analyzing relatively vague, subjective information. The subjectivity of the information used to value investments is likely to vary cross-sectionally. Consider, for example, a firm such as a REIT with stable existing operations and few growth options. The information used to value such a firm is likely to be fairly concrete, and thus the pricing bias resulting from investor overconfidence should be minimal. In contrast, to value a firm like Amazon.com or Microsoft, whose values are likely to be much more dependent on future growth options and

⁹Hong and Stein (1999) develop a different model in which the interaction of bounded-rationality agents ("newswatchers" and "momentum traders") and slow diffusion of information generates initial underreaction and subsequent overreaction.

¹⁰Barberis, Shleifer, and Vishny (1998) also discuss experimental evidence of what is known as a representativeness bias, which they argue can exist when investors observe larger amounts of new information. For example, if a value firm experiences three consecutive earnings increases, investors may overreact, reclassifying the former value firm as a growth firm. In other words, because of conservatism bias, investors overreact to the initial earnings increase, but because of their representativeness bias, they overreact to later earnings increases. We can explain conservatism bias within the context of evolutionary theories of overconfidence; however, we are unaware of how representativeness bias can be fit within this framework.

¹¹An interesting behavioral study demonstrating biases consistent with this phenomenon is Lord, Ross, and Lepper (1979), who find that individuals' views become more polarized as they process more (unbiased) information.

intangible assets, one relies on much more subjective information. For these type of firms the overconfidence-related mispricing effects should be stronger

To develop a portfolio strategy based on this hypothesis, we need some proxy for the concreteness of the information used in performing a valuation. One measure of this will be the fraction of a firm's value which is in growth options, and as a measure of this we can use a firm's book-to-market ratio. To the extent that lower book-to-market firms have more growth options, they should exhibit stronger overconfidence effects. In particular, this means that momentum effects should be stronger for growth stocks.¹² This is the motivation for the empirical work that follows.

3 Empirical Evidence

As we discussed in the previous section, if overconfidence affects stock prices, then we expect to observe a relation between stock returns and both the book-to-market ratio and momentum. Although, previous research has documented that stock returns are indeed related to these characteristics, we will be considering this evidence in more detail in this section.

We start our analysis by examining the performance of 125 size, book-to-market and momentum sorted portfolios. Our construction of these portfolios is close to that of Daniel, Grinblatt, Titman, and Wermers (1997), with some differences; details are given in the Appendix. In summary, to form these 125 portfolios we group the universe of listed common stocks from the NYSE, AMEX and NASDAQ into three quintile groupings based on market capitalization, book-to-market ratio and prior-year return.¹³ The 5x5x5 sorting

¹²Daniel, Hirshleifer, and Subrahmanyam (1999) also note that their overconfidence theory predicts that BM ratios should better forecast risk-adjusted returns among firms that are hard to value (e.g., R&D-intensive firms comprised largely of intangible assets); A low book/market ratio is itself an indicator of high intangible assets, but can also be low for other reasons such as a risk premium or market misvaluation. Thus, conditioning on other intangible measures provides a test of how intangible measures affect the misvaluation-induced relation between fundamental-price ratios and future returns. So far as we know, this is an as-yet untested implication.

¹³The common stocks in our portfolios must also meet certain data requirements, discussed in the appendix. Also we exclude negative book-equity firms

of the portfolios gives a total of 125 passive portfolios.¹⁴ To minimize transaction costs, each of these portfolios is value-weighted, and the portfolios are rebalanced once per year.

Table 1, Panel A presents the average returns of 25 size/book-to-market sorted portfolios over the 1963:07-1997:12 period. These 25 portfolios are formed by equally weighting the five corresponding size-sorted portfolios. So, for example, the upper-left return in this panel is the average return of a portfolio which equally weights the low book-to-market/low momentum portfolios constructed for the five size quintiles.

The evidence in this panel indicates that there is an independent book-to-market and momentum effect.¹⁵ Consistent with the predictions of the overconfidence model, the momentum effect is much stronger for low book-to-market stocks.¹⁶ Panel B of Table 1 examines similar strategies that exclude the largest and smallest quintile stocks. The largest stocks are presumably the most efficiently priced, so any behavioral effects might become more apparent when these stocks are excluded from the sample. The rationale for excluding the smallest stocks is that these have the least reliable prices, are the least liquid, and are likely to be the most difficult to short.

Panels C and D of Table 1 show the average difference in the returns of the various BM/momentum portfolios and the return of the CRSP value-weighted index (which tracks the S&P 500 closely). These panels reveal that the high-BM/high-momentum portfolio dramatically outperforms the value-weighted index, and the low-BM/low-momentum portfolio underperforms the value-weighted index by a similar amount.

¹⁴We use NYSE-based breakpoints to determine the size quintile breakpoints, and we value-weight the stocks within each of the 125 passive portfolios to avoid biases in computed returns from rebalancing and to avoid placing too much weight on extremely small stocks.

¹⁵Other studies have investigated the interaction of the momentum effect and other characteristics: Chan, Jegadeesh, and Lakonishok (1996) look at price momentum and earnings momentum; Hong, Lim, and Stein (1999) look at analyst coverage and momentum, and Lee and Swaminathan (1999) investigate momentum and trading volume.

¹⁶These results are also consistent with those previously presented by Asness (1997) in this Journal. However, there are some differences between our table and his Table 4. Specifically, his low-momentum, low-book-to-market (LL) portfolio has a considerably lower return than our corresponding portfolio. The reason for this is that in picking out the set of stocks with the lowest book-to-market and past return, one also selects a set of very low market value securities. Since our portfolios are first sorted on size, our average LL firm is considerably larger, and the return of our LL portfolio is not as low.

Table 1: **The Returns of Size, Book-to-Market and Momentum Sorted Portfolios 1963:07-1997:12**

Panel A: Raw Returns, All Quintiles							
	<i>Low</i>	BM			<i>High</i>	<i>H - L</i>	T-stat
<i>Low</i>	0.454	0.713	1.067	1.166	1.389	0.935	(5.286)
	0.728	0.980	1.137	1.288	1.455	0.727	(4.748)
M-m	0.922	1.058	1.174	1.298	1.369	0.447	(2.730)
	1.043	1.141	1.162	1.364	1.400	0.357	(1.930)
<i>High</i>	1.206	1.418	1.369	1.511	1.494	0.288	(1.449)
<i>H - L</i>	0.752	0.705	0.302	0.345	0.105	HH-LL	
T-stat	(3.838)	(4.027)	(1.866)	(2.180)	(0.587)	1.0398	(5.656)

Panel B: Raw Returns, Quintiles 2-4 Only							
	<i>Low</i>	BM			<i>High</i>	<i>H - L</i>	T-stat
<i>Low</i>	0.550	0.651	1.064	1.159	1.527	0.977	(5.019)
	0.686	0.966	1.159	1.174	1.506	0.820	(4.625)
M-m	0.900	1.025	1.120	1.330	1.419	0.519	(2.738)
	1.003	1.098	1.149	1.398	1.430	0.427	(1.998)
<i>High</i>	1.341	1.503	1.406	1.516	1.606	0.265	(1.113)
<i>H - L</i>	0.792	0.852	0.342	0.357	0.080	HH-LL	
T-stat	(3.329)	(4.255)	(1.959)	(2.130)	(0.424)	1.0566	(5.022)

Panel C: Excess of VW, All Quintiles							
	<i>Low</i>	BM			<i>High</i>		
<i>Low</i>	-0.571	-0.312	0.042	0.141	0.364		
	-0.298	-0.045	0.111	0.262	0.429		
M-m	-0.104	0.033	0.148	0.273	0.343		
	0.018	0.115	0.136	0.339	0.375		
<i>High</i>	0.180	0.392	0.344	0.485	0.469		

Panel D: Excess of VW, Quintiles 2-4 Only							
	<i>Low</i>	BM			<i>High</i>		
<i>Low</i>	-0.476	-0.374	0.039	0.134	0.501		
	-0.340	-0.059	0.134	0.149	0.481		
M-m	-0.126	-0.001	0.094	0.305	0.393		
	-0.022	0.073	0.123	0.372	0.404		
<i>High</i>	0.316	0.477	0.381	0.491	0.581		




Figure 1: Annual Returns of HH-LL Strategy (all size quintiles), 1964-1997

3.1 The time-series behavior of the HH-LL portfolio

This section provides an in depth analysis of the returns of the long-short strategy which buys the High-BM, High-momentum (HH) portfolio, and sell the Low-BM, Low-momentum (LL) portfolio. The mean returns and t-statistics for these HH-LL portfolios are given in the lower right corners of Panels A and B. In Figure 1 we plot the annual returns to the HH-LL strategy.

The performance of this portfolio is striking. It generated profits in 31 out of the 34 years in our sample period and realized an average profit of 1.04%/month (and 12.64%/year). A t-statistic testing whether this value is different from zero is 5.66. Based on the t-statistic, the one-tailed p-value testing whether the average return of this portfolio is greater than zero is $1 - (1.4 \times 10^{-8})$. Moreover, the (monthly) CAPM beta of this portfolio is -0.258 ($\alpha = 1.17\%/month$, with a $t = 6.62$); thus the portfolio even provides insurance against market fluctuations.

3.2 Data Mining

Over the past 20 years, financial economists and investors have spent considerable time searching for investment strategies that yield abnormal returns. With so much searching, it is likely, purely by chance, that someone will uncover what looks to be incredible opportunities. Thus, apparent profit opportunities, uncovered in this way, must be interpreted skeptically, and the statistical tests of significance for the profitability of the strategies must be corrected for the data mining bias.

There are several ways of addressing the data-mining question. Perhaps the most robust is to examine how the strategies perform in a sample that is independent of the sample in which the strategies were originally investigated. This has been done extensively for size, book to market, and momentum strategies, but has not been done for strategies that combine the three characteristics as we do here. Fama and French (1998) show that the book to market effect is strong internationally, and Davis, Fama, and French (1999) show that the book-to-market effect is even stronger in the 1929-63 period than in the post-1963 period originally studied. Rouwenhorst (1998) shows that the momentum effect is present in European stock markets, and Jegadeesh and Titman (1999) show that the effect is equally as strong in the 1990s as in the pre-1990 period originally studied.

We will take a somewhat different approach and address the data-mining question by adjusting the significance levels of the test statistics to account for the amount of searching that was done to uncover the strategy. To do this we apply the *Bonferroni inequality* which states that the the p-value which accounts for the fact that we have searched over N strategies is no greater than $N \times p^*$, where p^* , is the p-value obtained from a t-statistic on the significance of the return without correcting for the searching.

Based on the Bonferroni inequality, and the p-value cited earlier, we can say that even if researchers considered over one million separate strategies, an upper bound on the p-value for the best strategy would be 1.4×10^{-2} , or 1.4%, meaning that, if the mean return of all strategies were truly zero, even after searching over one million of them, the probability of getting a t-statistic this high would be still be very low.

However, the p-value for the t-statistic assumes that returns are normally distributed.

It is possible that this statistic is biased because the return distribution is "fat-tailed." Hence, we also used a bootstrap method to construct a p-value for this test. To implement the bootstrap, for each iteration we redraw 414 monthly returns (the length of our original time series) from the set of monthly HH-LL portfolio returns, with replacement. Assuming that the empirical distribution represents the true distribution of returns, and assuming that returns are serially uncorrelated, the fraction of the iterations that are less than zero gives us likelihood that we would ever see an average return over 414 months which was less than zero. We found that with 100,000 bootstrap iterations, the total return of the bootstrapped portfolio was never less than or equal to zero, and in less than 1% of the iterations was the average returns less than 0.5%/month.

As one additional robustness check, we also perform a block-bootstrap of the returns (see Hall (1985)). While standard bootstrap methods are valid only for independent data, block bootstrap methods allow for serial correlation up to a certain number of lags. We performed block bootstraps with block lengths of 12, 24, 36, and 48 months with nearly identical results.¹⁷

3.3 Could Mutual Funds Have Enhanced their Performance?

The previous subsections suggest that the returns of the HH-LL portfolio is negatively correlated with the CRSP value-weighted index, which is inconsistent with the hypothesis that the returns represent compensation for covariation with the market portfolio. However, it is possible that institutions hold portfolios that are exposed to an element of risk that is not captured by the CRSP value-weighted index, and that the HH-LL portfolio exacerbates the institutions' exposure to that risk.

To explore this possibility we examine the returns of mutual funds, which we think are representative of other institutions. Our analysis demonstrates that the return standard deviation of these portfolios would not have increased appreciably if they had tilted

¹⁷We again found that in none of the 100,000 iterations was the total return of the bootstrapped portfolio was less than or equal to zero, and in less than 1% of the iterations was the average returns less than 0.5%/month.

their portfolios towards the HH-LL portfolio. Indeed, a tilt towards the HH-LL portfolio would have substantially increased the Sharpe-ratios of the mutual funds.¹⁸ Moreover, the addition of the HH-LL portfolio would have allowed the mutual funds to increase the expected returns of their portfolios without increasing their tracking error relative to the value-weighted index.

We first examine the marginal effect of adding a small amount of the HH-LL portfolio to two composite mutual fund portfolios which include all growth and all income mutual funds that are available on the CRSP mutual fund data base compiled by Carhart (1997). By regressing the return of our HH-LL portfolio on the mutual fund returns net of the risk-free rate, we are able to discern whether the mutual funds could have achieved higher Sharpe ratios by tilting their portfolios towards our strategy. If we find that the alphas from this regression are positive, then the funds could have increased their Sharpe ratios by tilting towards the HH-LL portfolio.¹⁹

We also consider the fact that maximizing Sharpe ratios may not have been the primary objective of these funds. Rather, the relevant risk of the funds may instead be tracking error, *i.e.*, the risk of underperforming the benchmark. To address this issue, we run a similar regression, except now, the independent variable is the mutual fund return minus the CRSP value-weighted index return.²⁰

The results of these tests are reported in Table 2. Panel A presents a standard CAPM-like regression. The R^2 s indicate that both sets of mutual-fund returns are well diversified, and the intercepts (α s) indicate that, on average, their performance is close to what is predicted by the CAPM. The results in Panel B indicate that the Sharpe ratio of the average mutual fund could have been greatly enhanced by adding the HH-LL portfolio.

¹⁸These tests are equivalent to tests of whether the average growth or income mutual fund is mean-variance efficient. The reported improvements in their Sharpe ratios ignore the possibility that mutual funds may not wish to hold mean-variance efficient portfolios.

¹⁹The alpha will be zero for all portfolios if the mutual-fund is a mean-variance-efficient portfolio. If the mutual-fund is not mean-variance-efficient, then its Sharpe ratio can be increased through tilting towards positive alpha assets and away from negative alpha assets.

²⁰We are assuming that these fund managers believe that they will maximize the ratio of expected return to standard deviation or tracking error by holding growth or income portfolio. These tests would need to be altered if one wanted to examine funds that wish to maximize the Sharpe-ratio conditional on holding stocks in a particular asset class.

Table 2: **Aggregate Mutual Fund Performance**

This table presents the results of a number of regressions of aggregate mutual fund and zero-cost arbitrage portfolio returns. Each row gives the results of a different regression. For each regression, the column headed LHS gives the dependent (Left-Hand Side) variable, and the column headed RHS gives the independent (Right-Hand Side) variable. R_{IN} denotes the aggregate income-fund return, R_{GR} the aggregate growth-fund return, R_{VW} the return on the CRSP value-weighted index of all US common stocks, and R_f the risk-free return over the month (from Ibbotson). α is the regression intercept term, β the regression coefficient, and R^2 the regression R^2 (in percent). All return series are monthly in percent per month, over the period 1963:07-1997:12. t-statistics are in parentheses.

LHS	RHS	α	β	R^2
Panel A: CAPM-like Regressions				
$R_{IN} - R_f$	$R_{VW} - R_f$	0.00 (0.17)	0.8194 (139.73)	97.93
$R_{GR} - R_f$	$R_{VW} - R_f$	0.01 (0.245)	1.0514 (81.71)	94.19
Panel B: Sharpe Ratio Maximization				
HH-LL	$R_{IN} - R_f$	1.15 (6.45)	-0.274 (-5.48)	6.79
HH-LL	$R_{GR} - R_f$	1.18 (6.73)	-0.258 (-6.91)	10.39
Panel C: Tracking Error Maximization				
HH-LL	$R_{IN} - R_{VW}$	1.17 (6.84)	1.517 (8.30)	14.34
HH-LL	$R_{GR} - R_{VW}$	1.06 (5.89)	-0.631 (-4.01)	3.75

Moreover, Panel C of Table 2 indicates that the average fund could have enhanced its return without increasing its tracking error. The t-statistics on all of the intercepts in Panels B and C are highly significant.

The regression results presented in Table 2 show that by adding a small amount of the HH-LL portfolio, mutual funds could have increased their Sharpe ratios. In Table 3, we show exactly how much the Sharpe-ratios of the average growth mutual fund would have improved had they tilted their portfolios toward the HH-LL portfolio by various amounts. This table shows that, with sufficient tilting toward the HH-LL portfolio, the *ex-post* Sharpe ratio would have tripled. In addition, the ratio of the return premium relative to

Table 3: **Enhanced Growth Mutual Fund Performance**

We calculate the return for a portfolio which consists of a \$1 investment in the average growth mutual fund from the CRSP mutual fund database, a \$x investment in the HH portfolio, where x is the number given in column 1 of the Table, and an offsetting \$x short position in the LL portfolio.

This table presents a number of risk and return measures for this portfolio: the average portfolio return (\bar{r}); the annualized -return standard deviation ($\sigma(r-r_f)$); the tracking error or standard deviation of the portfolio return net of the VW market return ($\sigma(r-r_m)$); the Sharpe-ratio ($\overline{r-r_f}/\sigma(r-r_f)$); and the ratio of return in excess of the market to the tracking error ($\overline{r-r_m}/\sigma(r-r_m)$), all in annualized form. All results are for the period 1963:07-1997:12.

frac inv	\bar{r} (%)	$\sigma(r-r_f)$ (%)	$\sigma(r-r_m)$ (%)	Sharpe ratio	$\overline{r-r_m}/$ $\sigma(r-r_m)$
0.0	12.8	16.2	4.0	0.404	0.120
0.1	14.0	15.8	3.9	0.493	0.438
0.2	15.3	15.6	4.3	0.581	0.690
0.3	16.5	15.4	5.0	0.669	0.845
0.4	17.8	15.3	5.9	0.753	0.928
0.5	19.0	15.4	6.9	0.831	0.971
0.6	20.3	15.5	8.0	0.903	0.993
0.7	21.5	15.8	9.2	0.967	1.004
0.8	22.8	16.2	10.4	1.022	1.010
0.9	24.0	16.6	11.6	1.069	1.012
1.0	25.3	17.2	12.8	1.108	1.012
1.1	26.5	17.8	14.0	1.139	1.012
1.2	27.8	18.5	15.3	1.164	1.011
1.3	29.0	19.2	16.5	1.184	1.009
1.4	30.3	20.1	17.8	1.198	1.008
1.5	31.5	20.9	19.1	1.208	1.006
1.6	32.7	21.8	20.3	1.216	1.005
1.7	34.0	22.8	21.6	1.220	1.003
1.8	35.2	23.7	22.9	1.223	1.002
1.9	36.5	24.7	24.2	1.224	1.001
2.0	37.7	25.8	25.4	1.223	0.999
2.1	39.0	26.8	26.7	1.222	0.998
2.2	40.2	27.9	28.0	1.220	0.997

the market tracking error would have increased about eight times.²¹ For comparison, the Sharpe-ratio of the VW index over this period was 0.406. Of course, achieving the largest Sharpe-ratio in the table would probably have required short positions in a number of the securities in the LL portfolio, and hence was probably not feasible for most mutual funds. However, Table 3 shows that even relatively minor tilting would have produced impressive Sharpe-ratio increases.

4 Is The Stock Market Adaptive-Efficient?

The results in the previous section indicate that with the benefit of perfect hindsight a mutual fund investor could have earned significantly higher returns with little apparent additional risk by employing the HH-LL strategy. Based on this we have asserted that the evidence is inconsistent with our traditional view of efficient markets. Of course in reality, investors did not have the benefit of perfect hindsight, so we are somewhat premature in concluding that rational investors should have been able to earn abnormal returns.

In this section we ask whether a rational investor, without the benefit of perfect hindsight, would have been able to profit from the behavioral biases of his or her less rational counterparts. One could argue that with a sufficient understanding of investor overconfidence and its effect on asset prices, an investor could have determined that investing in a portfolio that resembled our HH-LL portfolio would generate abnormal returns.²² However, as we argued earlier, this is probably not realistic. To have the knowledge to implement such a strategy would have required information about the amount of money being managed by overconfident investors as well as the investment strategies of other rational investors. It is more realistic to assume that to benefit from the overconfidence of other investors, a rational investor would first need to learn about the tastes and ra-

²¹Table 3 reports results only for the average growth portfolio. The Sharpe ratio of the average income portfolio is 0.405; The maximized SR in combination with the HH-LL is 1.178, which occurs with \$1.55 investment in the HH portfolio per dollar invested in the mutual-fund holdings.

²²There should be considerable variation in the ability of investors to carry out this sort of reasoning. We would expect that some investors would be able to learn about investor irrationality and mispricing sooner than others, and would be able to earn rents as a result.

tionality of other investors by observing the historical pattern of stock returns.

To address these issues we consider the investment behavior of a hypothetical investor with the sophistication and data that would be required to examine the historical pattern of stock returns. We assume that the investor starts his trades in 1974. We use 1974 for the starting date of this analysis partly for COMPUSTAT data availability reasons. However, we also picked this starting date because an investor could actually have undertaken the type of analysis we perform at this time, while 10 years earlier such an analysis would have been extremely costly.²³

4.1 A Rational Response to an Irrational World

We assume that our hypothetical investor would have tilted his portfolio towards strategies that performed well over the previous 10 years. By doing so, he or she would benefit from the strategy if the excess returns persisted. If the investor had aggressively tilted toward such a strategy, then his or her profits would be quite large if the profits indeed persisted. However, if the investor investor, without the benefits of hindsight, had tilted very cautiously towards the strategy, then his realized profits would be fairly modest even if the strategy earned very large abnormal returns.

Unfortunately, we do not have a good metric for determining how much a rational

²³Prior to the mid-1960s, no one had the powerful computers, asset return data, or the analytical finance tools necessary to evaluate the risk and returns of diversified characteristic-sorted portfolios like those we analyze here. It was only in 1952 that Markowitz (1952) gave precise definitions to return and risk, and spelled out how to think about the tradeoff between them. Only in the 1960's did the CRSP stock data become widely available. Merton (1987) makes a similar point, noting that:

To avoid inadvertent positing of a 'Connecticut Yankee in King Arthur's Court,' empirical studies ... should take care to account for the evolution of institutions and information technologies during the sample period. It is, for example, common in tests of the weak-form of the Efficient Markets Hypothesis to assume that real-world investors at the time of their portfolio decisions had access to the complete prior history of all stock returns. When, however, investors' decisions were made, the price data may not have been in reasonably-accessible form and the computational technology necessary to analyze all these data may not even have been invented. In such cases, the classification of all prior price data as part of the publicly available information set may introduce an important bias against the null hypothesis.

investor, in the presence of a large number of irrational investors, should tilt his portfolio towards strategies that performed well in the past. If the returns that the investor had observed were drawn from a stationary distribution, it would be relatively straightforward to calculate optimal portfolio weights.²⁴ However, stationarity is not likely to be a good assumption in a market that consists of both irrational investors and rational investors that learn from past price movements. At the same time our hypothetical investor is discovering these return patterns, other rational investors are likely to be doing the same thing. If all of the rational investors attempt to exploit investment opportunities observed from historical data, they may push prices to a level where the investment opportunities disappear.

The above discussion implies that determining how much one should tilt towards the higher performing strategies is not at all trivial. An investor would want to have some kind of theory of inefficient markets to help him understand the extent to which the irrational behavior causing the pricing errors is likely to persist. He would also need to have some idea of the extent to which other rational investors uncover these same pricing anomalies and then alter their portfolios to exploit them.²⁵ An investor who believed that he was the only one doing this sort of analysis would strongly tilt his portfolio towards the strategies that performed well in the past. However, an investor who believed that inefficiencies are almost immediately corrected by other active investors might choose not to tilt at all towards the better performing strategies.²⁶ Ironically, if most rational investors believe that the market is very efficient, they will not exploit the strategies and the pricing anomaly is likely to persist. Alternatively, if the rational investors underestimate the

²⁴See, for example, Kandel and Stambaugh (1996) and Barberis (1999).

²⁵We are, of course, assuming here that investors can not perform firm by firm fundamental analysis to determine whether other rational investors have pushed prices to the level where the profit opportunities to have been eliminated. We are implicitly assuming that some of the overconfident investors have special information about fundamentals that are not available to our rational investor who is strictly a quantitative analyst.

²⁶For example, in a recent article in the *Journal of Portfolio Management*, Merton Miller (1999) states: “If you find some mechanical rule that seems to earn above-normal returns – and with thousands of researchers spinning through the mountains of tapes of past data, anomalies, like the currently fashionable ‘momentum effects,’ are bound to keep turning up – and imitators will enter and compete away those above-normal returns exactly as in any other setting in economics. Above-normal profits, wherever they are found, inevitably carry with them the seeds of their own decay.” (p. 99)

number and aggressiveness of other rational investors, they may as a group tilt too much towards these strategies causing the anomaly to be reversed.

4.2 The Profits From an Adaptive Strategy

While it is difficult to determine the optimal portfolio composition in this setting, we can examine the success of a purely empirically-based portfolio strategy that, in each year between 1974 and 1997, mechanically exploits the return anomalies observed in the prior ten years. As our starting point we assume that the investor observes the returns of the 125 characteristic-sorted portfolios described in Section 3 and forms a zero cost portfolio based on the characteristic-sorted returns observed over the prior ten years. Specifically, in each month, the portfolio consists of a combination of the 125 characteristic-sorted portfolios with weights that are proportional to their returns over the prior ten years. Implicitly, we are assuming that the past returns provide information about future returns and that we have no information about the covariance structure of returns.²⁷

If the stock market is efficient in the traditional sense, then this strategy should not be profitable. However, as we just mentioned, inferring behavioral biases from past returns and determining the extent to which they are likely to persist is not a trivial task, so we do not necessarily expect to observe zero initial profits from a strategy of this type. However, in a market that we call ‘adaptive-efficient,’ rational investors can learn and trade based on information inferred from the past price history. If there exist a sufficient number of these rational investors, the excess returns from mechanical strategies of this type should quickly dissipate.

The results in Table 4 appear to be inconsistent with this notion of adaptive efficiency. Specifically, the profits from the adaptive strategy are quite large and show no sign of dissipating over time. Indeed, the performance continues into the 1990’s, subsequent to the publication of academic papers on both the momentum and book-to-market effects.

²⁷Lehmann (1990) and Lo and MacKinlay (1990) previously examined this type of strategy. Again, we are not saying that a rational investor would implement such a strategy, only that a relatively agnostic investor would tend to tilt his or her portfolio towards this strategy.

Table 4: **Annual Returns to Empirically-Based Strategies**

*Returns in Percent/Year, Annualized Sharpe-ratio (SR)
CAPM α (in %) and β for monthly regressions*

year	All Size Quintiles			Quintiles 2-4		
	L-S	L-VW	S-VW	L-S	L-VW	S-VW
1974	9.58	6.53	-3.06	12.39	7.49	-4.90
1975	-2.44	14.26	16.70	-3.97	11.78	15.75
1976	16.52	20.86	4.35	16.67	20.61	3.94
1977	9.26	18.55	9.29	4.68	15.26	10.58
1978	0.33	5.87	5.54	-0.69	3.66	4.36
1979	-6.61	7.86	14.47	-12.60	7.21	19.82
1980	-2.50	-1.39	1.11	-6.36	-3.98	2.38
1981	5.22	10.01	4.79	4.69	10.01	5.33
1982	23.07	15.57	-7.50	21.31	15.14	-6.17
1983	16.13	14.57	-1.57	14.05	13.61	-0.45
1984	4.32	-3.56	-7.88	11.55	0.03	-11.52
1985	4.08	-0.24	-4.33	7.48	3.59	-3.90
1986	4.43	-2.74	-7.17	11.42	0.20	-11.23
1987	4.20	-6.44	-10.64	0.81	-6.61	-7.42
1988	17.11	7.51	-9.60	15.50	8.85	-6.65
1989	10.11	-5.22	-15.33	4.77	-4.90	-9.67
1990	12.82	-8.38	-21.20	2.03	-9.21	-11.24
1991	-2.99	6.53	9.52	0.52	9.09	8.57
1992	4.20	8.44	4.24	11.12	10.77	-0.35
1993	7.34	7.26	-0.08	13.96	9.30	-4.65
1994	11.42	-0.81	-12.23	2.93	-1.67	-4.61
1995	7.23	-0.25	-7.48	5.64	-0.62	-6.26
1996	9.00	-0.18	-9.18	6.05	-1.39	-7.44
1997	20.83	-0.95	-21.77	15.23	-3.72	-18.95
mean	7.61	4.74	-2.88	6.63	4.77	-1.86
t-stat	(4.57)	(3.09)	(-1.40)	(4.29)	(3.29)	(-0.81)
SR	0.932	0.630	-0.288	0.876	0.672	-0.166
β	-0.099	0.020	0.120	-0.155	0.021	0.176
t(β)	(-3.77)	(0.82)	(3.67)	(-6.98)	(0.90)	(5.93)
α	0.619	0.332	-0.287	0.564	0.335	-0.229
t(α)	(5.16)	(2.94)	(-1.94)	(5.58)	(3.13)	(-1.69)

The first three columns of Table 4 report the returns of this strategy implemented over the 1974 to 1997 period. We report the returns both to the overall long-short strategy (“L-S”), as well as the returns of the long side of the strategy net of the return on the CRSP value-weighted index (“L-VW”), and the return to the short side of the strategy, net of the CRSP VW return (“S-VW”). The last three columns report the returns of a similar strategy that excludes stocks with capitalizations that place them in the smallest and largest size quintiles.

Our evidence indicates that an investor would have benefited from tilting his or her portfolio towards either portfolio. Both are profitable in almost every year, have high Sharpe ratios (0.932 and 0.876), and, like the HH-LL strategy, have negative CAPM β s. In unreported regressions we show that if mutual funds had tilted their portfolios towards this strategy, they could have increased their average rates of return without increasing their variance or their tracking error.

5 Conclusion

The traditional notion of efficient markets described in the academic finance literature is quite strong and probably unrealistic. Essentially, the theory implies that an investor in the 1960s and 1970s, with access to the sophisticated computers and data bases that we now have, would not have been able to use his better technology to earn above average returns. The theory states that market prices are set ‘as if’ investors have access to, and can costlessly process, all available information.

The evidence in this paper rejects this notion of efficient markets in favor of an alternative theory which suggests that asset prices are influenced by investor overconfidence. Portfolio strategies that might be suggested by the overconfidence theory realize extremely high and persistent abnormal returns. The paper also discusses a somewhat weaker, and somewhat murkier definition of efficiency that we call adaptive-efficiency. The definition recognizes the behavioral biases of most market participants, but also assumes that there exist other investors who can detect and profit from these behavioral biases by examining

past price trends. Because these investors are likely to be risk averse and have limited capital, they are not likely to immediately eliminate the pricing biases generated by their less rational counterparts.²⁸ However, if a significant number of investors are rational in this sense, then the profits from their strategies should quickly dissipate.

To assess the extent to which the market is adaptive-efficient we examine a strategy that mechanically purchases (or shorts) stocks with characteristics that did well (or poorly) over the previous 10 years. We examine this strategy in the post 1974 period, when the technology and data required to implement the strategy was clearly available. If rational investors can learn from past price patterns, then the profits from such a strategy should be relatively modest and should diminish over time. Our evidence indicates that this is not the case. The strategy generates substantial profits over a long time period and continues to generate substantial profits through the 1990s.

There are, of course, a number of caveats that must be considered before one rejects our weaker view of market efficiency. The most important caveat is that our tests were limited to the analysis of only three characteristics, size, book-to-market and momentum. Perhaps, by basing our tests on characteristics that we knew to be associated with returns, we have biased our tests towards rejecting market efficiency. While this is probably true, this critique does not really apply to the evidence in the 1990s, since our knowledge of these trading strategies, if not this particular strategy, came from research on data prior to 1990.

However, if this paper had been written in 1980, our prime focus might have been on the small-firm effect and we might have presented an algorithm that endogenously selected small stocks based on the prior returns of this asset class. Of course, we now know that a small firm strategy has done poorly during the past 15 years. Perhaps this is evidence that, after sufficient time and publicity, investors do learn from past price patterns, and that the anomalies presented in this paper will not continue into the next millennium.

²⁸See Shleifer and Vishny (1997) for a discussion of these sorts of frictions.

References

- Asness, C. S., 1997, “The Interaction of Value and Momentum Strategies,” *Financial Analysts’ Journal*, vol. 53, no. 2, 29–36.
- Barberis, N., 1999, “Investing for the Long Run when Returns are Predictable,” *Journal of Finance*, *forthcoming*.
- Barberis, N., A. Shleifer, and R. Vishny, 1998, “A Model of Investor Sentiment,” *Journal of Financial Economics*, vol. 49, no. 3, 307–343.
- Carhart, M. M., 1997, “On Persistence in Mutual Fund Performance,” *Journal of Finance*, vol. 52.
- Chan, L. K., N. Jegadeesh, and J. Lakonishok, 1996, “Momentum Strategies,” *Journal of Finance*, vol. 51, no. 5, 1681–1714.
- Daniel, K. D., M. Grinblatt, S. Titman, and R. Wermers, 1997, “Measuring Mutual Fund Performance with Characteristic-Based Benchmarks,” *Journal of Finance*, vol. 52, no. 3, 1035–1058.
- Daniel, K. D., D. Hirshleifer, and A. Subrahmanyam, 1998, “Investor Psychology and Security Market Under- and Over-reactions,” *Journal of Finance*, vol. 53, no. 6, 1839–1886.
- , 1999, “Investor Psychology and Capital Asset Pricing,” Northwestern University Working Paper.
- Davis, J., E. F. Fama, and K. R. French, 1999, “Characteristics, Covariances, and Average Returns: 1929-1997,” *Journal of Finance*, *forthcoming*.
- De Long, J. B., A. Shleifer, L. Summers, and R. J. Waldmann, 1991, “The Survival of Noise Traders in Financial Markets,” *Journal of Business*, vol. 64, no. 1, 1–20.
- DeBondt, W. F. M., and R. H. Thaler, 1995, “Financial Decision-Making in Markets and Firms: A Behavioral Perspective,” in *Finance, Handbooks in Operations Research and Management Science*, ed. by R. A. Jarrow, V. Maksimovic, and W. T. Ziemba. North Holland, Amsterdam, vol. 9, chap. 13, pp. 385–410.
- Einhorn, H. J., 1980, “Overconfidence in Judgment,” *New Directions for Methodology of Social and Behavioral Science*, vol. 4, no. 1, 1–16.
- Fama, E. F., 1970, “Efficient Capital Markets: A Review of Theory and Empirical Work,” *Journal of Finance*, vol. 25, no. 2, 383–417.
- Fama, E. F., and K. R. French, 1993, “Common risk factors in the returns on stocks and bonds,” *Journal of Financial Economics*, vol. 33, 3–56.
- , 1998, “Value Versus Growth: The International Evidence,” *Journal of Finance*, vol. 53, no. 6, 1975–1999.

- Gervais, S., and T. Odean, 1998, "Learning to be Overconfident," Working paper, Wharton School.
- Griffin, D., and A. Tversky, 1992, "The Weighing of Evidence and the Determinants of Overconfidence," *Cognitive Psychology*, vol. 24, 411–435.
- Grossman, S. J., and J. E. Stiglitz, 1980, "On the impossibility of informationally efficient markets," *American Economic Review*, vol. 70, no. 3, 393–408.
- Hall, P., 1985, "Resampling a Coverage Pattern," *Stochastic Processes and their Applications*, vol. 20, 231–246.
- Hansen, L. P., and R. Jagannathan, 1991, "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, vol. 99, no. 2, 225–262.
- Hirshleifer, D., 1999, "An Adaptive Theory of Self-Deception," Ohio State University Working Paper.
- Hong, H., T. Lim, and J. Stein, 1999, "Bad News Travels Slowly: Size, Analyst Coverage and the Profitability of Momentum Strategies," *Journal of Finance*, forthcoming.
- Hong, H., and J. Stein, 1999, "A Unified Theory of Underreaction, Momentum Trading and Overreaction in Asset Markets," *Journal of Finance*, forthcoming.
- Jegadeesh, N., and S. Titman, 1999, "Profitability of Momentum Strategies: An Evaluation of Alternative Explanations," University of Texas, Austin working paper.
- Kandel, S., and R. F. Stambaugh, 1996, "On the Predictability of Stock Returns: An Asset-Allocation Perspective," *Journal of Finance*, vol. 53, no. 2, 385–424.
- Lee, C. M., and B. Swaminathan, 1999, "Price Momentum and Trading Volume," Cornell University Working Paper.
- Lehmann, B. N., 1990, "Fads, Martingales, and Market Efficiency," *Quarterly Journal of Economics*, vol. 105, no. 1, 1–28.
- Lo, A. W., and A. C. MacKinlay, 1990, "When Are Contrarian Profits Due to Stock Market Overreaction?," *Review of Financial Studies*, vol. 3, no. 2, 175–205.
- Lord, C. G., L. Ross, and M. R. Lepper, 1979, "Biased assimilation and attitude polarization: the effects of prior theories on subsequently considered evidence," *Journal of Personality and Social Psychology*, vol. 37, no. 11, 2098–2109.
- Markowitz, H. M., 1952, "Portfolio Selection," *Journal of Finance*, vol. 7, no. 1, 77–91.
- Merton, R. C., 1987, "A Simple Model of Capital Market Equilibrium with Incomplete Information," *Journal of Finance*, vol. 42, no. 3, 483–510.
- Miller, M. H., 1999, "A History of Finance," *Journal of Portfolio Management*, vol. 25, no. 4, 95–101.

- Odean, T., 1998a, "Are Investors Reluctant to Realize Their Losses?," *Journal of Finance*, vol. 53, no. 5, 1775–1798.
- , 1998b, "Volume, Volatility, Price and Profit When All Traders are Above Average," *Journal of Finance*, vol. 53, no. 6, 1887–1934.
- Roberts, H., 1967, "Statistical versus Clinical Predictions of the Stock Market," unpublished manuscript, University of Chicago.
- Rouwenhorst, K. G., 1998, "International Momentum Strategies," *Journal of Finance*, vol. 53, no. 1, 267–284.
- Shleifer, A., and R. Vishny, 1997, "The Limits to Arbitrage," *Journal of Finance*, vol. 52, no. 1, 35–55.
- Simon, H. A., 1982, *Models of Bounded Rationality*. MIT Press, Cambridge, Mass.
- Trivers, R., 1985, *Social Evolution*. Benjamin/Cummings, Menlo Park.
- , 1991, "Deceit and Self-Deception," in *Man and Beast Revisited*, ed. by R. Robinson, and L. Tiger. Smithsonian Press, Washington D.C.
- Waldman, M., 1994, "Systematic Errors and the Theory of Natural Selection," *American Economic Review*, vol. 84, no. 3, 482–497.