May 18, 2004

Self-Deception and Deception in Capital Markets

Josh Coval*
David Hirshleifer**
Siew Hong Teoh***

This paper was written for presentation at the December 2001 Austrian Ministry of Science Workshop on Exchange, Deceptions and Self-deceptions in Paris, France. We thank Caroline Gerschlager and anonymous reviewers for very helpful comments.

*Harvard Business School, Soldiers Field, Boston, MA 02163; (617) 495-5056; *jcoval@hbs.edu*; http://dor.hbs.edu/fi_redirect.jhtml?facInfo=bio&facEmId=jcoval

,*Fisher College of Business, The Ohio State University, 2100 Neil Avenue, Columbus, OH 43210-1144; Hirshleifer: http://www.cob.ohio-state.edu/fin/faculty/hirshleifer/, (614) 292-5174; *hirshleifer_2@cob.osu.edu*; Teoh: (614) 292-6547; *teoh_2@cob.osu.edu*

Self-Deception and Deception in Capital Markets

Abstract

We argue that self-deception underlies various aspects of the behavior of investors and of prices in capital markets. We examine the implications of self-deception for investor overconfidence, and how firms and financial institutions can exploit the overconfidence of investors in a predatory fashion. These ideas link self-deception to deception by others. We also examine how investor self-deception and overconfidence can affect financial reporting and disclosure policy.

1 Introduction

In recent years a body of evidence on security returns has presented a sharp challenge to the traditional view, known as the efficient markets hypothesis, that securities are rationally priced to reflect all publicly available information. Some of the more pervasive anomalies can be classified as follows:

1. Event-based return predictability (average stock returns on the days of publicly-announced events are of the same sign as average subsequent long-run abnormal performance)

2. Short-term momentum (stocks that have done well (poorly) over the past 6-18 months continue to do well/poorly over the next 6-18 months)

3. Long-term reversal (stocks that have done well (poorly) in the past 6-18 months, "overreact" and perform poorly (well) over the subsequent 3-5 years)

4. High volatility of asset prices relative to fundamentals.

5. Abnormal returns to stocks with a book value of assets that is large relative to the market value of their assets.

6. Short-run post-earnings announcement stock price 'drift' in the direction indicated by the earnings surprise, but abnormal stock price performance in the opposite direction of long-term earnings changes.

There remains disagreement over the interpretation of the above evidence of predictability. One possibility is that these anomalies are chance deviations to be expected under market efficiency (Fama (1998)). We believe the evidence does not accord with this viewpoint, because some of the return patterns are strong and regular. The book-to-market and momentum effects are present both internationally and in different time-periods. Also, the pattern mentioned in (1) above obtains for the great majority of event studies.

Alternatively, these patterns could represent the effects of variations in rational risk premia. That is, the abnormal returns may, in fact, be appropriate average compensation for the risks their investors are expected to bear. However, these patterns seem to imply that very high average stock returns can be achieved with simple trading strategies. Thus, for these patterns to be consistent with a frictionless, fully rational model of asset prices, individuals would have to perceive these strategies to be extremely risky.

A rational perception of high risk could arise if the payoffs from these strategies tended to be particularly poor exactly when some other important component of investors' wealth (such as the stock market as a whole) was also doing badly. Alternatively, investors could care about more than just the probability distribution of their possible wealth and consumption outcomes; they might care about the process by which uncertainty about these outcomes is resolved.

Either explanation for these patterns amounts to a claim that investors place very widely varying values on an additional dollar of investment payoff in different states of the world (MacKinlay (1995)). Most research indicates that the outcome correlations required to deliver such variation in the value placed upon consumption are not apparent in the data; and that the investor preferences needed to achieve such variation are implausible.[1] Given this evidence, it seems reasonable to consider explanations for the observed return patterns based on imperfect rationality.

Moreover, there are important corporate financing and payout patterns that seem potentially related to misvaluation in the stock market. Firms tend to issue equity (rather than debt) after rises in market value, and when the firm or industry book/market ratio is low. In fact, firms seem to be able to engage in market-timing by issuing stock before the stock market as a whole is going to fall, and issuing debt instead prior to rises in the market (see Baker and Wurgler (2002)). There are industry-specific financing and repurchase booms, perhaps designed to exploit industry-level mispricing. Transactions such as takeovers that often rely on securities financing are also prone to industry booms and quiet periods, and there is evidence that market misvaluation influences the volume and character of takeover activity (Dong, Hirshleifer, Richardson, and Teoh (2003)), and the extent of investment activity more generally (Polk and Sapienza (2003)).

Proponents of rational markets contend that the above anomalies either are linked to risks that are poorly captured in current asset pricing models or are simply the result of data-mining. To support their view, they point to the enormous body of evidence that the average active mutual fund manager cannot outperform the index.[2] Moreover,

[1] See, e.g., Chen (2003). Campbell and Cochrane (1999) find that a utility function with very strong 'habit persistence' is required to explain the predictable variation in market returns. To be consistent with cross-sectional predictability findings (on size, book-to-market, and momentum, for example), a model would presumably require even more extreme variation across states of the world in the marginal value placed upon consumption (i.e. in marginal utilities). Also, the model would require that marginal utilities covary strongly with the returns on portfolios formed based upon firm size, book-to-market ratios and past return momentum. No such correlation is obvious in examining the data.

[2] The finding by Jensen (1968), (1969) that actively managed mutual funds do not outperform the market has been confirmed in dozens of subsequent empirical studies of mutual funds; see also Malkiel

while some mutual funds do outperform the market, little if any evidence indicates this occurs any more frequently would be expected by pure chance. Proponents of rational markets contend that if the above anomalies are pervasive, economically significant, and unrelated to risk, surely we should see evidence that professional money managers take advantage of them. The fact that we see scant evidence of outperformance indicates that they do not. As Rubinstein (2001) puts it, the evidence that the average mutual fund cannot beat the market is "a nuclear bomb against [the behavioralists'] puny sticks." On the other hand, Berk and Green (2003) (no behavioralists themselves) counter that even if managers have the ability to take advantage of such anomalies we should not necessarily expect to see this ability show up strongly in their performance. If new capital flows to managers who have strong track records and, due to declining economies of scale, restricts their ability to outperform going forward, little persistence will appear in manager return data.

In our view, the failure of the average mutual fund to beat the market is interesting but not particularly supportive of market efficiency. The efficient markets hypothesis does not say that the average investor cannot beat the market; it posits that *no* investor can beat the market (unless he or she possesses private information unavailable to other investors. And in fact, there is evidence that some investors do persistently beat the market, and that this abnormal performance is unlikely to derive from inside information (for perspectives on the persistence of mutual fund abnormal performance, see, e.g., Grinblatt and Titman (1992), Carhart (1997), and Daniel, Grinblatt, Titman, and Wermers (1997); on the persistence of individual investor abnormal trading performance, see, e.g., Coval, Hirshleifer, and Shumway (2003)). From a psychological perspective suggests that the funds that attract investors will be those that appeal to investors' emotions and beliefs, however biased. So there is no presumption that the average fund will do well. For example, if at some point investors are irrationally exuberant about the biotech sector, cash will flow to funds heavy in biotech firms. More rational portfolios that place lower weight upon biotech will on average earn high subsequent returns. However, such funds may not attract many investors— for the same reason that the mispricing is present in the first place.

The fact that vast amounts of wealth are invested in actively managed mutual funds that are unable to achieve superior performance does not support the view that investors are good at choosing funds. Nor does it support the view that funds make good choices

(1995).

on behalf of investors.[3] There is some dissonance between the views that investors trade foolishly in a fashion that creates market inefficiencies, and yet are smart enough to select those mutual funds that follow strategies designed to exploit these same inefficiencies.

Of course, one should not swing to the opposite extreme of dismissing the efficient markets hypothesis as a mere misconception. All serious students of financial markets agree that there are forces acting to make prices reflect information accurately. The question is what *degree* of efficiency is achieved. For example, investors' observation of historical performance should push them toward better mutual funds. This is just an instance of the general argument that when there is a profit opportunity, smart economic agents ought to exploit it. The general obstacle is that investors may be biased in their assessments. Such bias, in the context of mutual funds, can be hard to eliminate because of inattention, noise, sample size, post-selection/reporting biases, and fund manager turnover. Nevertheless, any intelligent critical judgment that investors bring to bear in evaluating investments will tend to make the market more efficient.

In any case, it is natural to consider whether failures of rationality explain the puzzling empirical patterns in securities markets. Richard Thaler and colleagues starting in the 1980s led the application of psychological evidence to problems in economics. Psychological ideas relating to overreaction stimulated the collection of a great deal of empirical evidence on stock price serial correlation and volatility. The publication of Fischer Black's (1986) presidential address to the American Financial Association on "Noise," and the 1987 stock market crash further increased academic interest in imperfectly rational trading. This set the stage for analytical modeling of the pricing effects of imperfectly rational trading.

Some pioneering models captured imperfect rationality in asset markets by including mechanistic traders who either make pure noise trades,[4] or positive feedback trades in which new purchases are an increasing function of past price moves.[5] This was an

[3]Rubinstein (2001) offers overconfidence as an explanation for the empirical findings; it is somewhat surprising to see overconfidence playing a pivotal role in a brief *in support* of efficient markets. He maintains that overconfidence causes managers and investors to work too hard to eliminate profit opportunities, making the market in a sense too efficient.

It is plausible that overconfidence will cause individuals to generate more information. But Rubinstein's reasoning does not address the possibility that overconfident investors and portfolio managers may take actions that *generate* rather than correct mispricing. This is, however, a natural implication of overconfidence, as developed in several models (see *e.g.*, Odean (1998) and Daniel, Hirshleifer and Subrahmanyam (1998, 2001)).

[4]That is, they place trades that are purely random in size and direction; see Kyle (1985).

[5]See Cutler, Poterba, and Summers (1990), DeLong et al (1990a, 1990b), and Frankel and Froot (1986, 1990).

efficient way to illustrate some crucial insights about survival, arbitrage, and pricing. However, in full generality, the mechanistic modeling approach is very elastic. If noise trades can be arbitrarily correlated with other economic variables, *any* return pattern can be explained. It is not enough for behavioral finance to offer as an alternative to rational finance the model "everything else". The economic content of mechanistic trader models comes from the choice of assumptions on trades to reflect facts about psychology or trading. In the hope of being more accurately predictive, recent research has explicitly modeled how decision-making occurs in a way that reflects psychological biases.

This paper discusses the potential role of a particular psychological bias – the tendency of individuals to deceive themselves – in explaining the above securities market phenomena. We speculate that self-deception is consistent with a variety of patterns in the experimental psychology literature (see, e.g., Trivers (1985) and Arkes and Hirshleifer (2004)). Individuals consistently convince themselves that their ability or judgment is above the population average. They talk themselves out of decision options that might call their past judgments into question. At times, they choose to disregard new information that is inconsistent with their priors. At other times, they are excessively quick to cast aside old evidence in favor of new.

The following section discusses the evidence on self-deception in the psychology literature. Section 3 discusses a variety of models of securities market behavior that are based on self-deception. Section 5 assesses the promise of self-deception in contributing to our understanding of securities market phenomena.

2 Self-Deception in Experimental Psychology

Psychological evidence suggests that people are remarkably skillful at deceiving themselves. In our view, a wide variety of evidence from experimental psychology can be interpreted in terms of self-deception. Most people believe that their ability or judgment is better than it is (overconfidence), actively avoid decisions that risk calling their judgment into question (regret avoidance), and learn in ways that are inconsistent with Bayes' theorem (conservativeness and representativeness). Before examining these phenomena in greater detail, an important question to address is why self-deception exists in the first place.

Self-deception is a trait that some evolutionary theorists believe was actively selected for when humans were evolving. The argument is that sometimes it is easier for an indi-

vidual to deceive others if he genuinely is fooled himself. Trivers (1985, 1991) discusses evidence that people cannot perfectly control indicators of their true internal states. This creates selection for the ability to read subtle cues to infer the mental states of other individuals. In Trivers' self-deception theory, individuals are designed to think they are better (smarter, stronger, better friends) than they really are. Truly believing this helps them fool others about these qualities. Thus, an individual may be able to posture more effectively to potential competitors, friends and mates by genuinely believing that his abilities are greater than they are.

An alternative approach, suggested by Bernardo and Welch (2001), is that the presence of overconfident individuals confers an informational benefit upon the group. Overconfident individuals are more willing to make use of their own information rather than blindly following the actions of predecessors in what has been called an informational cascade (see, e.g., Bikhchandani, Hirshleifer, and Welch (1998)). More generally, overconfident individuals may be more willing to generate new information or undertake untested activities that the group can learn about. These considerations suggest that natural selection (via kin selection or group selection) may sometimes favor self-deception and overconfidence.

We now proceed to discuss, in turn, several apparent manifestations of self-deception that may have influence behavior in financial markets.

2.1 Overconfidence

Self-deception can explain overconfidence (a tendency to overestimate ones ability or judgment accuracy), and dynamic processes that support overconfidence such as biased self-attribution (a tendency to attribute successes to one's own ability and failure to bad luck or other factors), confirmatory bias (a tendency to interpret evidence as consistent with one's preexisting beliefs), hindsight bias (a tendency to think you 'knew it all along'), rationalization (straining to come up with arguments in favor of one's past judgments and choices), and action-induced attitude changes of the sort that motivate cognitive dissonance theory (becoming more strongly persuaded of the validity of an action or belief as a direct consequence of adopting that action or belief); see Cooper and Fazio (1984). Casual observation suggests that all these phenomena affect the decisions of both amateur and professional investors.

An extensive literature on calibration shows that people believe their knowledge is

more accurate than it really is.[6] For example, their predictions of probabilities of events are too extreme (too high relative to the true frequency when they think the event probably will occur, too low when they think it will not). The confidence intervals they provide for quantities are too narrow, e.g., 98% confidence intervals contain the true quantity only 60% of the time (Alpert and Raiffa (1982)). Experts are well-calibrated in some contexts but not others (see Camerer (1995) p. 592-3). Experts can be more prone to overconfidence than non-experts when predictability is low and evidence is ambiguous (Griffin and Tversky (1992)). Overconfidence is greater for challenging judgment tasks, and individuals tend to be more overconfident when feedback on their information or decisions is deferred or inconclusive.[7]

Overconfidence is sometimes reversed for very easy items (Lichtenstein and Fischoff (1977)).[8] Overconfidence implies *overoptimism* about the individual's ability to succeed in his endeavors. Such optimism has been found in a number of different settings (Miller and Ross (1975)).

2.2 Dynamic Processes that Support Overconfidence

Self-deception also explains why there are *action-induced attitude changes* related to overconfidence that motivate the theory of *cognitive dissonance*.[9] In one experiment people who chose between two products downgraded their assessments of the one they did not pick. In another, women who had to exert greater effort to gain entry to a group subsequently liked the group more. In other experiments, people who were induced with mild incentives or by request to express opinions became more sympathetic to those opinions. A tendency to be excessively attached to activities for which one has expended resources, the *sunk cost effect*, has been confirmed in several contexts (Arkes and Blumer (1985)). The self-deception theory suggests that a tendency to adjust attitudes to match past actions is a mechanism designed to persuade the individual that he is a skillful

[6]See, e.g., Keren (1991), Lichtenstein, Fischoff, and Phillips (1982), McClelland and Bolger (1994), and Yates (1990).

[7]Einhorn (1980), Fischoff, Slovic, and Lichtenstein (1977), Griffin and Tversky (1992), Lichtenstein, Fischoff, and Phillips (1982), Yates (1990).

[8]This is not surprising on mechanical grounds; in the extreme case of perfect decision accuracy, it is possible to be under- but not over- confident about accuracy. It has been suggested that apparent overconfidence could be an artifact of the choice of questions that are not a "representative sample of the knowledge domain" (e.g., Gigerenzer, Hoffrage, and Kleinbolting (1991)), but overconfidence remains when questions are randomly selected from the knowledge domain, and has been documented in many practical choice settings (Griffin and Tversky (1992), Brenner et al. (1996), Soll (1996)).

[9]Festinger and Carlsmith (1959), Cooper and Fazio (1984) and Harmon-Jones and Mills (1999) ch.1).

decisionmaker (see also Nel, Helmreich, and Aronson (1969) and Steele and Liu (1983)).

People tend to interpret ambiguous evidence in a fashion consistent with their own prior beliefs. They give careful scrutiny to inconsistent facts and explain them as due to luck or faulty data-gathering (see Gilovich (1991) ch.4). This form of *confirmatory bias* can help maintain self-esteem, consistent with self-deception. Exposure to evidence should tend to cause rational Bayesians with differing beliefs to converge, whereas the attitudes of experimental subjects exposed to mixed evidence tend to become more polarized (e.g., Isenberg (1986), Lord, Ross, and Lepper (1979)). Forsythe et al (1992) find that individuals more subject to this confirmation bias lose money in an experimental market to those who are less subject to it. Confirmatory bias may cause some investors to stick to unsuccessful trading strategies, causing mispricing to persist.

Overconfidence should cause people fail more often than they expect to. But this suggests that rational learning over time would tend to eliminate overconfidence. So for self-deception to succeed, nature must provide mechanisms that bias the learning process. This is consistent with self-enhancing *biased self-attribution*. People tend to credit themselves for past success, and blame external factors for failure.[10] As Langer and Roth (1975) put it, 'Heads I win, tails it's chance'; see also the discussion of DeLong, Shleifer, Summers, and Waldmann (1991). Overconfidence and biased self-attribution are static and dynamic counterparts; self-attribution causes individuals to learn to be overconfident rather than converging to an accurate self-assessment. Such a phenomenon may help explain the escalation by individuals of poor investments.[11] For example, in the disposition effect of Shefrin and Statman (1985), individuals are reluctant to sell stocks in which they have lost money relative to their winners; for supporting evidence, see, e.g., Grinblatt and Han (2003).

In a financial market setting, this suggests that the confidence of an investor will tend to grow when public information is in agreement with his past investment choices, but does not fall commensurately when public information contradicts his private information. A possible sequence in the financial market setting is as follows: (1) If an investor is overconfident about his ability to dig up useful information, then he will sacrifice leisure to do more digging. (2) When he digs up a given piece of information, his ego is on the line (i.e., overconfidence-inducing psychological mechanisms are triggered); (3) When an investor is in the habit of digging things out and trading actively on this

[10] See Fischoff (1982), Langer and Roth (1975), Miller and Ross (1975), and Taylor and Brown (1988).

[11] There are alternative explanations for sunk cost and escalation behavior; see, e.g., Staw (1976), Arkes and Blumer (1985), Arkes (2002), and Hirshleifer and Welch (2002).

basis, he becomes more overconfident (biased self-attribution is induced).

2.3 Anticipation and Regret

Anticipation and regret represent additional dimensions along which individuals may be prone to self-deception. According to expected utility theory, instantaneous utility derives solely from current consumption resulting from current and past decisions. However, people also seem to derive instantaneous utility also from anticipating future consumption and from regret or elation if current consumption falls differs from that which was anticipated or experienced in the past, or what could have been obtained from alternative decisions. (e.g., Josephs et al. (1996), Ritov (1996)). An individual enjoys their day at work more on a Friday than on a Monday (habit), on the Friday prior to a vacation than on other Fridays (anticipation), and on the Monday after a one-week vacation than on the Monday after a two-week vacation that was accidentally cut short by a week (regret).

Chacko, Chowdhry, Cohen, and Coval (2002) model anticipation and regret in a time-consistent setting in which individuals rationally choose to deceive themselves. That is, after having deceived themselves and paid the consequences, if individuals are reminded of their self-deception, they "would not have done things differently." Individuals highly averse to regret will rationally convince themselves to be pessimistic. Individuals who derive utility from anticipation will rationally choose to be optimistic. Individuals will also optimally choose to learn in a conservative manner (see the discussion below on conservativeness), so that they smooth the effects of anticipation and regret through time.

An efficient heuristic method of comparing decision alternatives may be to line up and compare possible outcomes by state of the world (rather than evaluating the expected utility of each alternative separately and then comparing). Thus, having feelings be triggered by comparison of outcomes may be an effective mechanism for motivating good choices. Regret avoidance may also reflect a self-deception mechanism designed to protect self-esteem about decision-making ability (Josephs et al (1996)).

Regret theory (Bell (1982); Loomes and Sugden (1982)) was developed to explain patterns of choices that violated expected utility theory. The idea is that the utility of an option depends not only upon the utility of its potential outcomes, but also on the disutility of not obtaining the outcomes associated with other options were they to prove to have been the correct choice. Research has demonstrated that anticipated re-

gret increases regret-minimizing choices, changes attitudes about behavior, and changes future behavior (Josephs et al (1996); Ritov (1996)). Anticipation of regret about bad choices has been used to explain why consumers purchase higher priced but well-known brands (Simonson, 1992) and why they are reluctant to trade equivalent lottery tickets, even with an added cash incentive (Bar-Hillel and Neter (1996)).

Another line of choice research has focused on regret for past actions. The general format of these studies is to describe two people who made choices that resulted in bad outcomes and ask, "Who feels worse (or more regret)?" Kahneman and Tversky (1982) presented the now classic stockholder problem in which two investors each lose $1,200, but through different choices. Mr. Paul 'loses' his $1,200 by choosing to keep his stock in Company A, which remains unchanged in value, instead of availing himself of the opportunity to sell and buy stock in Company B, which increases in value by $1,200. Mr. George loses his $1,200 by choosing to sell his stock in Company B to buy stock in Company A. The majority of subjects state that Mr. George, who choose to act and lost, would feel more regret than Mr. Paul, who choose not to act but still lost. This result, that loss resulting from action is regretted more than loss resulting from inaction, has been called the 'omission bias' (Spranca, Minsk, and Baron (1991); Ritov and Baron (1990)).

Regret aversion can also explain the *endowment effect*, a preference for people to hold on to what they have rather than exchange for a better alternative, as with the refusal of individuals to swap a lottery ticket for an equivalent one plus cash.[12] The endowment effect has been confirmed in experimental securities markets.

2.4 Conservatism and Representativeness

A final dimension to self-deception that relates to financial markets is the propensity of individuals to depart from Bayes' theorem in predictable ways. Edwards (1968) identified and later studies confirmed the phenomenon of *conservatism*, that under appropriate circumstances individuals do not change their beliefs as much as would a rational Bayesian in the face of new evidence. The more useful the evidence, the greater is the shortfall between actual updating and rational updating.

Edwards indicated that it takes approximately two to five observations to cause a subject to change his opinions as much as he ought based on a single observation. A tendency toward conservatism makes sense as a means of maintaining self-esteem. An

[12]See Bar-Hillel and Neter (1996), Kahneman, Knetsch, and Thaler (1991), Knez, Smith, and Williams (1985).

individual who changes his beliefs is implicitly recognizing that his previous beliefs were not entirely correct. To put this differently, as Barberis, Shleifer, and Vishny (1998) mention (p. 315), such an individual can be viewed as overconfident about his prior information. Such overconfidence fits with self-deception: I believe I am very good at acquiring good information, so I stick with it.

Finance models designed to capture overconfidence often focus on confidence about a new signal, and therefore involve moving too far from the prior toward the new information. This discussion of self-deception and the psychological evidence suggests that this is not necessarily the case. Individuals should hold too close to the prior if the prior is strongly associated with the individual's self-esteem, and move with the new information if the new information is more strongly linked to self-esteem. For example, Daniel, Hirshleifer, and Subrahmanyam (1998) argue that a private signal will tend to be tied closely to self-esteem, so that overconfidence should be about private but not public signals. Presumably the individual's prior belief reflects a combination of past private and public signals, so it should be more closely tied to self-esteem than a new public signal, and less tied to self-esteem than a new private signal.

The representativeness heuristic (Kahneman and Tversky (1973)) involves decisions about whether a specific item belongs to a general class. The heuristic involves judging the item as probably belonging to the class if it shares some similarity to the class. This can be inaccurate because greater similarity is not always associated with higher probability. For example, the son of a man who is 6'7" tall is more likely to be 6'6" tall than 6'7" tall; big effects are not always the result of big causes; a fair coin is less likely to produce TTH than TT.

Since similarity often is associated with higher probability, this heuristic helps economize on limited cognitive capacity. This heuristic can help people detect patterns more quickly by looking for similarities. Furthermore, overconfidence should cause such a bias. The ability to discern meaningful patterns from complex inputs is an important ability. An overconfident individual who overestimates this ability should tend to treat identified patterns as meaningful that are actually coincidental. The financial industry of technical analysis is arguably based upon the excessive faith of individuals in their own abilities to detect meaningful patterns.

One explanation for conservatism is that processing new information and updating beliefs is costly. There is evidence that information that is presented in a cognitively costly form is weighed less: information that is abstract and statistical, such as sample size and probabilistic base-rate information. Furthermore, people may overreact to

information that is easily processed, i.e., scenarios and concrete examples.

The costly-processing argument can be extended to explain base rate underweighting. If an individual underweights new information received about population frequencies (base rates), then base rate underweighting is really a form of conservatism. Indeed, base rates are underweighted less when they are presented in more salient form or in a fashion which emphasizes their causal relation to the decision problem (see Koehler (1996)). This costly-processing-of-new-information argument does not suggest that an individual will underweight his pre-existing internalized prior belief. On the other hand, if base rate underweighting is a consequence of the use of the representativeness heuristic, there *should* be underweighting of priors.

Griffin and Tversky (1992) suggest that base-rate underweighting and conservatism, interpreted as under- versus over-reaction to signals, can be understood as results of excessive reliance on the *strength* of information signals and underreliance on the *weight* of information signals. The strength of an information realization is how 'extreme' the evidence is (in some sense), and the weight of evidence is its reliability or precision. For example, a large sample of conditionally independent and identically distributed information signals has high weight. But if the preponderance of favorable over unfavorable signals is modest, it has low strength. Conservatism arises when people rely too little on high weight evidence such as a long sample, and base rate neglect when people rely too heavily on high-strength evidence such as a few signals all in one direction.

In reality, people try to recognize their own psychological failings and often take steps to adjust for them. For example, Shefrin and Statman (1984) suggested that individual investor demand for high-dividend-yield stocks is motivated by a desire for an external means of disciplining consumption. Such adjustment is imperfect for many reasons, but a notable one is overconfidence. An overconfident investor may be prone to deny his biases rather than recognize and compensate for them.

In summary, different experimental settings can lead to under- or over-reliance on new signals; people seem to make judgments differently in different situations (see Grether (1992), Payne, Bettman, and Johnson (1992)). Given the different possible effects, invoking the name of a bias does not provide compelling support for assuming under- or over-reaction in a financial model. Further support can be provided by comparing the economic decision environment of the model with the specific experimental decision setting in which the bias was documented, and especially by running new experiments that match closely the decision environment in the financial model.

3 Self-Deception in Financial Markets

We now discuss how the above notions of self-deception may help us understand the behavior of financial markets. Fama's famous early survey (Fama (1970)) contrasted 'semi-strong form efficiency', wherein market prices fully reflect all available public information, with 'strong form efficiency', wherein market prices fully reflect even private information. It is well accepted that markets are not strong form efficient. But a bit of thought suggests that there is no absolute division between public and private information. Information in the newspaper is public, and information known only to an individual is private. But what if only a few people receive that newspaper? If only one person has the information, but he is happy to tell anyone who asks, is the information private? What if the individual shouts the information far and wide, but no one will listen? Is it very different for an individual to expend time and resources to read something in the newspaper (the contents of which are generally viewed as public information by financial economists) versus talking to another individual?

To consider such issues clearly we must take into account costs of conveying and processing information. Private information is just the extreme case of information that is hard to dig out. But the expense of digging it out may come not just from incentives for secrecy, but from the form it is presented (consider a pile of balance sheets versus a well-written analyst's report), the outlets at which the information has been certified, and so on. Information that is conveyed briefly and lucidly will have higher impact than information that is harder to make sense out of. In other words, some forms of information are more salient than others. So far we have talked about some good reasons for differences in salience: brevity, clarity, certification. However, an individual with limited information processing capacity will imperfectly optimize the decision of what information to process. Going beyond this, it is reasonable to think that there are some systematic biases in the way that irrelevant considerations affect salience– as with the use of sexually attractive models in product advertisements. Most studies of price forecasts find biases that are predictable using current observables. For example, forecasts are often found to be adaptive, i.e., they respond partially to past forecast errors.[13] Such biases are potentially consistent either with Bayesian learning with an unknown distribution, or with overconfidence. Experimental studies involving a fixed distribution generally also yield biases, and forecasts are adaptive in most forecast

[13]See the discussions in Lovell (1986) and Williams (1987), but see also Keane and Runkle (1990). There is a similar finding for survey forecasts of macroeconomic variables (e.g., Aggarwal, Mohanty, and Song (1995)).

experiments involving endogenously determined prices as well (see Camerer (1995)). Consistent with overconfidence, forecasters seem to put too little weight on the known forecasts of other forecasters (Batchelor and Dua (1992)).

Analyst forecasts of earnings are over-optimistic at long time horizons and pessimistic at short horizons (e.g., Richardson, Teoh, and Wysocki (2004)). Such biases may come from misperceptions or from agency incentives. However, we would normally expect rational agents to provide at least a positive incremental value in their forecasting activities. A large literature shows that real-world decisionmakers such as PhD admission committees or doctors do not predict outcomes as well as mechanical decision rules based on simple linear combinations of objective input measures (see Camerer (1991)). This suggests that the rise of arbitrage based upon modern statistical analysis in securities markets will indeed reduce mispricing.

3.1 Overconfidence and Asset Prices

In a model with multiple securities, Daniel, Hirshleifer, and Subrahmanyam (2001) provide a model of the pricing of assets and securities (analogous to the traditional Capital Asset Pricing Model, or CAPM, of financial theory) when investors are overconfident. In the traditional CAPM, a security's expected return is increasing in the security's risk. This reflects the notion that investors must be compensated for bearing high risk. The appropriate measure of risk in the model is the marginal contribution an extra unit of the security adds to the variability of a typical investor's portfolio. It turns out that this marginal contribution to risk is equal to the theoretical regression coefficient of the security's return on the return on the 'market portfolio,' a portfolio consisting of all invested wealth in the asset market. This regression coefficient is generally referred to as the security's 'beta.'

In the overconfidence-based model, security terminal cash flows satisfy a linear factor model, and each investor observes signals about the factors and the idiosyncratic component of security payoffs. Risk-averse investors form what they perceive to be mean-variance efficient portfolios. Overconfident individuals trade with risk-averse arbitrageurs who form rational beliefs. A security's equilibrium expected return is linearly increasing in the security's beta with the market, and the security's current mispricing. Variables containing market price are proxies for the security's misvaluation. For example, a fundamental/price ratio such as book/market is driven down when favorable news drives a stock up. Since there is overreaction, this is when the stock is overvalued.

Thus, a high fundamental/price ratio predicts high future returns. Aggregate value measures such as the market dividend yield or book/market positively predict future market returns.

Fundamental/price ratios mix two sources of predictive power. A fundamental/price ratio (*e.g.*, high book/market) tends to be high if either risk is high or if the market has overreacted to a highly adverse signal. In either case, price on average rises. Since high book/market reflects both mispricing and risk, whereas beta reflects only risk, book/market tends to be a better predictor of returns. These two sources of predictive power are unequal. Beta helps disentangle these cases, so beta and book/market are joint predictors of future returns.

However, when overconfidence becomes very strong, and if the proxy for the unconditional expected value (e.g., book value) is perfect, then the incremental ability of beta to predict future returns vanishes. The fundamental/price ratio dominates beta even though risk is priced. This is an extreme case, but it helps explain why empirical findings on the incremental effect of beta have been weak and inconsistent. The model also implies that in univariate regressions beta should predict future returns. The model further describes the tradeoffs in constructing optimal price-related proxies for misvaluation.

Another puzzling pattern for fully rational asset pricing theory is the forward discount puzzle. High bond yields in a country relative to another country should be an indicator that that country's currency is likely to have higher inflation, and therefore is expected by investors to depreciate relative to the other country's currency. The higher bond yield (in nominal terms) then provides investors with appropriate compensation for the greater expected real decline in the value of the first country's currency. This bond yield differential is called the 'forward discount.' The puzzle is that a higher than usual forward discount between two countries is associated on average with *appreciation* rather than *depreciation* in the first country's currency. In other words, investing in high-interest-rate countries appears to be a 'good deal' for investors.

Several authors have commented on the difficulties that fully rational models have in explaining this puzzle (see, e.g., Engel (1996)). A possible explanation for this puzzle is provided by investor overconfidence (see Hirshleifer and Wang (2001)). A large forward discount may reflect an overreaction by overconfident investors to information (perhaps spurious) about likely high inflation. This overreaction drives down the value of the first country's currency too much relative to the other country's currency (i.e., it drives up the exchange rate). Eventually correction occurs as information arrives that forces

investors to revise their beliefs. This correction leads to appreciation of the first country's currency.

Static models provide simple generalizations of the insights of the traditional Capital Asset Pricing Model that can encompass the effect of risk as well as mispricing. However, a static approach has no hope of capturing the distinction between short-term continuation and long-term reversals.

In both static and dynamic models, long-run reversal occurs when there is an overreaction to an impulse such as the arrival of good news. In a dynamic setting, short-run positive autocorrelation is consistent with long-run reversal so long as the process of overreaction and correction is sufficiently smooth. Such smoothness implies that when an impulse sets price rising, it will probably rise some more; that on average the last up-move to the peak of the impulse response function is not followed by a precipitous drop; and when the price is falling, it tends to fall some more. In contrast, a long-lag autocorrelation tends to associate positive returns during the overreaction process with negative returns arising during the correction process.

The subsections that follow describes the effects of *pure* (independent) noise trading, mechanistic models based on correlated trading (positive feedback), the effects of mistaken beliefs, and the effects of alternative preferences.

3.2 Pure Noise Trading

Pure noise trading and positive feedback trading cause overreaction, and hence negative autocorrelations in long-run returns. When a stock rises too high, it needs to correct back down. Equivalently, this overreaction causes excess volatility in returns. Furthermore, Campbell and Kyle (1993) showed that overreaction can cause aggregate stock market value measures such as dividend yield to predict future market returns, so that contrarian investment strategies are on average profitable.

DeLong, Shleifer, Summers, and Waldmann (1990a) model the consequences of unpredictable random trades. Two securities pay identical, riskless dividends. The price of one asset is exogenously fixed. The other asset is risky because pure noise trades cause stochastic mispricing. Rational arbitrageurs with exogenous short time horizons limit their arbitrage trades for fear that the mispricing will get worse before it gets better. On average the risky asset trades at a discount, the risk premium demanded by the rational investors.

The noise trading approach provides an explanation for the existence and behavior

of closed-end fund discounts and their correlations with stock returns. According to DeLong, Shleifer, Summers, and Waldmann (1990a), noise traders buy and sell closed-end funds in a correlated fashion, causing discounts or premia relative to net asset value to fluctuate. The mispricing risk this creates makes these funds less attractive to rational investors, so on average there is a discount. This theory implies that fund discounts move together based on a systematic noise-trading factor; such comovement exists (e.g., Lee, Shleifer, and Thaler (1991)). The theory also explains why such funds are created: to exploit optimistic noise traders.

Lee, Shleifer, and Thaler (1991) suggest that shifts in fund discounts reflect shifts in noise trader sentiment toward all small stocks. This is consistent with their evidence that narrowing of closed-end fund discounts is associated contemporaneously with high small stock returns. This implies that discounts predict small stock returns. If discounts were a consequence of *pure* noise trading, they would be uncorrelated with future fundamentals such as as accounting performance. Swaminathan (1996) finds that at lags of greater than one year high discounts predict both low future accounting profits and high future stock returns. This is consistent with fund investors overreacting to genuine information.

The comovement in small stock returns documented in Fama and French (1993) may come from correlated imperfectly rational trades. The DeLong, Shleifer, Summers, and Waldmann (1990a) approach then suggests that small stocks, including closed-end fund shares, will earn high expected returns in compensation for their high mispricing risk. Alternatively, low market-value stocks may earn high returns because a stock's low market value on average derives partly from its being undervalued (see, e.g., Daniel, Hirshleifer, and Subrahmanyam (2001)). The U.S. small firm effect has been weak or absent in the last 15 years, yet closed-end fund discounts remain.

3.3 Positive Feedback Trading

Positive feedback trading has several possible motivations, one being that investors form expectations of future prices by extrapolating trends, consistent with the representativeness heuristic. DeLong, Shleifer, Summers, and Waldmann (1990b) offer a model with a risky asset and riskfree cash, in which information arrives sequentially. The exogenous date 2 demand of the positive feedback traders is linearly increasing in the preceding price trend. Forseeing this, rational speculators buy into price trends, exaggerating trends and overshooting. As a result there is excess volatility, and long-term negative

autocorrelations in returns.

In Cutler, Poterba, and Summers (1990), there are two types of imperfectly rational traders (i.e. traders overconfident in their understanding of the market), positive feedback traders, and fundamental traders who ignore price and trade based upon a signal about the security's payoff. Some fundamental traders observe this signal with a lag. This lag creates price trends that are profitably exploited by feedback traders. The gradual process of overshooting and correction induces both short-lag positive autocorrelation and long-lag negative autocorrelation. In the model of DeLong, Shleifer, Summers, and Waldmann (1990b), positive feedback trading loses money. In contrast, in Cutler et al positive feedback traders can potentially make money from their activities. These potential profits come from being able to exploit underreaction.

More recent models with endogenous decisions have found things akin to pure noise trading— a limiting case of overconfidence, and positive feedback trading. But endogenously derived positive feedback is conditional and statistical, which seems more realistic than the older models. For reasons of both descriptiveness and predictiveness, explicit modeling of the psychology of investors is likely to supersede the mechanistic approach (except perhaps in otherwise-intractable applications).

Two recent papers provide models with a single risky security that reflect the fact that people learn about their own abilities in a biased, self-promoting fashion. In these models, investors do not know the precision of their private information signals, which reflects their information-gathering ability. They learn about their precision through time by observing whether later public news confirms or disconfirms their previous signal. The analyses assume the dynamic complement of overconfidence, biased self-attribution. When an investor receives confirming news his confidence in his precision rises too much, and when there is disconfirming news his confidence declines too little.

In Daniel, Hirshleifer, and Subrahmanyam (1998), the impulse response function to a favorable initial shock, the private information signal, is hump-shaped. Price on average rises further as public information arrives, because confidence about the private signal on average grows. Eventually, however, accumulating evidence forces investors back to a more reasonable self-perception. This smooth hump-shaped impulse response implies positive short-lag and negative long-lag return autocorrelations. Daniel, Hirshleifer, and Subrahmanyam (1998) also numerically simulate the correlation of a public information surprise (such as favorable accounting performance) with future returns with self-attribution bias. At short lags this correlation is positive, but at long lags the correlation can be negative.

Gervais and Odean (2001) provide a model that accommodates analytical solution for the learning process under biased self-attribution. As traders become overconfident trading volume and market return volatility increase. Since equity is in positive net supply, the model also predicts that trading volume will be higher after market rises than market falls, consistent with Statman and Thorley (1998).[14]

3.4 Excessive Trading

It has been argued that the volume of trade in speculative markets is too large, and overconfidence of traders has been advanced as an explanation (e.g., DeBondt and Thaler (1995)). Whether volume is too large is hard to establish without a benchmark rational level of volume. Rational dynamic hedging strategies, in principle, can generate enormous volume with moderate amounts of news.

However, as DeBondt and Thaler point out, support for overconfidence is provided by evidence suggesting that more active investors earn lower returns as a result of incurring higher transaction costs (Odean (1999), Barber and Odean (2000)).[15] Also consistent with overconfidence, traders in experimental markets do not place enough weight on the information and actions of others (Bloomfield, Libby, and Nelson (1999)). In experimental markets, investors also tend to overreact more to unreliable than to reliable information (Bloomfield, Libby, and Nelson (2000)).

Barber and Odean (2002) find that investors who have experienced the greatest past success in trading are the most likely to switch to online trading, and will trade the most in the future. This evidence is consistent with self-attribution bias, meaning that the investors have likely attributed their past success to skill rather than to luck. Also, there is some evidence that access to Internet trading appears to encourage more active trading (Choi, Laibson, and Metrick (2002)). There is also evidence of overconfidence even by experienced dealers in an 'opaque' experimental asset market, but not in a transparent one (Lamoureux and Schnitzlein (2001)).

[14]The implication of attribution/overconfidence models for whether there should be something akin to a disposition effect (holding winners, selling losers) is not obvious. When a stock is first becoming a winner, rational arbitrageurs who foresee further price rises should drive the price up even higher than the overconfident think is justified. This encourages the overconfident to sell, consistent with the disposition effect. However, for a stock that has been a winner for some time, the arbitrageurs will sell to the overconfident as the price peaks. Other recent models of momentum and reversal have similar opposing effects.

[15]Coval, Hirshleifer, and Shumway (2003), who find that a subset of active investors are able to consistently outperform the market, point out that 'excessive' trading may be rational if it allows individuals to learn whether they have ability.

3.5 The Dynamics of Representativeness and Conservatism in Financial Markets

Barberis, Shleifer, and Vishny (1998) offer an explanation for under- and over-reactions based on a model in which actual earnings for a risky asset follow a random walk, but investors do not understand this. They mistakenly believe that the earnings process stochastically fluctuates between a regime with mean-reverting earnings, and a regime with expected earnings growth. Apart from misunderstanding the structure of the earnings process, the representative investor updates as a rational Bayesian. This structure is designed to capture the representativeness heuristic and conservatism bias.

If recent earnings changes reverse, investors erroneously believe the firm is in a mean-reverting state, and underreact to recent news, consistent with conservatism. If investors see a sequence of growing earnings, they tend to conclude (wrongly) that the firm is in a growth regime, and overextrapolate trends, which is arguably reminiscent of representativeness. Overreaction to a long enough trend implies subsequent low returns during the process of correction. Thus, there can be long-term overreaction and correction, implying negative long-lag return autocorrelation. Yet the average response to an initial impulse can be smooth, implying positive short lag autocorrelation. Similarly, the model can accommodate a positive short-term correlation between the asset return and an earnings change, and a negative long-term correlation. If sporadic events such as dividend initiations are viewed as isolated from earnings patterns, a single-event version of the model applies implies, under appropriate parameter values, underreaction.

Cross-sectional effects (such as a value effect) are simulated with earnings that are independently distributed across stocks. This implies a nearly riskfree arbitrage opportunity for a rational investor who buys and sells stocks based on return predictors. Such arbitrage would be risky in a setting where investors update their beliefs about systematic factors in earnings trends or reversals. The psychological literature on multiple cue learning may provide guidance for such a model.

Hong and Stein (1999) analyze a market in which, as in Cutler, Poterba, and Summers (1990), some traders react sluggishly, and others trade based on positive feedback. Each group of traders is risk averse, and is able to process only a subset of available information. Information about the liquidating dividend dribbles into the hands of different groups of newswatchers. Newswatchers condition on their own private signals but ignore market prices, causing underreaction.

Momentum traders, in contrast, condition on the cumulative price change over the

last k periods. Each trader takes a fixed position for a given number of periods. Momentum traders exploit the underreaction of newswatchers by buying in response to price increases. This accelerates the reaction to news, but also causes overshooting. The smoothness of the overreaction process causes positive short-lag and negative long-lag autocorrelation. Slower information diffusion tends to launch a more powerful an overreaction, leading to more negative long-lag autocorrelations.

4 Self-Deception and Deception

One might expect the tendency to deceive oneself to open the door to deception by others. Both self-deception and the deception of others may be important in financial markets. For instance, the tendency to self-deception can make it easier for predators to deceive and exploit investors.[16] Overconfidence in one's ability to interpret signals can lead to excessive credulity in interpreting the signals transmitted by others. As a consequence, accounting and reporting policy can differ widely across firms and can influence actual firm valuation and performance.

In simple, fully rational settings, if disclosure is costless, all information is disclosed. Intuitively, in equilibrium investors are highly skeptical of a decision by a firm to withhold information.[17] There are some qualifications to this conclusion based upon proprietary costs, firms that do not receive information, and signaling incentives.[18] When there is a non-verifiable information signal as well as a verifiable one, it can be an equilibrium for the firm to withhold (fail to disclose) favorable verifiable information as a 'macho' signal that it has favorable unverifiable information (Teoh and Hwang (1991)).

Excessive credulity among investors further suggests the "unraveling" results of Grossman (1981) and Milgrom (1981) may no longer hold. Individuals who are overconfident may underestimate the importance of information they do not possess, thereby causing them to underweight the importance of non-disclosure. As Jonathan Swift put it, "The most positive men are the most credulous." [*Thoughts on Various Subjects*,

[16] Coval, Hirshleifer, and Thakor (2004) analyze the incentives of financial advisers and intermediaries to exploit imperfectly rational investors.

[17] In the most basic possible setting, there is rationally extreme skepticism of failure to disclose–the 'unraveling' results of Grossman (1981) and Milgrom (1981).

[18] If there are costs of disclosure (for example, if the information that can be disclosed has value for competitors) the unraveling may be limited; there remains a pool of low-quality types that do not disclose (Verrecchia (1983)). Similarly, if there is a probability that the firm does not possess information, then observers will not be completely skeptical about a firm that does not disclose any information (Dye (1985)).

(1711).]" In this regard, it is interesting that experts in some experiments are prone than non-experts to ignore decision tools that would be useful for both experts and non-experts alike for accurately performing an assigned task (Arkes, Dawes, and Christensen (1986)).

A variety of evidence indicates that individuals are insufficiently skeptical about firms' motives for refraining from disclosing information. For example, Mathios (2000) examined the effect of the Nutrition Labeling and Education Act on purchases of salad dressing, which made mandatory the labeling of information about fat content. He found that even though there was voluntary labeling (mostly of low-fat brands) prior to the regulation, mandatory disclosure caused high-fat dressings to lose market share. Hanson and Kysar (1999) review a literature in consumer psychology and marketing on the ability of sellers to manipulate consumer perceptions of their products.

Investors also seem to be insufficiently skeptical of firms that refrain from disclosing information, or that disclose in a non-salient fashion. For example, there is evidence that firms tend to release good news early and bad news late,[19] the exception being that the possibility of litigation can induce disclosure of bad news (Skinner (1994)). Such behavior is consistent with a fully rational equilibrium with proprietary disclosure costs (e.g., Verrecchia (1990), Darrough and Stoughton (1990), Feltham and Xie (1992)). However, this raises the question of whether such costs are high enough to explain this bias. Excessive investor credulity strengthens the incentive of firms to behave in such a fashion.

As a consequence of excessive credulity, it is plausible that a partial disclosure equilibrium analogous to that of the Verrecchia model will obtain– see Hirshleifer, Lim, and Teoh (2003). Firms with more favorable information disclose. But firms with sufficiently adverse information (below some cutoff) withhold information and delay revelation.

The combination of limited attention with overconfidence may be more potent in generating credulity than limited attention alone. If individuals are overconfident, they will fail to discount fully for the fact that they are ignoring some of the considerations relevant for their decisions. In consequence investors may be fooled by actions firms can take to manipulate perceptions.

The example of new equity issues and repurchases illustrates how limited attention and overconfidence may affect firms' incentives to take informative actions. A manager who has an incentive to maintain a high stock price will try to make profits in his firm's

[19]See Chambers and Penman (1984), McNichols (1989), Begley and Fischer (1998), and Haw, Qi, and Wu (2000)).

share trading. This encourages issuance of new shares when the stock is overvalued (owing either to information asymmetry or to market irrationality) and repurchase when the stock is undervalued. If the market is credulous and fails to discount fully for this incentive, then the manager will indeed get a good price for the shares it buys and sells through this procedure. The market may understand that new issues are an adverse indicator of value, consistent with a negative stock price reaction, but being insufficiently skeptical, the price does not fall enough. This leads to a long-run negative return. Similarly, consistent with the evidence, this story suggests a positive price reaction to repurchase and a long-run positive average abnormal post-event return.

Similarly, there is evidence that positive earnings adjustments of firms (accounting accruals) are associated with low subsequent abnormal stock returns (Sloan (1996)). This suggests that firms are able to induce overvaluation through cosmetic accounting adjustments. Consistent with this hypothesis, this effect comes mainly from the discretionary component of accruals (Teoh, Welch, and Wong (1998a), Teoh, Welch, and Wong (1998b)), and is particularly strong just before initial public offerings of equity as well as seasoned equity offerings (Teoh, Welch, and Wong (1998a), Teoh, Welch, and Wong (1998b), Teoh, Wong, and Rao (1998)). Even the forecasts of earnings made by professional stock analysts are biased by an amount that depends on the level of the firm's accruals (Teoh and Wong (2002)).

Managers generally like high stock prices, so stocks that are more subject to investor credulity should on the whole tend to be overvalued. The problem of credulity is likely to be greater for firms that are able to weave hard-to-refute stories to tell investors about future prospects. Thus, empirical findings of inferior performance of stocks with low book/market ratios, and the stronger relation of book/market to returns among high R&D firms (Chan, Lakonishok, and Sougiannis (2001)), are consistent with credulity.[20]

Firms also seem to have a particular distaste for required disclosure of liabilities. Fully rational investor skepticism should force voluntary revelation of such information. Indeed, the fact that a firm or industry organization would campaign for secrecy would itself seem to reveal bad news. However, if investors are excessively credulous, rules forcing more disclosure can be helpful. This reasoning provides a motivation for mandatory

[20]Daniel, Hirshleifer, and Subrahmanyam (2001) provide an overconfidence-based explanation for such effects based upon overreaction to private information signals rather than credulity about managerial incentives. These accounts are reconcilable if management is able to manipulate not just public information sets but also information which investors perceive to be 'private.' For example, investors may trade upon information provided to them by analysts, even though this information is fed to analysts by management.

disclosure rules.[21]

More broadly, Hirshleifer and Teoh (2003) provide a model in which investor interpretations of financial reports and market prices are influenced by limited attention. For example, limited attention causes credulity about so-called pro forma earnings disclosures (extra earnings disclosures outside of financial statements that do not have to satisfy Generally Accepted Accounting Principles). It also causes the market to overvalue firms that have high commitments of employee stock option compensation, as investors fail to take into account the dilution in stock value that would occur if these options are eventually exercised. Furthermore, the market misvalues diversified firms as it fails to disentangle properly the differing growth rates of the the components of earnings derived from different business segments.

Finally, a recent study supports the widely-held view that analysts and investors may have been excessively credulous in interpreting information released by Enron prior to its 2001 collapse. Examining posts to Enron message boards on the Internet, Felton and Kim (2002) found a variety of detailed warnings regarding Enron's off-balance sheet liabilities that turned out to be extremely prescient.[22] Of course, it is not obvious that even appropriately skeptical analysts should be able to have the foresight of the best-informed individual posting to an Internet message board.

5 Discussion

Although it is not obvious how the empirical securities market phenomena can be captured plausibly in a model based on perfect investor rationality, no psychological ("behavioral") theory for these phenomena has won general acceptance. Why is this the case? The criticisms of psychology-based theories can be classified into four major categories. First, the use of psychological theories in explaining securities market behavior

[21]Coffee (1984) (p. 745-6) argues that critical adverse information was withheld from investors in municipal bond markets in which the Securities and Exchange Commisssion disclosure requirements do not apply, and that bond rating agencies were ineffective alternative sources of information to investors, leading to such problems as the problems with New York City's bond offerings in the 1970's, and the Washington Public Power System's failure in the 1980s. Nevertheless, Palmiter (1999) argues that in private placement offerings (which are exempt from Securities Act of 1933 disclosure requirements) issuers generally disclose information similar to or going beyond what is required for registered offerings, and many foreign issuers voluntarily choose to list on the New York Stock Exchange and reconcile its financial reports with Generally Accepted Accounting Principles (GAAP). On the other hand, in many countries it is evident that the level of disclosure voluntarily achieved is much less than that provided in the U.S.

[22]For example, one post from March 2000 warned, "Dig deep behind the Enron financials and you'll see a growing mountain of off-balance sheet debt which will eventually swallow this company."

is often viewed as an *ad hoc* exercise. Because of the wide array of evidence in the experimental psychology literature, skeptics feel that psychology-based researchers can easily "psycho-mine" the literature to find whatever evidence they need to account for the particular securities market phenomenon they are studying.

Second, although the array of experimental evidence in psychology is wide, it is not clear that the laboratory setting is sufficiently close to that of financial market. For instance, little direct evidence exists as to whether psychological biases are important in influencing investor behavior when significant capital is on the line. On the other hand, the general finding from psychological studies has been that increased pecuniary incentives tend to moderate, but do not eliminate biases.

Third, even if psychological biases are important at the individual level, it is often argued that they should cancel out in aggregate. What is the mechanism that causes systematic patterns?

A final criticism of a psychology-based approach is that one would expect that, with time, psychological biases in financial markets disappear. What prevents individuals with psychological biases from either learning about and eliminating their biases or from losing so much money that they drop out of the market? We address each of these criticisms in turn.

A general criticism often leveled by economists against psychological theories is that, in a given economic setting, the universe of conceivable irrational behavior patterns is essentially unrestricted. Why overconfidence, why not underconfidence? What does this have to do with the representative heuristic, or cognitive dissonance? Does summary of empirical patterns from psychology experiments deserve the name 'theory'? Thus, it is sometimes claimed that allowing for irrationality opens a Pandora's box of *ad hoc* stories that will have little out-of-sample predictive power. In a specific investment setting, it can be hard to judge which documented psychological bias is relevant. This creates an extra degree of freedom for model-mining not present in the purely rational approach. However, DeBondt and Thaler (1995) argue that a good psychological finance theory will be grounded on psychological evidence about how people actually behave. We concur, and also believe that such a theory should allow for the rational side of investor decisions. To deserve consideration a theory should be parsimonious, explain a *range* of anomalous patterns in different contexts, and generate new empirical implications.

Financial economists are familiar with criticisms of psychological experiments: that the stakes are low, that subjects have little experience with the experimental setting, that there is weak incentive to pay attention or tell the truth, and that publication depends

on finding an effect. What may not be as familiar is that there is data addressing these issues. On the whole training and increasing rewards and number of repetitions often reduces, but does not eliminate biases. Lessons learned through repetition often do not carry over well across seemingly similar tasks. The well-known biases have been subjected to replication.[23] Many (though not all) of the cognitive biases are stronger for individuals with low cognitive ability or skills than for those with high ability or skills, consistent with biases being genuine errors (see Stanovich and West (2000)).

Economists also argue that errors are independent across individuals, and therefore cancel out in equilibrium. Why should we expect any of these phenomena to matter in the aggregate? People share similar heuristics, those that worked well in our evolutionary past. So on the whole we should be subject to similar biases. Systematic biases (common to most people, and predictable based upon the nature of the decision problem) have been confirmed in a vast literature in experimental psychology. The evidence simply does not support the idea that the errors people make are unrelated, or independent of the type of decision problems.

One of the reasons that financial economists have traditionally adhered to a fully rational framework is that it has a powerful motivation. People who behave this way will tend to do well, and so in an evolutionary sense should dominate the economy. More deeply, at a longer time scale the human mind evolved under pressure of natural selection, which should design a mind that works well. However, as discussed in Section 2, there are strong arguments for the survival of self-deception in the prehistoric environment in which humans evolved.

In the financial market setting, one might expect self-deception to be severely penalized and selected against. However, several papers argue the opposite: that self-deception can help individuals prosper in financial markets. Kyle and Wang (1997) argue that overconfident traders, by committing to trade aggressively on their information, can obtain an above-average share of informed-trading profits. Similarly, DeLong, Shleifer, Summers, and Waldmann (1990a) and Hirshleifer and Luo (2001) argue that overconfident traders, because they create mispricing risk and because they underestimate the level of risk that they assume, may earn greater expected returns than rational traders. Although they assume correspondingly high levels of risk, if survival in financial markets depends on terminal wealth levels, under appropriate conditions overconfident traders thrive relative to rational traders. Other papers that conclude that irrationality may persist in financial market settings include Blume and Easley (1990), Benos (1998),

[23]See, e.g., Camerer (1995), Rabin (1998) and Hertwig and Ortmann (2001).

and Wang (2001).

Why don't people simply learn their way out of biased judgments? To some extent they do. One barrier is that learning is just too hard. Some of these biases are the side effects of heuristics that are necessary to cut the Gordian knot and come to decision in the face of a complicated environment. The other barrier arises from self-deception itself. Individuals who think they are already competent may be slow to adjust their decision procedures (e.g., Einhorn and Hogarth (1978)).

In general, psychology-based theories of securities market phenomenon are not necessarily direct substitutes for traditional finance theory. Much of traditional finance theory is compatible with psychological effects. Indeed, central to any robust psychology-based theory of financial markets will be the traditional analysis of how individual decisions aggregate, through market clearing, to determine equilibrium prices. It will be at the individual level, with the introduction of judgment and decision biases, where the traditional and psychological approaches depart.

While the degree of efficiency of securities prices remains a controversial issue, we have argued that self-deception holds promise as a basis for understanding capital market behavior. Self-deception unifies several regularities in psychology and finance in an appealing way, provides a variety of empirical predictions, and survives generic critiques of psychology-based financial theories. Biases consistent with self-deception have emerged in a range of experimental settings involving many subjects, which suggests that the effects of self-deception are unlikely to cancel out in the aggregate. There are plausible arguments for why self-deceiving beliefs and learning processes could have survived and proliferated in the prehistoric environment of human evolutionary adaptation. In consequence, it is reasonable to expect such beliefs will tend to persist in modern financial market settings. Several intriguing issues remain to be explored. These include the questions of when self-deception is most likely to operate unchecked (e.g. when penalties are low and potential benefits are large), what other implications does self-deception have for financial behavior, and whether the further empirical predictions of theories based upon self-deception are borne out in the data.

References

Aggarwal, R., S. Mohanty, and F. Song, 1995, Are survey forecasts of macroeconomic variables rational?, *Journal of Business* **68**, 99–119.

Alpert, M. and H. Raiffa, 1982,. A progress report on the training of probability assessors, . In D. Kahneman, P. Slovic, and A. Tversky, eds., *Judgement under Uncertainty: Heuristics and Biases*. (Cambridge University Press, Cambridge).

Arkes, H. R. and C. Blumer, 1985, The psychology of sunk costs, *Organizational Behavior and Human Decision Processes* **35**, 124–140.

Arkes, H. R., R. M. Dawes, and C. Christensen, 1986, Factors influencing the use of a decision rule in a probabilistic task, *Organizational Behavior and Human Decision Processes* **37**, 93–110.

Arkes, H. R., 2002, Incremental investment decisions and the diagnosticity of feedback, working paper, The Ohio State University.

Arkes, H. and D. Hirshleifer, 2004, Integrating heuristics and biases: An evolutionary approach, Ohio State University, in progress.

Baker, M. P. and J. Wurgler, 2002, Market timing and capital structure, *Journal of Finance* **57**, 1–32.

Barberis, N., A. Shleifer, and R. Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* **49**, 307–343.

Barber, B. and T. Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *Journal of Finance* **55**, 773–806.

Barber, B. and T. Odean, 2002, Online investors: Do the slow die first?, *Review of Financial Studies* **15**, 455–488.

Bar-Hillel, M. and E. Neter, 1996, Why are people reluctant to exchange lottery tickets?, *Journal of Personality and Social Psychology* **70**, 17–27.

Batchelor, R. and P. Dua, 1992, Conservatism and consensus-seeking among economic forecasters, *Journal of Forecasting* **11**, 169–181.

Begley, J. and P. E. Fischer, 1998, Is there information in an earnings announcement delay?, *Review of Accounting Studies* **3**, 347–363.

Bell, D. E., 1982, Regret in decision making under uncertainty, *Operations Research* **30**, 961–981.

Benos, A., 1998, Aggressiveness and survival of overconfident traders, *Journal of Financial Markets* **1**, 353–83.

Berk, J. and R. C. Green, 2003, Mutual fund flows and performance in rational markets, NBER Working Paper No. W9275.

Bernardo, A. and I. Welch, 2001, On the evolution of overconfidence and entrepreneurs, *Journal of Economics and Management Strategy* **10**, 301–330.

Bikhchandani, S., D. Hirshleifer, and I. Welch, 1998, Learning from the behavior of others: Conformity, fads, and informational cascades, *Journal of Economic Perspectives* **12**, 151–170.

Black, F., 1986, Noise, *Journal of Finance* **41**, 529–543.

Bloomfield, R., R. Libby, and M. Nelson, 1999, Confidence and the welfare of less-informed investors, *Accounting, Organization and Society* **24**, 623–647.

Bloomfield, R., R. Libby, and M. Nelson, 2000, Underreactions, overreactions and moderated confidence, *Journal of Financial Markets* **3**, 113–137.

Blume, L. and D. Easley, 1990, Evolution and market behavior, *Journal of Economic Theory* **58**, 9–40.

Brenner, L. A., D. J. Koehler, V. Liberman, and A. Tversky, 1996, Overconfidence in probability and frequency judgments: A critical examination, *Organizational Behavior and Human Decision Processes* **65**, 212–219.

Camerer, C. F., 1991,. The process-performance paradox in expert judgment: How can experts know so much and predict so badly?, . In K. A. Ericsson and J. Smith, eds., *Towards a general theory of expertise: Prospects and Limits*, pp. 195–217. (Cambridge University Press, New York).

Camerer, C., 1995,. Individual decision making, . In J. H. Kagel and A. E. Roth, eds., *The Handbook of Experimental Economics*, Chapter 8, pp. 587–703. (Princeton University Press, Princeton, NJ).

Campbell, J. Y. and J. H. Cochrane, 1999, By force of habit: A consumption based explanation of aggregate stock market behavior, *Journal of Political Economy* **107**, 205–251.

Campbell, J. Y. and A. S. Kyle, 1993, Smart money, noise trading, and stock price behaviour, *Review of Economic Studies* **60**, 1–34.

Carhart, M. M., 1997, On persistence in mutual fund performance, *Journal of Finance* **52**, 57–82.

Chacko, G., B. Chowdhry, R. Cohen, and J. Coval, 2002, Rational exuberance, HBS working paper.

Chambers, A. E. and S. H. Penman, 1984, Timeliness of reporting and the stock price reaction to earnings announcements, *Journal of Accounting Research* **22**, 21–47.

Chan, L. K., J. Lakonishok, and T. Sougiannis, 2001, The stock market valuation of research and development expenditures, *Journal of Finance* **56**, 2431–2456.

Chen, J., 2003, Intertemporal CAPM and the cross-section of stock returns, EFA 2002 Berlin Meetings Discussion Paper.

Choi, J. J., D. Laibson, and A. Metrick, 2002, How does the internet affect trading? Evidence from investor behavior in 401(k) plans, *Journal of Financial Economics* **64**, 397–421.

Coffee, Jr., J. C., 1984, Market failure and the economic case for a mandatory disclosure system, *Virginia Law Review* **70**, 717–753.

Cooper, J. and R. H. Fazio, 1984,. A new look at dissonance theory, . In L. Berkowitz, ed., *Advances in Experimental Social Psychology*, Volume 17, pp. 229–264. (Academic Press, Orlando, FL).

Coval, J. D., D. Hirshleifer, and T. Shumway, 2003, Can individual investors beat the market?, Harvard NOM Working Paper No. 02-45.

Coval, J., D. Hirshleifer, and A. Thakor, 2004, Exploitative intermediaries, in progress, Harvard Business School, University of Michigan and Ohio State University.

Cutler, D. M., J. M. Poterba, and L. H. Summers, 1990, Speculative dynamics and the role of feedback traders, *American Economic Review* **80**, 63–68.

Daniel, K. D., M. Grinblatt, S. Titman, and R. Wermers, 1997, Measuring mutual fund performance with characteristic-based benchmarks, *Journal of Finance* **52**, 1035–1058.

Daniel, K. D., D. Hirshleifer, and A. Subrahmanyam, 1998, Investor psychology and security market under- and over-reactions, *Journal of Finance* **53**, 1839–1886.

Daniel, K. D., D. Hirshleifer, and A. Subrahmanyam, 2001, Overconfidence, arbitrage, and equilibrium asset pricing, *Journal of Finance* **56**, 921–965.

Darrough, M. and N. Stoughton, 1990, Financial disclosure policy in an entry game, *Journal of Accounting and Economics* **12**, 219–244.

DeBondt, W. F. M. and R. H. Thaler, 1995,. Financial decision-making in markets and firms: A behavioral perspective, . In R. A. Jarrow, V. Maksimovic, and W. T. Ziemba, eds., *Finance, Handbooks in Operations Research and Management Science*, Volume 9, Chapter 13, pp. 385–410. (North Holland, Amsterdam).

DeLong, J. B., A. Shleifer, L. Summers, and R. J. Waldmann, 1990, Noise trader risk in financial markets, *Journal of Political Economy* **98**, 703–738.

DeLong, J. B., A. Shleifer, L. Summers, and R. J. Waldmann, 1990, Positive feedback investment strategies and destabilizing rational speculation, *Journal of Finance* **45**, 375–395.

DeLong, J. B., A. Shleifer, L. Summers, and R. J. Waldmann, 1991, The survival of noise traders in financial markets, *Journal of Business* **64**, 1–20.

Dong, M., D. Hirshleifer, S. Richardson, and S. H. Teoh, 2003, Does investor misvaluation drive the takeover market?, Working paper, Ohio State University, York University and Wharton School.

Dye, R. A., 1985, Disclosure of nonproprietary information, *Journal of Accounting Research* **23**, 123–145.

Einhorn, H. J. and R. Hogarth, 1978, Confidence in judgment: Persistence in the illusion of validity, *Psychological Review* **85**, 395–416.

Einhorn, H. J., 1980, Overconfidence in judgment, *New Directions for Methodology of Social and Behavioral Science* **4**, 1–16.

Engel, C., 1996, The forward discount anomaly and the risk premium: A survey of recent evidence, *Journal of Empirical Finance* **3**, 123–192.

Fama, E. F. and K. R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* **33**, 3–56.

Fama, E. F., 1970, Efficient capital markets: A review of theory and empirical work, *Journal of Finance* **25**, 383–417.

Fama, E. F., 1998, Market efficiency, long-term returns and behavioral finance, *Journal of Financial Economics* **49**, 283–306.

Feltham, G. A. and J. Xie, 1992, Voluntary financial disclosure in an entry game with continua of types, *Contemporary Accounting Research* **9**, 46–80.

Felton, J. and J. Kim, 2002, Warnings from the enron message board, *Journal of Investing* **11**, 29 – 52.

Festinger, L. and J. M. Carlsmith, 1959, Cognitive consequences of forced compliance, *Journal of Abnormal and Social Psychology* **58**, 203–211.

Fischoff, B., P. Slovic, and S. Lichtenstein, 1977, Knowing with certainty: the appropriateness of extreme confidence, *Journal of Experimental Psychology* **3**, 552–564.

Fischoff, B., 1982,. For those condemned to study the past: Heuristics and biases in hindsight, . In D. Kahneman, P. Slovic, and A. Tversky, eds., *Judgement under Uncertainty: Heuristics and Biases*. (Cambridge University Press, Cambridge).

Forsythe, R., F. Nelson, G. Neumann, and J. Wright, 1992, Anatomy of an experimental political stock market, *American Economic Review* **82**, 1142–1161.

Frankel, J. A. and K. A. Froot, 1986, Understanding the US dollar in the eighties: The expectations of chartists and fundamentalists, *Economic Record, Supp.* **62**, 24–38.

Frankel, J. A. and K. A. Froot, 1990, Chartists, fundamentalists, and trading in the foreign exchange market, *American Economic Review* **80**, 181–185.

Gervais, S. and T. Odean, 2001, Learning to be overconfident, *Review of Financial Studies* **14**, 1–27.

Gigerenzer, G., U. Hoffrage, and H. Kleinbolting, 1991, Probabilistic mental models: A brunswikian theory of confidence, *Psychological Review* **98**, 506–528.

Gilovich, T. R., 1991, *How We Know What Isn't So: The Fallibility of Human Reason in Everyday Life* (Free Press, New York).

Grether, D. M., 1992, Testing bayes rule and the representativeness heuristic: Some experimental evidence, *Journal of Economic Behavior and Organization* **17**, 31–57.

Griffin, D. and A. Tversky, 1992, The weighing of evidence and the determinants of overconfidence, *Cognitive Psychology* **24**, 411–435.

Grinblatt, M. and B. Han, 2003, Disposition effect and momentum profit, Working paper, UCLA and The Ohio State University.

Grinblatt, M. and S. Titman, 1992, The persistence of mutual fund performance, *Journal of Finance* **47**, 1977–1984.

Grossman, S. J., 1981, The informational role of warranties and private disclosure about product quality, *Journal of Law and Economics* **24**, 461–483.

Hanson, J. D. and D. A. Kysar, 1999, Taking behavioralism seriously: Some evidence of market manipulation, *Harvard Law Review* **112**, 1420–1572.

Harmon-Jones, E. and J. Mills, 1999,. An introduction to cognitive dissonance theory and an overview of current perspectives on the theory, . In E. Harmon-Jones and J. Mills, eds., *Cognitive Dissonance: Progress on a Pivotal Theory in Social Psychology.* (American Psychological Association, Washington, DC).

Haw, I., D. Qi, and W. Wu, 2000, Timeliness of annual report releases and market reaction to earnings announcements in an emerging capital market: The case of China, *Journal of International Financial Management and Accounting* **11**, 108–131.

Hertwig, R. and A. Ortmann, 2001, Experimental practices in economics: A challenge for psychologists?, *Behavioral and Brain Sciences* **24**, 383–403.

Hirshleifer, D., S. S. Lim, and S. H. Teoh, 2003, Disclosure to a credulous audience: The role of limited attention, Fisher College of Business, Ohio State University.

Hirshleifer, D. and G. Y. Luo, 2001, On the survival of overconfident traders in a competitive security market, *Journal of Financial Markets* **4**, 73–84.

Hirshleifer, D. and S. H. Teoh, 2003, Herd behavior and cascading in capital markets: A review and synthesis, *European Financial Management* **9**, 25–66.

Hirshleifer, D. and Y. Wang, 2001, Investor overconfidence and the forward discount puzzle, Fisher College of Business, Ohio State Univeristy.

Hirshleifer, D. and I. Welch, 2002, An economic approach to the psychology of change: Amnesia, inertia, and impulsiveness, *Journal of Economics and Management Strategy* **11**, 379–421.

Hong, H. and J. C. Stein, 1999, A unified theory of underreaction, momentum trading and overreaction in asset markets, *Journal of Finance* **54**, 2143–2184.

Isenberg, D. J., 1986, Group polarization: A critical review and meta-analysis, *Journal of Personality and Social Psychology* **50**, 1141–1151.

Jensen, M. C., 1968, The performance of mutual funds in the period 1945-1964, *Journal of Finance* **23**, 389–416.

Jensen, M. C., 1969, Risk, the pricing of capital assets, and the evaluation of investment portfolios, *Journal of Business* **42**, 167–247.

Josephs, R., R. P. Larrick, C. M. Steele, and R. E. Nisbett, 1996, Protecting the self from the negative consequences of risky decisions, *Journal of Personality and Social Psychology* **62**, 26–37.

Kahneman, D., J. L. Knetsch, and R. H. Thaler, 1991, The endowment effect, loss aversion, and status quo bias: Anomalies, *Journal of Economic Perspectives* **5**, 193–206.

Kahneman, D. and A. Tversky, 1973, On the psychology of prediction, *Psychological Review* **80**, 237–25l.

Kahneman, D. and A. Tversky, 1982,, The psychology of preferences,, *Scientific American* **246**, 160–173.

Keane, M. P. and D. E. Runkle, 1990, Testing the rationality of price forecasts: New evidence from panel data, *American Economic Review* **80**, 714–735.

Keren, G., 1991, Calibration and probability judgments: Conceptual and methodological issues, *Acta Psychologica* **77**, 217–273.

Knez, P., V. L. Smith, and A. Williams, 1985, Individual rationality, market rationality, and value estimation, *American Economic Review Papers and Proceedings* **52**, 397–402.

Koehler, J. J., 1996, The base rate fallacy reconsidered: Normative, descriptive and methodological challenges, *Behavioral and Brain Sciences* **19**, 1–53.

Kyle, A. S., 1985, Continuous auctions and insider trading, *Econometrica* **53**, 1315–1335.

Kyle, A. and F. A. Wang, 1997, Speculation duopoly with agreement to disagree: Can overconfidence survive the market test?, *Journal of Finance* **52**, 2073–2090.

Lamoureux, C. and C. Schnitzlein, 2001, Mechanism-dependent behavior: Overconfidence in experimental asset markets, Working paper, College of Busineess, University of Arizona.

Langer, E. J. and J. Roth, 1975, Heads I win tails it's chance: The illusion of control as a function of the sequence of outcomes in a purely chance task, *Journal of Personality and Social Psychology* **32**, 951–955.

Lee, C., A. Shleifer, and R. Thaler, 1991, Investor sentiment and the closed-end fund puzzle, *Journal of Finance* **46**, 75–109.

Lichtenstein, S., B. Fischoff, and L. Phillips, 1982,. Calibration of probabilities: The state of the art to 1980, . In D. Kahneman, P. Slovic, and A. Tversky, eds., *Judgement under Uncertainty: Heuristics and Biases*, pp. 306–334. (Cambridge University Press, Cambridge).

Lichtenstein, S. and B. Fischoff, 1977, Do those who know more also know more about how much they know?, *Organizational Behavior and Human Performance* **20**, 159–183.

Loomes, G. and R. Sugden, 1982, Regret theory: An alternative theory of rational choice under uncertainty, *Economic Journal* **92**, 805–824.

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* **37**, 2098–2109.

Lovell, M. C., 1986, Tests of the rational expectations hypothesis, *American Economic Review* **76**, 110–124.

MacKinlay, A. C., 1995, Multifactor models do not explain deviations from the CAPM, *Journal of Financial Economics* **38**, 3–28.

Malkiel, B. G., 1995, Returns from investing in equity mutual funds 1971 to 1991, *Journal of Finance* **50**, 549–572.

Mathios, A. D., 2000, The impact of mandatory disclosure laws on product choices: An analysis of the salad dressing market, *Journal of Law and Economics* **43**, 651–677.

McClelland, A. G. and F. Bolger, 1994,. The calibration of subjective probabilities: Theories and models 1980-1993, . In G. Wright and P. Ayton, eds., *Subjective Probability*, Chapter 18, pp. 453–482. (Wiley, Chichester; New York).

McNichols, M., 1989, Evidence of informational asymmetry from management earnings, *The Accounting Review* **64**, 1–27.

Milgrom, P. R., 1981, Good news and bad news: Representation theorems and applications, *Bell Journal of Economics* **12**, 380–391.

Miller, D. T. and M. Ross, 1975, Self-serving bias in attribution of causality: Fact or fiction?, *Psychological Bulletin* **82**, 213–225.

Nel, E., R. Helmreich, and E. Aronson, 1969, Opinion change in the advocate as a function of the persuasibility of his audience: A clarification of the meaning of dissonance, *Journal of Personality and Social Psychology* **12**, 117–124.

Odean, T., 1998, Volume, volatility, price and profit when all traders are above average, *Journal of Finance* **53**, 1887–1934.

Odean, T., 1999, Do investors trade too much?, *American Economic Review* **89**, 1279–1298.

Palmiter, A. R., 1999, Toward disclosure choice in securities offerings, *Columbia Business Law Review* **1**, 1–135.

Payne, J. W., J. R. Bettman, and E. J. Johnson, 1992, Behavioral decision research: A constructive processing perspective, *Annual Review of Psychology* **43**, 87–131.

Polk, C. and P. Sapienza, 2003, The real effects of investor sentiment, Working paper, Northwestern University.

Rabin, M., 1998, Psychology and economics, *Journal of Economic Literature* **36**, 11–46.

Richardson, S., S. H. Teoh, and P. D. Wysocki, 2004, The walkdown to beatable analysts' forecasts: The roles of equity issuance and insider trading incentives, *Contemporary Accounting Research, forthcoming*.

Ritov, I. and J. Baron, 1990, Reluctance to vaccinate: omission bias and ambiguity, *Journal of Behavioral Decision Making* **3**, 263–277.

Ritov, I., 1996, Probability of regret: Anticipation of uncertainty resolution in choice, *Organizational Behavior and Human Decision Processes* **66**, 228–236.

Rubinstein, M., 2001, Rational markets: Yes or no? the affirmative case, *Financial Analysts Journal* **57**, 15–29.

Shefrin, H. and M. Statman, 1984, Explaining investor preference for cash dividends, *Journal of Financial Economics* **13**, 253–82.

Shefrin, H. and M. Statman, 1985, The disposition to sell winners too early and ride losers too long: Theory and evidence, *Journal of Finance* **40**, 777–790.

Skinner, D., 1994, Why firms voluntarily disclose bad news, *Journal of Accounting Research* **32**, 38–60.

Sloan, R., 1996, Do stock prices fully reflect information in accruals and cash flows about future earnings?, *Accounting Review* **71**, 289–315.

Soll, J. B., 1996, Determinants of overconfidence and miscalibration: The roles of random error and ecological structure, *Organizational Behavior and Human Decision Processes* **65**, 117–137.

Spranca, M., E. Minsk, and J. Baron, 1991, Omission and comission in judgment and choice, *Journal of Experimental Social Psychology* **27**, 76–105.

Stanovich, K. E. and R. F. West, 2000, Individual differences in reasoning: Implications for the rationality debate?, *Brain and Behavioral Sciences* **22**, 645–665.

Statman, M. and S. Thorley, 1998, Overconfidence and trading volume, Santa Clara University.

Staw, B. M., 1976, Knee-deep in the big muddy: A study of escalating commitment to a chosen course of action, *Organizational Behavior and Human Performance* **16**, 27–44.

Steele, C. and T. Liu, 1983, Dissonance processes as self-affirmation, *Journal of Personality and Social Psychology* **45**, 393–397.

Swaminathan, B., 1996, Time-varying expected small firm returns and closed-end fund discounts, *Review of Financial Studies* **9**, 845–888.

Taylor, S. E. and J. D. Brown, 1988, Illusion and well-being: A social psychological perspective on mental health, *Psychological Bulletin* **103**, 193–210.

Teoh, S. H. and C.-Y. Hwang, 1991, Non-disclosure and adverse disclosure as signals of firm value, *Review of Financial Studies* **4**, 283–313.

Teoh, S. H., I. Welch, and T. J. Wong, 1998, Earnings management and the long-term market performance of initial public offerings, *Journal of Finance* **53**, 1935–1974.

Teoh, S. H., I. Welch, and T. J. Wong, 1998, Earnings management and the underperformance of seasoned equity offerings, *Journal of Financial Economics* **50**, 63–99.

Teoh, S. H., T. J. Wong, and G. Rao, 1998, Are accruals during an initial public offering opportunistic?, *Review of Accounting Studies* **3**, 175–208.

Teoh, S. H. and T. Wong, 2002, Why do new issues and high accrual firms underperform? The role of analysts' credulity, *Review of Financial Studies* **15**, 869–900.

Trivers, R., 1985, *Social Evolution* (Benjamin/Cummings, Menlo Park).

Trivers, R., 1991,. Deceit and self-deception, . In R. Robinson and L. Tiger, eds., *Man and Beast Revisited*. (Smithsonian Press, Washington D.C.).

Verrecchia, R. E., 1983, Discretionary disclosure, *Journal of Accounting and Economics* **5**, 179–194.

Verrecchia, R. E., 1990, Endogenous proprietary costs through firm interdependence, *Journal of Accounting and Economics* **12**, 245–250.

Wang, F. A., 2001, Overconfidence, investor sentiment, and evolution, *Journal of Financial Intermediation* **10**, 138–170.

Williams, A. W., 1987, The formation of price forecasts in experimental markets, *Journal of Money, Credit, and Banking* **19**, 1–18.

Yates, J. F., 1990, *Judgement and Decision Making* (Prentice Hall, Englewood Cliffs, New Jersey).