Who's the Winner? An Experimental Examination of Trader Performance

Lucy F. Ackert
Department of Economics and Finance
Michael J. Coles College of Business
Kennesaw State University
1000 Chastain Road
Kennesaw, GA 30144
770-423-6111
lackert@kennesaw.edu
and
Research Department
Federal Reserve Bank of Atlanta
1000 Peachtree Street N.E.
Atlanta, GA 30309-4470

Bryan K. Church
800 West Peachtree Street
College of Management
Georgia Tech
Atlanta, GA 30308-0520
404-894-3907
bryan.church@mgt.gatech.edu

Ping Zhang
Rotman School of Management
105 St. George Street
University of Toronto
Toronto, ON
Canada M5S 3E6
416-946-5655
pzhang@rotman.utoronto.ca

September 2009

The views expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. The authors thank Ann Gillette, Jeff Hales, Brian Kluger, Pete Kyle, and Steve Smith for helpful comments and gratefully acknowledge the financial support of the Federal Reserve Bank of Atlanta and Social Sciences and Humanities Research Council of Canada.

Who's the Winner? An Experimental Examination of Trader Performance

Abstract

This paper reports the results of 18 experimental asset markets designed to investigate the association between overconfidence and trader performance. Prior to trading, participants compete in a sealed-bid auction to acquire a forecast of asset value. The auction procedures are similar to those used in studies that look at the winner's curse and allow us to endogenously determine overconfident individuals each trading period – overconfident individuals are those willing to pay the most for private information. After competing for private information, participants trade a one-period asset in a double auction market. We find that participants pay too much for private information, which indicates that they overestimate the usefulness of private information. As a consequence, informed agents generally are unable to recover the cost of private information and, thus, fare poorly in our asset markets. Further analyses indicate that informed agents are reasonably well calibrated, but that they misjudge their trading skills - they believe that they are more skilled than others. Lastly, our findings suggest that informed agents behave strategically in our markets (to avoid tipping their hand), but such behavior does not permit them to recoup the cost of private information.

Who's the Winner? An Experimental Examination of Trader Performance

1. Introduction

Research has long recognized that people are often overconfident (e.g., Oskamp 1965; Fischhoff, Lichtenstein, and Slovic 1977). Overconfident individuals can be miscalibrated (overestimating the precision of their private information), believe they are more skilled than others ("better than average"), or may be overly optimistic about the future. Extant evidence suggests that overconfident individuals are a presence in securities markets (e.g., Barber and Odean 2000; 2001).

Prior studies that consider the association between overconfidence and trader performance produce mixed results. Some empirical findings suggest that overconfidence is detrimental to trader performance. Odean (1998) shows that with costly information and overconfidence, informed traders perform worse than uninformed traders. Informed traders overestimate the value of their signal so that "its benefits are spread too thin." Barber and Odean (2000) provide evidence that is indicative of overconfident individuals faring poorly in securities markets. The traditional view in finance is that such traders will disappear in the long run because of trading losses. Yet other studies conjecture that overconfident individuals can persist and prosper in securities markets. For instance, Benos' (1998) model shows that overconfident traders can earn excess profits because active trading allows them to make the most of a "first mover advantage." Hirshleifer and Luo (2001) assert that overconfident individuals are better able to exploit mispricing that arises due to noise/liquidity traders.

We design a series of experimental asset markets to investigate the association between overconfidence and trader performance. The asset markets are organized as

double oral auctions, with the supply of private information fixed per market. Prior to trading, participants submit a sealed bid indicating the price at which they are willing to acquire an imperfect forecast of asset value. Through these sealed bids, we are able to assess the extent of trader overconfidence (see also Allen and Evans 2005). Excessive bidding in common value auctions has been documented in many contexts, including corporate takeovers and oil leases, under the rubric of the winner's curse (Thaler 1992). Individuals who are more confident in the usefulness of private information will submit higher bids.

We manipulate the number of traders who acquire private information across experimental markets. In one-third of the markets, one trader acquires an imperfect forecast of asset value; in another one-third, two traders acquire the forecast; and in the remaining one-third, four traders acquire the forecast. All markets are comprised of seven or eight traders in total. By design, the winning bids should be highest in markets with one informed trader (a second-price auction), lower in markets with two informed traders (a third-price auction), and lowest in markets with four informed traders (a fifth-price auction). We proffer that the process of submitting a sealed bid and competing to acquire information impacts winners' assessment of the information. More specifically, winners are subject to a self-serving bias, which leads them to overestimate the usefulness of private information (e.g., Gervais and Odean 2001). The effect intensifies as the cost of winning increases and, in turn, the number of winners decreases. Research on the winner's curse suggests that bids are more extreme as the number of bidders increases – or as the proportion of winners decreases (e.g., Kagel 1995; Hong and Shum

2002). Accordingly, overconfidence is magnified as the number of informed traders (winners) decreases.

Notably our design allows for the endogenous determination of overconfidence. Gervais and Odean (2001) argue that "it is a common feature of human existence that we constantly learn about our own abilities by observing the consequences of our actions." Thus, a person's confidence level is not constant, neither over time nor across domains. Rather, the level of confidence changes depending on personal experience. In our experiment, participants bid for information. Those who submit the highest bids *win* the auction and purchase private information. The very act of winning the auction promotes overconfidence. While other experimental studies have examined the link between overconfidence and performance (Kirchler and Maciejovsky 2002; Biais, Hilton, Mazurier, and Pouget 2005), the endogenous determination of overconfidence is a distinguishing feature of our study.[1]

Another important feature is that we are able to delve into the nature of overconfidence. By conducting experimental asset markets, we can collect data to gain insight into *how* overconfidence manifests itself. We investigate whether informed traders overestimate (1) the quality of private information (its precision and accuracy) and/or (2) their trading skill (i.e., ability to prosper using private information). The former is referred to as miscalibration (overconfidence in an absolute sense) and the latter as "better than average" (overconfidence in a relative sense).

We elicit traders' assessment of asset value each period to identify whether informed traders are overconfident in the quality of private information. Using traders'

[1] In prior experimental studies on overconfidence, traders are endowed with private information so that they do not actively choose whether to become informed on a period-by-period basis.

predictions, we are able to determine whether their assessments are accurate and sufficiently precise (well calibrated). In addition, we examine whether traders are able to undo the bias contained in the forecast of asset value. We manipulate the forecast of asset value across experimental markets. In one-half of the markets the forecast is an unbiased estimate of asset value, and in the other half it is systematically biased upward. We focus on an upward bias because previous studies document that financial analysts' 6- to 12-month ahead earnings forecasts on average are overly optimistic (e.g., Barefield and Comiskey, 1975; O'Brien, 1988; Richardson, Teoh, and Wysocki, 2004).

We also collect data on informed traders' activity per period, including their quotes, trades, and profit. These data provide evidence of *how* overconfidence affects behavior in the double auction market setting. In our markets, informed traders (i.e., winners in the sealed-price auction) have the opportunity to generate excess profit by exploiting an informational advantage.[2] Prior experimental findings indicate that private information (perfect or imperfect) is not reflected immediately in asset price (e.g., Ackert, Church, and Shehata 1997). Rather the price-adjustment process is gradual. Hence, informed traders can make money by acquiring private information. Moreover, the ability to generate excess profit should increase as the number of informed traders decreases, in which case fewer traders compete to exploit the informational advantage. But overconfidence may have an offsetting effect on informed traders' welfare. As the number of informed traders decreases, the cost of information escalates due to overconfidence, making it more difficult to earn excess profit.

Empirical evidence suggests that overconfidence may be associated with excessive trading (e.g., Barber and Odean 2001; Glaser and Weber 2007; Grinblatt and

[2] Excess profit refers to making more than the cost of information.

Keloharju 2009). In our setting, however, informed agents may temper their trading behavior in an attempt to exploit the informational advantage. Kyle (1984; 1989) contends that informed traders behave strategically to avoid revealing private information.[3] We analyze whether informed agents trade aggressively or strategically. In either case, they run the risk of failing to recover the cost of private information. If informed agents trade aggressively, they may reveal their private information, eliminating the informational advantage. If they trade strategically, on the other hand, they may not engage in enough activity to generate a profit. Our design allows us to probe the association between overconfidence and trading activity and the resultant effect on performance.

Our results indicate that the median bid for private information is approximately three to four times greater with one or two informed traders than with four informed traders. The winning bid (i.e., price of information) is approximately five to six times greater with one or two informed agents! What is more, the observed differences do not diminish over time. Subsequent analyses suggest that informed traders' are reasonably well calibrated: their predictions of asset value are usually within the range of feasible values. Informed traders properly assess the accuracy of unbiased, private information. Their assessment of biased forecasts, however, is problematic. They appear to anchor on the biased forecasts such that their adjustment to undue bias is insufficient. Our data suggest that, as may be expected, informed traders *try to* exploit their informational advantage: the average informed trader engages in more activity (quotes and trades) than

[3] Hellwig (1980) and Kyle (1989) recognize the trader schizophrenia that can arise in competitive rational expectations models because traders are assumed to be price takers yet they influence prices through their trades. In our markets informed traders may affect prices through their trades and may also take this impact into consideration when formulating their trading strategy.

the average uninformed trader. We also find evidence of sufficient opportunity to generate excess profits, particularly in markets with one or two informed traders: period-end asset price does not fully reflect private information in these markets. Nevertheless, informed traders generally fare worse than uninformed traders across all markets. They pay too much to become informed and appear to be overconfident in their trading skill (i.e., ability to exploit their informational advantage).

The remainder of this paper is organized as follows. In section 2, we provide a framework to assess trader performance (informed versus uninformed). In section 3, we describe the experimental method. In section 4, we present the experimental results. Lastly, we offer concluding remarks and discuss the implications of our findings for empirical research.

2. Framework

2.1. Bidding, acquiring information, and overconfidence

When individuals compete for a good (information in our experimental setting), the winning bidder invariably overpays, commonly referred to as the winner's curse (e.g., Bazerman and Samuelson 1983; Thaler 1992). By definition, the winner values the good more than others, paying the most for it. But, chances are the winner has overestimated its value (i.e., pays more than the expected value) and is overconfident in its usefulness. Experimental and field evidence suggest that the winner's curse is an often observed and robust phenomenon (e.g., Thaler 1992; Grosskopf, Bereby-Meyer, and Bazerman 2007). What is more, the effect is magnified as competition increases. With more participants,

individuals have to bid more aggressively to win, but such behavior increases the chances of overbidding (i.e., the winning bid is more likely to be too extreme).

Ku, Malhotra, and Murnighan (2005) suggest that overbidding occurs because of competitive arousal, which can arise due to individuals' desire to win and their rivalry with other participants. Simply put, bidders get caught up in the moment. Winning bidders may exaggerate value to justify extreme bids, as such behavior allows them to maintain cognitive consistency and preserve self concept (e.g., Festinger 1957; Aronson 1969; Baumeister 1998). But the process can feed itself iteratively, leading to increased overconfidence in the good's value. For our purposes, winning the auction and acquiring private information can reflect overconfidence. Winning bidders reason that private information is well worth its price and will allow them to thrive.

We contend that the extent of individuals' overconfidence is associated with the number of winning bidders. In our setting, the competitive nature of the bidding process is a function of the supply of information. As the number of winners decreases (supply falls), competitiveness increases. With fewer winners, competitive arousal is intensified due to the uniqueness of winning. With heightened arousal, bidders engage in more risk taking (Mano 1994) and less deliberate processing of information (Lewinsohn and Mano 1993) and, as a consequence, the winning bid can be quite extreme.

The process underlying overconfidence is characterized as follows. With fewer informed traders (i.e., winners), the price of information is determined by a more extreme bid: traders acquiring private information pay a steeper price. In this case, the winners are more likely to overvalue private information, meaning that they are overconfident. The mere act of winning the auction also may lead to elation, fostering overconfidence.

In other words, the winner(s) become even *more* overconfident in the usefulness of private information.

The effects described above are likely strongest with one or two informed traders and dampened, to some extent, with four informed traders. With one or two informed traders, the price of information will be more extreme, likely reflecting overconfidence. Traders recognize that the potential benefits of becoming informed increase as the number of winners decrease. But the uniqueness of winning the auction (one or two winners versus four winners) fuels more extreme bids, which promotes overconfidence.

2.2. Trader behavior and overconfidence

Winning bidders may be overconfident in the quality of private information and/or their trading skill. The former suggests that informed traders overestimate the closeness of the forecast and asset realization (accuracy and/or precision). The latter suggests that informed traders overestimate their ability to exploit the informational advantage. In either case, informed traders are overconfident in the usefulness of private information.

2.2.1. Quality of private information. If informed traders misjudge the quality of private information, they may have difficulty using the information to predict asset value. With unbiased private information, informed traders may overestimate the information's precision, which indicates miscalibration. Informed traders may overrate the diagnosticity of private information and underrate the conditional uncertainty surrounding asset value, which can lead to unprofitable trades. Theoretical results suggest that miscalibration leads to poor performance (e.g., Benos 1998; Daniel,

Hirshleifer, and Subrahmanyam 1998; Odean 1998). Experimental findings also provide evidence that miscalibrated traders fare worse than others (Biais, Hilton, Mazurier, and Pouget 2005).

In addition to miscalibration, informed traders may face further difficulty if the forecast of asset value is biased. Now, they must adjust the forecast of asset value to undo bias. Although informed traders may recognize bias (e.g., that the forecast of asset value, on average, is greater than asset realization), their adjustment may be insufficient. Instead, they may *anchor* on the forecast when assessing asset value. Much research suggests that individuals fail to adjust adequately away from an anchor, even when they know that the anchor is biased or manipulative (e.g., Hastie, Schkade, and Payne 1999; Galinsky and Mussweiler 2001). Accordingly, informed traders may rely too heavily on the forecast of asset value and systematically err in assessing value, which can negatively impact their trading profit.

2.2.2. Trading skill. Prior research suggests that individuals often have exaggerated beliefs of their ability, viewing themselves as *above average* (e.g., Svenson 1981; Taylor and Brown 1988; Odean 1998). If informed traders overestimate their trading skill, they may find it difficult to exploit their informational advantage. In our setting, market participants have common knowledge that a subset of the market is informed.[4] Uninformed traders know that they are disadvantaged and, thus, may behave cautiously.[5] They may be keenly attuned to others' bid-ask behavior as well as price

[4] In a related study, Schnitzlein (2002) examines trader behavior when traders are aware of the presence of insiders. In his markets information arrives randomly and traders must infer the number of informed traders. Our design differs in that participants know that there is a market for information and all are free to participate.

[5] As explained later, traders have an incentive to transact in our markets. Some traders pay a tax on dividend earnings (but not on capital gains), whereas others are not taxed. As such, the asset is worth more to some traders (i.e., those not taxed) than others.

signals. Therefore, informed traders are confronted with a tricky problem. They must transact to recoup the cost of information, but their bid-ask behavior may reveal private information. Informed traders may seek to behave strategically so as not to tip their hand (e.g., Kyle 1984; 1989). But their ability to do so is affected by the number of informed agents (i.e., the competitiveness of exploiting the informational advantage).

In markets with one informed trader, the opportunity to exploit the informational advantage is greatest. The informed trader can behave strategically and not worry about competition from others. But, in this case, the price of information is likely very extreme – overconfidence implies that the information is overvalued. The informed trader faces the challenge of engaging in sufficient activity (profitable trades) to recover the cost of information.

In markets with two informed traders, competition is introduced: the two informed traders compete with one another to exploit the informational advantage. Though they may behave strategically, their opportunities to prosper are reduced (because two agents are informed instead of one). Worse, the price of information may still be extreme. Again, the challenge for informed traders is to engage in enough activity to recover the cost information.

In markets with four informed traders, competition is intensified. Informed traders may have little opportunity to trade strategically because the situation is so competitive – at least one-half of the market participants are informed. Under such conditions, private information is likely to be fully disseminated (Sunder 1992). Informed traders may transact more hurriedly, compared to other market conditions, in an

effort to generate profit before their informational advantage disappears (i.e., before private information is fully disseminated).

We examine traders' welfare across market conditions. If informed traders are overconfident in their ability to generate profit, they may fare poorly. Specifically, they may not recoup the cost of private information and they may perform worse than uninformed traders. To gain additional insight, we investigate informed traders' activity (the extent to which they seek to exploit their informational advantage) and asset price at period end (the extent to which private information is disseminated in the marketplace).

3. Experimental method

3.1. Overview

Eighteen experimental asset markets are conducted in which a forecast of asset value (i.e., the period-end dividend) is sold to a subset of the market. We manipulate the number of traders that is allowed to acquire the forecast. In six markets, only one trader is permitted to buy the forecast, in six others two traders buy the forecast, and in the remaining markets four traders buy the forecast. Participants submit bids indicating the price they are willing to pay for the forecast and those with the highest bids acquire the information (described below). We also manipulate the nature of the forecast. In one-half of the markets, the forecast is an unbiased estimate of asset value. In the other half, the forecast is systematically biased upward.

3.2. Procedures

We recruited 139 students from a large university to participate in our experimental asset markets. Students were third and fourth-year undergraduates, primarily in business and economics. Each market included seven or eight student traders.[6] Students earned from $12.30 to $66.65 dollars, with an average of $35.81, for participating approximately two hours.

At the beginning of each session, participants receive a hard copy of the instructions and follow along as the instructions are read aloud by the experimenter. These instructions are included in the Appendix. Each market consists of 12 periods and participants are not informed beforehand of the number of periods. During each period participants trade certificates with one-period lives. A dividend is received for each certificate held at period end. Participants are instructed that the period-end dividend is determined by drawing from a discrete, bell-shaped distribution (refer to Figure 1). The period-end dividend ranges from $300 to $2,700, in $100 increments, with a mean of $1,500. The dividend draws were conducted prior to administering the experimental sessions and the same sequence was used across markets. The primary benefit of using a pre-selected sequence is that it facilitates comparisons across markets conducted under similar as well as different experimental conditions.[7]

Each period, participants are endowed with two certificates and $50,000. Participants are informed that they may have to pay a tax on their dividend earnings, but not on capital gains. The tax rate on dividend earnings can be either zero or 20 percent. The instructions indicate that the tax rates differ across traders and across periods. At the

[6] Seven markets include seven traders and 11 markets include eight traders.

[7] See Cason and Friedman (1996) for further discussion of the merits of using a pre-selected sequence.

beginning of each period, three or four participants are assigned each tax rate and across the 12 periods each participant is assigned each rate the same number of times. We introduce different tax rates to create incentives to trade. The use of different tax rates is comparable to varying traders' state-dependent dividend preferences (see Cason and Friedman 1996, note 5). In general, participants with a zero tax rate have incentives to buy certificates, whereas those with a 20 percent tax rate have incentives to sell. The certificates are worth more to those with a zero tax rate on dividend earnings.

Before trading commences, participants submit a sealed bid indicating the price at which they are willing to acquire a forecast of the period-end dividend. Participants are informed that the process of generating the forecast is unique and constant across periods. To allow them to assess the usefulness of the forecast, a forecast history, collected over 10 practice periods, is provided. The forecast history includes the forecast, period-end dividend, and forecast error. The unbiased forecast is drawn from a discrete distribution that includes the dividend and its four nearest neighbors. The distribution, shown below, was not disclosed to participants.[8]

Forecast	Probability
Dividend - $200	0.14
Dividend - $100	0.23
Dividend	0.26
Dividend + $100	0.23
Dividend + $200	0.14

The biased is forecast is constructed by adding a constant of $200 to the unbiased forecast.

An auction determines the participant(s) acquiring the forecast and the number of available forecasts varies across markets. In six markets, one participant acquires the

[8] The observed period-end dividend draws ranged from $800 to $2,400. Thus, the four nearest neighbors always included two values above the dividend and two below it.

forecast at the second highest bid (second-price auction). In another six markets, two participants acquire the forecast at the third highest bid (third-price auction). In the remaining six markets, four participants acquire the forecast at the fifth highest bid (fifth-price auction). Similar auction procedures have been used elsewhere (e.g., Sunder 1992).

Logistically, participants are provided with an index card on which they submit a bid to acquire the forecast. The cards are collected by the experimenter, who determines the price of the forecast. The forecast is handwritten on the card(s) of the winning bidder(s) and "00" is handwritten on the other cards. The cards are returned to participants and the price of the forecast is publicly announced. The identity of the participant(s) acquiring the forecast is never revealed.

The 18 experimental markets are organized as oral, double auctions. Each period, traders are free to make verbal offers to buy or sell one certificate at a designated price and all offers are publicly announced and recorded. Outstanding offers stand until accepted or replaced by a better bid or ask price, and we do not allow short sales. Each market period lasts four minutes.

Just after trading ends for a period, participants are asked to record a prediction of the period-end dividend (asset value). Participants are awarded $0.25 for recording a prediction and this amount is not available to finance trading in future periods. The prediction represents participants' assessment of dividend (asset) value after observing trading for the period. Subsequently, the dividend value is announced and the same dividend is received for all certificates held by a participant.

We elicit predictions at period end for two reasons. First, uninformed traders do not have a basis to predict asset value before trading commences. Second, in markets

with more than one informed agent, informed traders can observe the behavior of others (which includes at least one other informed agent) and, in turn, further reflect on the accuracy of private information.

Period-end cash balances are computed as follows. The number of certificates on hand is multiplied by the dividend per certificate to determine dividend earnings. This amount is converted to an after-tax figure by multiplying by one minus the tax rate. Participants add the after-tax dividend earnings to their cash balance and then subtract the initial endowment of $50,000. The net amount represents participants' profits for the period (the $0.25 for recording a dividend prediction is in addition to this amount). Endowments are reinitialized at the beginning of the next period.

The instructions indicate that participants will be paid 0.1 percent of their total after-tax profit in cash. At the conclusion of each session, participants compute the amount of cash to be received and complete a post-experiment questionnaire.

4. Results

We investigate participants' behavior to assess whether it is indicative of overconfidence, be it in the quality of private information and/or their trading ability. First, we scrutinize the price of private information, along with the bidding process, to gauge the potential for overconfidence in our experimental markets. Second, we examine traders' assessment of private information to ascertain whether they correctly judge the quality of information. We test whether informed traders' predictions are well calibrated and accurate. Lastly, we look at traders' performance to determine if it pays to acquire private information. We compare the gross and net profit of informed and uninformed

traders. We also analyze information dissemination and participants' trading behavior, both of which impact traders' ability to prosper.

4.1. Acquiring private information

We peruse the data from the sealed-bid auction to obtain an understanding of the price of information across the experimental conditions. We are interested in how the number of winning bidders affects information price. We also look at the median bid per period, across markets, to gain insight into the bidding process. Finally, we examine the winning bidders to determine if private information is acquired by those who can benefit the most by being informed.

4.1.1. Price of information. Figure 2 provides a bar chart that depicts the average information price per period by experimental condition. Panel A shows the average price for markets with one, two, and four informed traders in markets with an unbiased forecast of asset value. Panel B shows the same data for markets with a biased forecast of asset value. The data suggest that price drops somewhat in the early periods, after which it remains fairly constant. As may be expected, the price is much lower in markets with four informed traders than in other markets. In addition, price does not appear to vary much between markets with unbiased and biased forecasts.

To conduct formal statistical tests, we perform a linear mixed model analysis using maximum likelihood estimation (see Greene, 1997, Ch. 14). The approach expands the general linear model by allowing the data to exhibit correlated and non-constant variability. We use an AR(1) covariance structure in the analysis, though inferences are similar using alternative specifications. The independent variables include the number of

informed traders (one, two, or four), forecast bias (unbiased or biased), and an interaction term. The dependent variable is the price of information per period repeated over time (i.e., period is a repeated measure). In all analyses, we omit periods 1-3 to allow participants time to familiarize themselves with the experimental procedures.[9]

The results indicate that the number of informed traders is statistically significant ($F = 11.52$, $p = 0.001$), whereas forecasts bias ($F = 0.02$, $p = 0.896$) and the interaction term ($F = 0.15$, $p = 0.860$) are insignificant. The average price of information is 1,122 in markets with one informed trader, 1,330 with two informed traders, and 232 with four informed traders. Bonferroni pair-wise comparisons indicate that the price is significantly lower in markets with four informed traders than in the other markets ($p < 0.015$). By comparison, price does not differ between markets with one and two informed traders ($p > 0.50$). Accordingly, the price of information is much more extreme with one or two winners – roughly five to six times more.

To provide a reference point for interpreting the price of information, we compute the expected incremental value of being informed assuming that asset price does not reflect the forecast: this assumption is more likely to hold for transactions that occur early in a period. A similar approach has been used by others (e.g., Copeland and Friedman 1992; Ackert, Church, and Shehata 1997). In determining the expected value, we assume that informed traders sell their certificates if the unbiased forecast is below the mean of the dividend distribution and buy if it is above the mean. The mean represents the uninformed or prior price expectation. Uninformed traders may be willing, at least

[9] In some cases, the dependent variable is quite extreme in periods 1-3, which may potentially distort the results. Inspection of the data indicates that the price of information and period-end asset price are far more extreme in periods 1-3 than in other periods.

initially, to buy at prices below this amount and sell at prices above this amount. The expected profit to an informed trader, of engaging in one transaction, is

$$E(\pi) = \sum_{D<1,500} [(\mu - D_i) * \rho_i] + \sum_{D>1,500} \{[(D_i * (1 - \tau) - \mu] * \rho_i\} \quad (1)$$

where μ is the mean of the dividend distribution ($1,500), D_i refers to period-end dividend values, ρ_i refers to the probabilities corresponding to D_i, and τ is agents' tax rate on dividend earnings. Briefly, the first term in (1) represents the expected profit from selling a certificate. The second term represents the expected profit from buying a certificate.

Substituting in the experimental parameters, we find that $E(\pi)$, assuming informed traders engage in one transaction, is $280 for traders with a zero tax rate. The $E(\pi)$ is $208 for traders with a 20 percent tax rate. Based on the average information price (ranging from 232 to 1,330), informed traders with a zero tax rate must transact from one to five times, without any information dissemination, to recover their cost. Those with a 20 percent tax rate must transact one to seven times. Informed traders' ability to recoup cost is directly affected by the speed of information dissemination and uninformed traders' willingness to transact (e.g., uninformed traders may require a premium to trade because they are disadvantaged). We return to this issue later in the results section.

4.1.2. Median bid. To gain insight into the bidding process, which can propel overconfidence, we examine the median bid per period by experimental condition. Once again we perform a linear mixed model analysis. We use the same independent variables as before and the median bid per period as the dependent variable. We find that the number of informed traders is statistically significant ($F = 7.84$, $p = 0.003$), whereas forecast bias and the interaction term are insignificant. The average median bid is 819 in

markets with one informed trader, 1,040 with two informed traders, and 272 with four informed traders.[10] Bonferroni pair-wise comparisons indicate that the median bid is significantly lower in markets with four informed traders than in markets with one ($p = 0.053$) or two ($p = 0.003$) informed traders. The median bid does not differ significantly between markets with one and two informed traders ($p > 0.60$). The findings suggest that participants, on the whole, bid more aggressively (higher) with fewer winners. In turn, the potential for overconfidence is magnified in markets with one or two informed traders as compared to markets with four informed traders.

4.1.3. Winning bidders. Next, we investigate whether winning and non-winning bids are associated with participants' tax rate. As described earlier, participants are assessed a tax rate of zero or 20 percent on dividend earnings (approximately one-half with each rate per period). Because short sales are not allowed, private information is potentially more beneficial to those with a zero tax rate. As such, traders with a zero tax rate should be more likely to submit higher bids and, in turn, win the sealed-price auction (i.e., become informed).

Across the six experimental conditions (looking at periods 4-12), informed traders have a zero tax rate 54 percent of the time. A binomial test indicates that informed traders are no more likely to have a zero tax rate than a 20 percent tax rate ($p = 0.151$). For each experimental condition, we perform a chi-square test to determine whether the breakdown of informed and uninformed traders per market is independent of tax rate. In five of six conditions (not tabulated), the chi-square tests provide evidence of independence ($p > 0.10$). The one exception is markets with four informed traders and

[10] With four informed traders, the average median bid (272) is more than the average information price (232). The reason is that in two of six of these markets, the market is comprised of only seven participants (rather than eight).

unbiased forecasts, in which case informed traders are more likely to have a zero tax rate and uninformed traders a 20 percent tax rate ($\chi^2 = 5.88$, $p = 0.015$). Competitive arousal appears to compel participants to bid more aggressively, particularly in markets with one or two auction winners, regardless of tax rate (e.g., Ku, Malhotra, and Murnighan 2005). The fact that participants with a 20 percent tax rate win the sealed-price auction (and acquire private information) is consistent with them being overconfident in the usefulness of the information.

4.2. Participants' assessment of private information

We examine whether participants misjudge the quality of private information. We consider two facets: whether informed traders' predictions are calibrated (within the range of feasible predictions) and whether their predictions are accurate (i.e., mean zero prediction error).

4.2.1. Calibration. Based on the experimental parameters, the unbiased forecast of asset value provides a point estimate within (±) 200 of the actual value. We examine whether informed traders' predictions of asset value (elicited at period end) fall within the interval [unbiased forecast – 200, unbiased forecast + 200], which provides evidence of being calibrated. Looking at the nine markets with an unbiased forecast, we find that informed traders' predictions fall within the specified interval 86 percent of the time. The breakdown by the number of informed traders per market is shown in Panel A of Table 1. The data suggest that participants in markets with one or two informed traders are well calibrated (over 90 percent of the time). In contrast, miscalibration appears to be more prevalent in markets with four informed traders. Further inspection of the data,

however, suggests that this result is attributable to a small subset of participants. Three participants, all in markets with four informed traders, are miscalibrated 16 of 23 periods. Aside from these three, other participants in markets with four informed traders are calibrated 80 of 85 periods (94 percent of the time).

To gain added insight, we investigate the frequency that participants acquire the forecast of asset value and the corresponding number of times that their predictions fall within the interval [unbiased forecast – 200, unbiased forecast + 200]. We find that 45 traders acquire the unbiased forecast at least once over periods 4-12. Panel B of Table 1 presents the breakdown of forecast acquisitions by the number of times that predictions are calibrated or miscalibrated (once, twice, and more than twice). We find that 33 of 45 traders (73 percent) are perfectly calibrated: their predictions are within ± 200 of asset value 119 of 119 periods. Seven traders are miscalibrated once (their predictions are calibrated 36 of 47 periods), and two traders are miscalibrated twice (their predictions are calibrated 6 of 10 periods). Overall, traders who acquire the unbiased forecast appear to be well calibrated.

Next, we turn to markets with biased forecasts. In these markets, informed traders must undue forecast bias to make predictions that are calibrated. The forecast bias is a constant of 200. So, informed traders must subtract 200 from the (biased) forecast of asset value to produce an unbiased estimate. As before, informed traders' predictions are calibrated if the predictions are within (±) 200 of the unbiased forecast. We find that predictions fall within the specified interval 80 percent of the time. The breakdown by the number of informed traders per market is shown in Panel A of Table 2. The data suggest that participants in markets with one informed trader are well calibrated (85

percent of the time), whereas miscalibration occurs more often in markets with two or four informed traders. Further examination indicates that the majority of miscalibration is due to four participants (one in markets with two informed traders and three in markets with four informed traders). These four participants are miscalibrated 20 of 32 periods. Apart from these four, other participants in markets with two informed traders are calibrated 35 of 45 periods (78 percent of the time) and others in markets with four informed traders are calibrated 80 of 85 periods (94 percent of the time).

As before, we examine the frequency that participants acquire the forecast of asset value. We identify 44 traders who acquire the biased forecast at least once over periods 4-12. Panel B of Table 2 gives the breakdown of forecast acquisitions by the number of times that predictions are calibrated/miscalibrated. Twenty-five of 44 traders (57 percent) are perfectly calibrated (93 of 93 periods). Eleven traders are miscalibrated once (their predictions are calibrated 32 of 43 period), and four traders are miscalibrated twice (their predictions are calibrated 11 of 19 periods). The data are indicative of slightly more miscalibration in markets with biased forecasts as compared to markets with unbiased forecasts. We shed additional light on this issue below.

4.2.2. Accuracy. We investigate informed traders' predictions of asset value to determine whether their predictions are systematically biased. First we examine informed traders' prediction errors in markets with unbiased forecasts. We compute a normalized prediction error as follows.

$$NPE_{it} = (P_{it} - UE_t)/ UE_t$$

where P_{it} is informed trader i's prediction of asset value in period t and UE_t is an unbiased estimate of asset value in period t conditioned on the unbiased forecast. In

computing UE_t, we must make an assumption as to whether informed traders know the forecast distribution (i.e., the forecast-generating process). If the distribution is not known, UE_t is the unbiased forecast per period. If it is known, a Bayesian adjustment can be used to compute UE_t. The unbiased estimate per period, computed using each approach, along with the period-end dividend (asset value), is shown in Table 3. As may be expected, the values are very close. In performing tests, we use the unbiased forecast, though results are unaffected using the Bayesian adjustment.

We compute NPE_{it} for informed traders in markets with an unbiased forecast and determine whether the mean is significantly different from zero. As shown in Panel A of Table 4, across the nine asset markets, the mean NPE (−0.0013) is not significantly different from zero ($p = 0.925$). Looking at each experimental condition though, we find that the mean NPE is negative (mean = −0.0291) and significantly different from zero ($p = 0.061$) in markets with one informed trader. The mean NPE is not significantly different from zero in the other two experimental conditions ($p > 0.50$).

Subsequently, we examine the number of times that informed traders' prediction error ($P_{it} - UE_t$) is negative, zero, and positive. Panel B of Table 4 indicates that the prediction error is negative more often than positive. More strikingly though, the prediction error is most often zero (i.e., informed traders' prediction exactly matches the unbiased forecast). This result is more pronounced in markets with one or two informed traders than in markets with four informed traders (56 percent and 69 percent versus 31 percent).[11] We suggest that informed traders have a reasonable assessment of the accuracy of unbiased forecasts.

[11] We perform separate chi-square tests and find that informed traders' prediction error are more likely to be zero in markets with one or two informed traders than in markets with four informed traders ($p \leq 0.015$).

Next, we look at informed traders' prediction errors in markets with biased forecasts. Panel A of Table 5 presents the average normalized prediction error (NPE) across the nine asset markets as well as separately for each experimental condition. In every case, the average normalized prediction error is positive and significantly different from zero ($p < 0.01$). We also inspect the number of times that informed traders' prediction error is negative, zero, and positive. Panel B of Table 5 indicates that the prediction error is positive a majority of the time.

To investigate further, we assess whether informed traders' prediction errors are less than 200 (the systematic bias) divided by UE_t. Such a finding is indicative of *some* adjustment for bias. The results of paired t-tests, shown in Panel C of Table 5, reveal that NPE is less than the unadjusted prediction error in every case at $p < 0.03$. Hence, informed traders' adjust biased forecasts downward, but the adjustment is insufficient.

4.2.3. Summary of informed traders' assessment. Our findings suggest that informed traders, for the most part, are well calibrated. The majority of participants who acquire private information give predictions of asset value that are within the range of feasible values: that is, within the interval [unbiased forecast – 200, unbiased forecast + 200]. We observe a few exceptions, but the number is small. We identify seven informed traders (out of 89) who are the source of much miscalibration, and six of them are from markets with four informed traders. This finding is noteworthy because we expect overconfidence to be greatest in markets with one or two informed traders – yet participants are well calibrated in these markets.

Our results suggest that forecast bias affects the accuracy of informed traders' predictions of asset value. With unbiased forecasts, informed traders' predictions are

generally accurate: that is, their prediction errors, for the most part, are not significantly different from zero. By comparison, with biased forecasts, informed traders' prediction errors are positive. They seem to recognize that forecasts need to be adjusted downward, but their adjustments are insufficient. Hence, their prediction errors are significantly greater than zero.

In sum, informed traders appear to properly judge the quality of unbiased forecasts. But they misjudge the quality of biased forecast – particularly the accuracy of the information. Improper assessment of information quality may hinder informed traders' ability to earn excess profit. We address this issue in the next subsection.

4.3. Participants' trading ability

We scrutinize the performance of informed traders to determine whether they prosper in our experimental markets. We compare the profit of informed and uninformed traders, looking at both gross and net profit. We delve into what affects profit-making ability, including information dissemination (i.e., whether asset price fully reflects private information) and trading activity (i.e., transactions and quotes that involve informed versus uninformed traders).

4.3.1. Trading profit. We compute the average profit per period of informed and uninformed traders, including periods 4-12. For informed traders, we compute gross profit (excluding the cost of information) and net profit (subtracting the cost of information). We perform parametric paired t-tests and nonparametric Wilcoxon signed ranks tests to compare the trading profit of informed agents (gross and net) with that of uninformed agents. We perform separate tests for each experimental condition.

Table 6 presents the results for markets with unbiased forecasts. Informed agents earn significantly more gross profit in markets with one informed trader (refer to Panel A of Table 6), but not in other markets ($p > 0.30$). Taking the cost of information into account, informed agents earn *less* than uninformed agents, with the difference typically being statistically significant.[12] This finding, coupled with the earlier result that informed agents properly gauge the quality of unbiased forecasts, suggest that informed agents are overconfident in their trading ability.

Table 7 presents the results of paired tests for markets with biased forecasts. We find that informed agents earn significantly more gross profit than uninformed agents in markets with one informed trader ($p < 0.015$); the difference hovers around significance with two informed traders ($p = 0.063$ with a parametric test and $p = 0.107$ with a nonparametric test); and the difference is insignificant with four informed traders ($p > 0.45$). Subtracting the cost of information, informed agents earn significantly *less* than uninformed agents in markets with one or two informed traders ($p < 0.10$). Overall, the findings for trading profit are comparable across markets with unbiased and biased forecasts.

4.3.2. Information dissemination. We investigate information dissemination to assess the opportunity afforded informed agents to generate excess profit. We examine the closeness of period-end asset price with a benchmark price that fully reflects private information. We use the unbiased estimate of asset value conditioned on private information (UE_t) as the benchmark. For UE_t, we use the unbiased forecast in period t.

[12] We repeat the analyses separating informed agents by their tax rate on dividend earnings (i.e., zero or 20 percent). We want to ensure that the results are not being driven by informed agents' with a 20 percent tax rate (i.e., those who have less to gain by being informed). We find that the results are qualitatively similar. In no case do informed agents earn more than uninformed agents after subtracting the cost of information.

A Bayesian adjustment can be applied (as discussed earlier), although it does not affect the results. Hence, the data reported and the tests performed use the unbiased forecast as the benchmark price.

For each asset market, we compute the absolute deviation in the last price (LP) per period from the benchmark price (UE_t), normalized by the benchmark price. That is,

$$NAPD_t = |(LP_t - UE_t)| / UE_t.$$

Figure 3 presents the average NAPD per period across markets with one, two, and four informed traders. Panel A includes data for markets with unbiased forecasts, and Panel B includes data for markets with biased forecasts. A cursory inspection of the data suggests that (1) the NAPD is considerably smaller in markets with four informed traders and (2) the patterns are similar across markets with unbiased and biased forecasts.

To conduct formal statistical tests, we perform a linear mixed model analysis using maximum likelihood estimation. As before, we use an AR(1) covariance structure, though the results are robust to alternative specifications. The independent variables include the number of informed traders, forecast bias, and the interaction term. The dependent variable is the NAPD per period, repeated over periods 4-12.[13]

We find that the number of informed traders is statistically significant ($F = 28.32$, $p < 0.001$), whereas the other two variables are insignificant ($p > 0.65$). The average NAPD is 0.262, 0.216, and 0.077 in markets with one, two, and four informed traders, respectively. Bonferroni pair-wise comparisons indicate that the average NAPD is significantly less ($p < 0.001$) in markets four informed traders than in other markets and not significantly different between markets with one and two informed traders.

[13] We also compute NAPD using the average and median of the last three prices per period and repeat the analyses. Inferences are unaffected.

Accordingly, informed agents have *better* profit-making opportunities in markets with one or two informed traders than in markets with four informed agents. Yet our earlier analysis shows that informed agents are unable to generate excess profit in any markets. Moreover, they typically cannot even recoup the cost of information.

To further analyze information dissemination, we examine uninformed agents' predictions of asset value. If information dissemination is weaker in markets with one or two informed traders, then uninformed agents should have more difficulty predicting asset value. We repeat the linear mixed model analysis focusing on uninformed agents' prediction error. For each market, we compute the average absolute prediction error of uninformed agents per period, normalized by the unbiased forecast. We use this measure as the dependent variable in the analysis.

The results (not tabulated) show that the number of informed traders is statistically significant ($F = 16.67$, $p < 0.001$). Bonferroni pair-wise tests indicate that the average absolute normalized prediction error is significantly smaller ($p < 0.001$) in markets with four informed traders (mean = 0.179) than in markets with one (mean = 0.354) or two (mean = 0.291) informed traders. The difference is insignificant comparing markets with one or two informed traders ($p > 0.10$). Again, the findings suggest that profit-making opportunities are better in markets with one or two informed traders.[14] Nevertheless, informed agents are unable to reap the rewards.

4.3.3. Trading activity. We examine trading activity (trades and quotes) to shed light on trader performance. We compute the average activity (traders and quotes) by

[14] As might be expected, the average absolute normalized prediction error is smaller for informed traders than uninformed traders. To assess statistical significance, we perform parametric and nonparametric paired tests for each experimental condition. We find that the difference is always significant at $p \leq 0.001$ in markets with one or two informed traders. The difference is generally significant at $p < 0.10$ in markets with four informed traders.

trader per period over periods 4-12. Table 8 presents the average activity, partitioned by informed and uninformed agents, for markets with unbiased forecasts. The descriptive data indicate that informed agents engage in more activity, on average, than uninformed agents. To assess statistical significance, we perform parametric paired t-tests and nonparametric Wilcoxon signed rank tests.

We find that informed agents transact more frequently than uninformed agents in markets with one or two informed traders ($p < 0.01$), but not in markets with four informed traders. The average number of transactions per informed trader, though, is small – ranging from around one to two and one-half transactions per trader. We also find some evidence that informed agents post more quotes than uninformed agents (in markets with one or four informed traders), but again the number is not particularly large – the average number of quotes per trader type is roughly two to three times the number of transactions.

Table 9 presents the average activity, separated by informed and uninformed agents, for markets with biased forecasts. The descriptive data are similar to the data for markets with unbiased forecasts. We find that informed agents engage in more activity than uninformed agents, with the raw numbers being fairly small per trader. Most paired comparisons are statistically significant at conventional levels.

What we glean from the data on trading activity is that informed agents do not engage in enough trades to outperform uninformed agents. In markets with one or two informed traders the cost of information is often in excess of 1,000 (see Figure 2). Recall that earlier we computed an expected value of private information of 280 or 208 per transaction, conditioned on informed agents' tax rate. This figure represents the expected

profit per transaction assuming (1) that information is not disseminated and (2) that uninformed agents are willing to transact at the mean asset value (i.e., the uninformed prior). But even with these naïve assumptions, informed agents do not trade enough to recoup the cost of information, which is why they fare worse than uninformed agents (see Tables 6 and 7). We suggest informed agents' concerns over trading strategically and suppressing information dissemination causes them to behave too conservatively (limiting their activity too much).

In markets with four informed traders, the cost of information is much less, approaching the expected value of private information – almost always less than 500 and very often below 300 (see Figure 2). On the other hand, information is more widely disseminated in these markets (see Figure 3). Thus, the assumptions underlying the expected value of private information per transaction are unlikely to hold. Because informed agents do not transact much in these markets (just over one transaction per trader), they earn less than uninformed agents (refer to Table 6).

Conclusion

We use an experimental approach to investigate the effect of overconfidence on traders' performance. Our use of an experimental approach allows us to carefully control the supply of forecasted information and its properties, an investigation that cannot be conducted in naturally occurring markets. Importantly, overconfidence is not static across traders or time in our markets as confidence evolves endogenously. Prior to trading an asset in a double auction market, traders bid to acquire a forecast of the asset's

value and those who "win" the auction appear to overestimate its usefulness. The effect of winning intensifies as the numbers of winners decreases.

Daniel, Hirshleifer, and Subrahmanyam (1998, p. 1841) define an overconfident trader "as one who overestimates the precision of his private information signal, but not of information signals publicly received by all." To investigate whether informed traders in our markets overestimate the quality of private information we elicit assessments of asset value each period. Our analysis suggests that traders are reasonably well-calibrated and understand the accuracy of private information. However, informed traders appear to misjudge their skill at generating trading profit. While informed traders attempt to take advantage of their information advantage, they do not trade enough to generate superior earnings net of the cost of information.

Our findings provide important directions for theoretical research. A fundamental question surrounds the impact of overconfidence on performance. Clearly some level of self-confidence is requisite for success in any domain but caution is warranted as success endogenously fuels overconfidence. Additional analytical research along these lines may yield valuable insights into the behavior of investors.

Figure 1: Asset value distribution

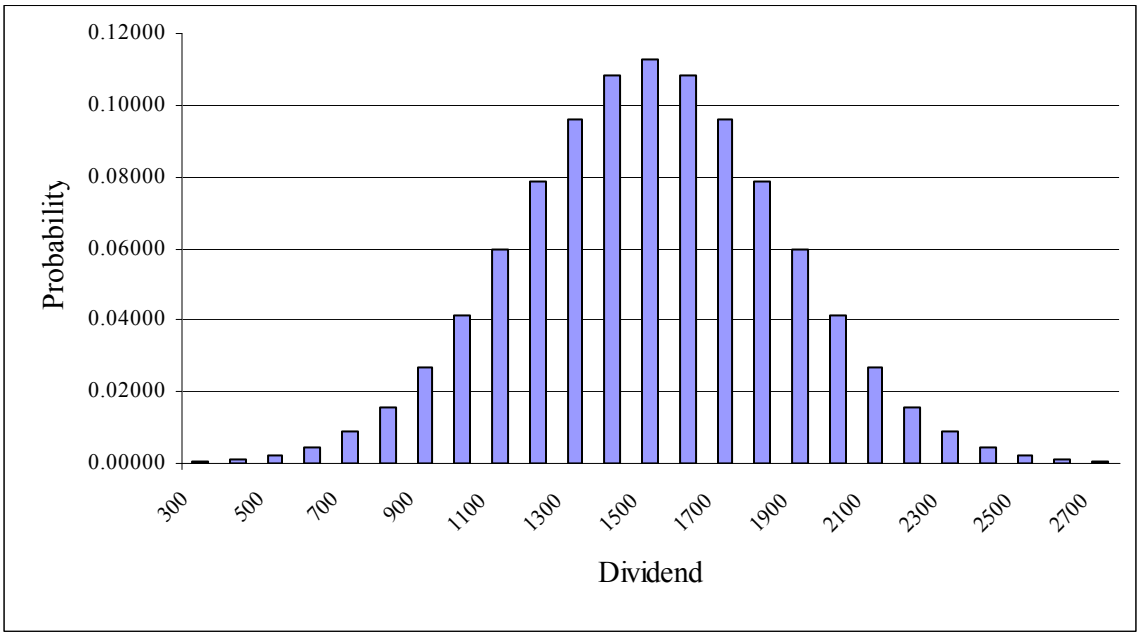

Notes: The discrete distribution shown above is used to determine period-end asset value.

Figure 2: Average price of information per period

Panel A: Markets with unbiased forecast of asset value

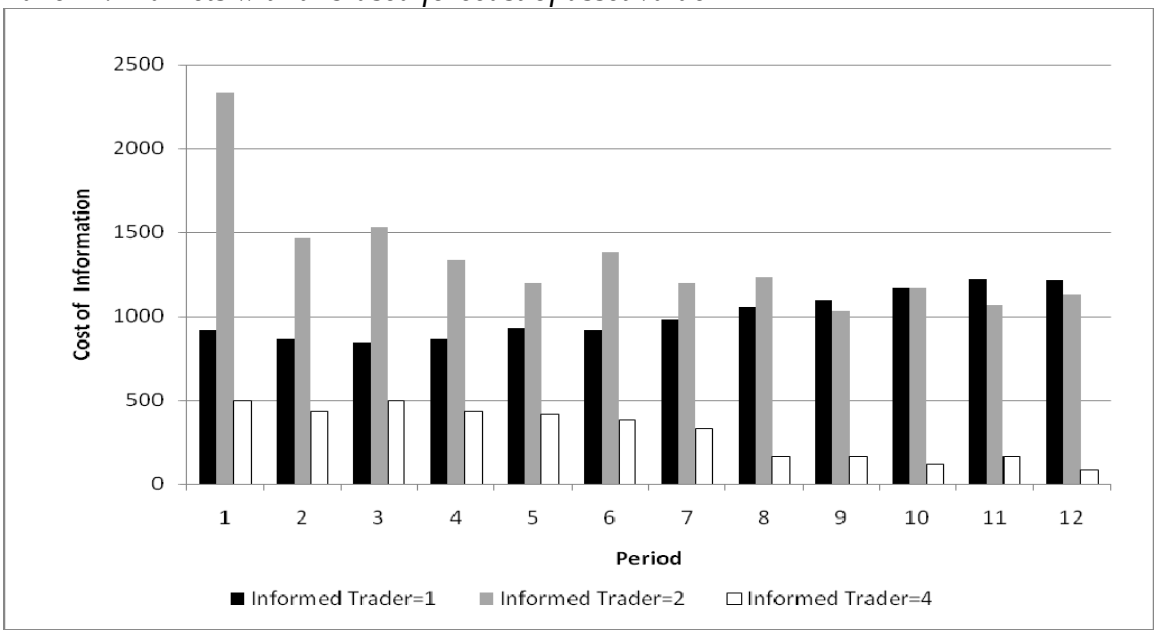

Panel B: Markets wih biased forecast of asset value

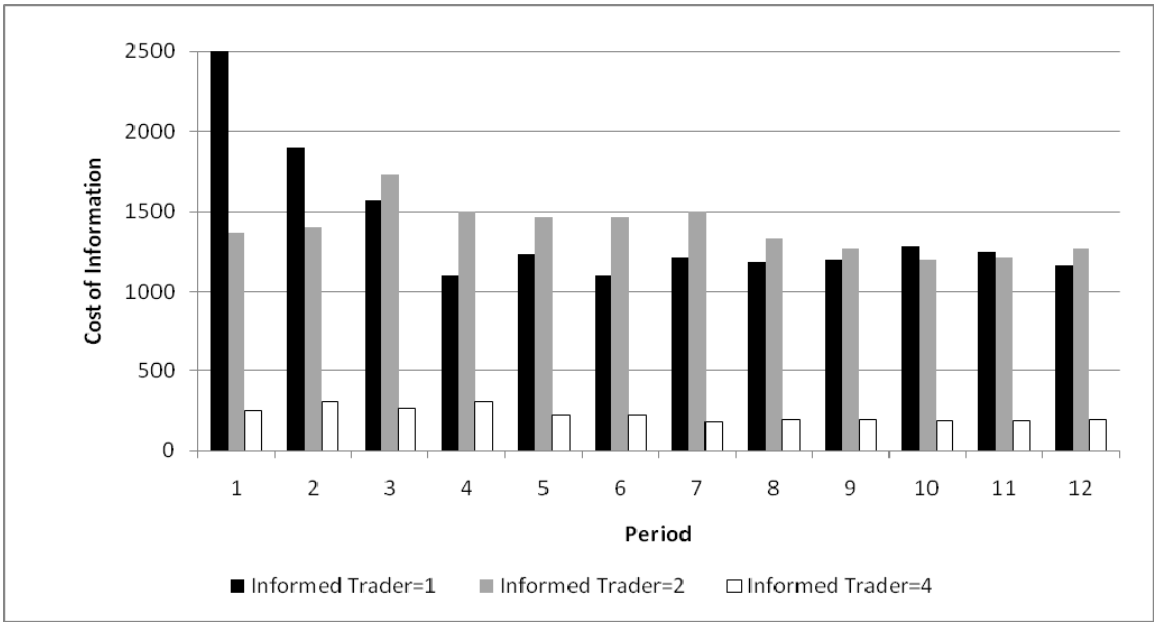

Notes: Three asset markets are conducted under each experimental condition. The average price per period is computed averaging over prices over the three markets.

Figure 3: Normalized absolute price deviations

Panel A: Markets with unbiased forecast of asset value

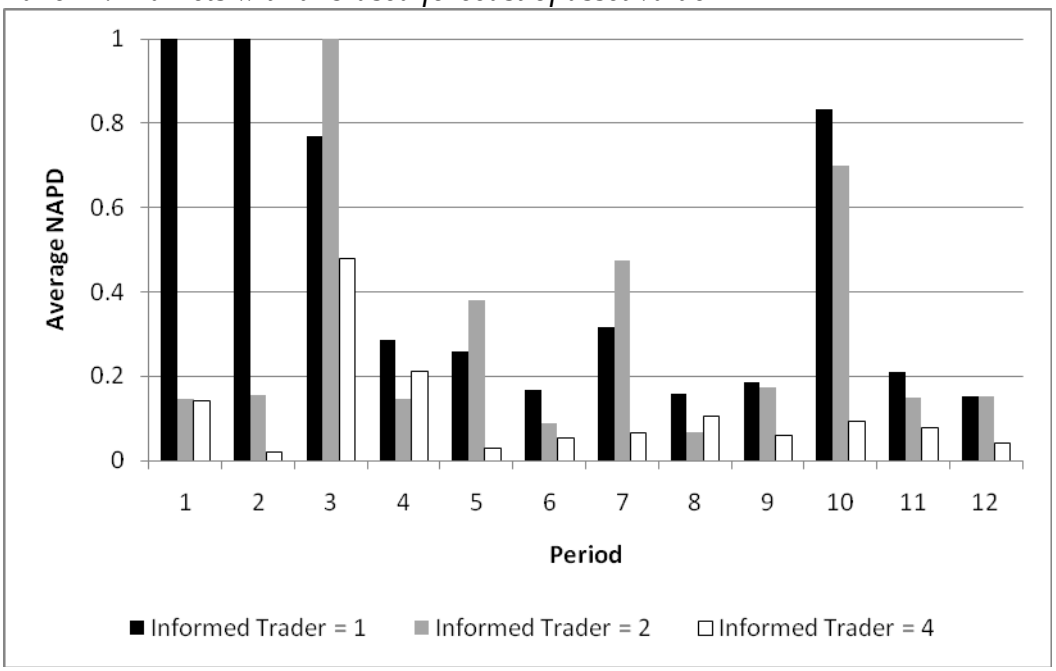

Panel B: Markets with biased forecasts of asset value

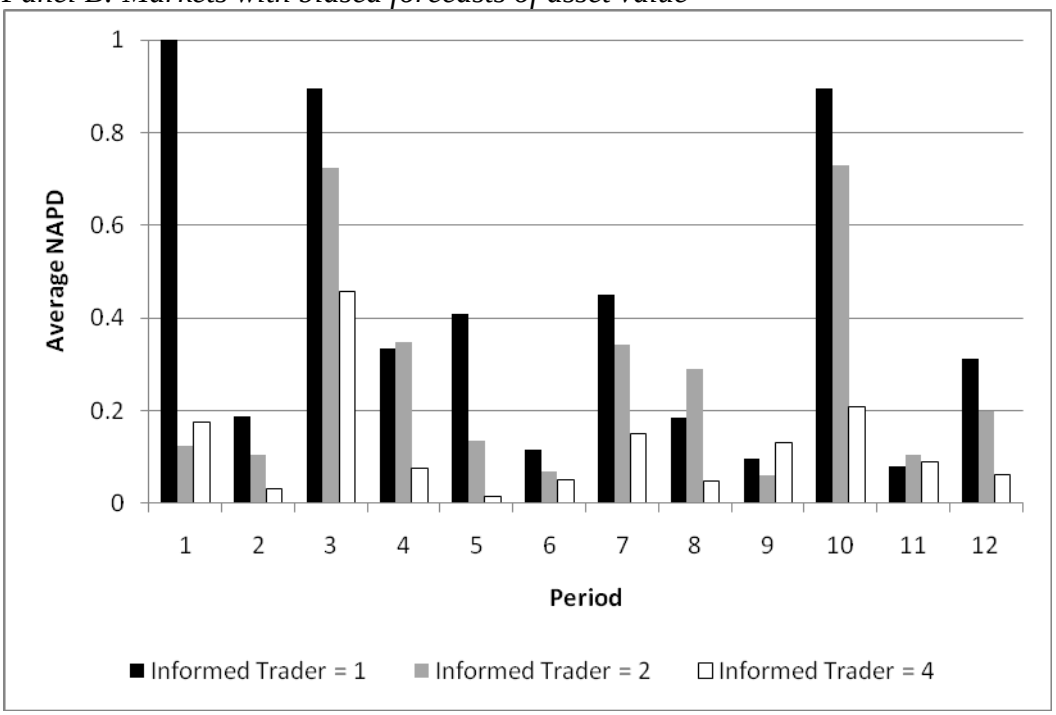

Notes: For each market, we compute the normalized absoluted price deviation per period (NAPD), defined as (last price − unbiased forecast) dividided by unbiased forecast. Three asset markets are conducted under each experimental condition. We compute the average NAPD per period across the the three markets conducted under the same experimental conditions.

Table 1
Unbiased forecasts and informed traders' calibration

Panel A: Frequency of calibrated predictions by number of informed traders per market

Number of informed traders per market	Calibrated	Not Calibrated
One	25 (0.93)	2 (0.07)
Two	50 (0.93)	4 (0.07)
Four	87 (0.81)	21 (0.19)

Panel B: Frequency of forecast acquisitions and calibrated predictions

Frequency of forecast acquisition	Calibrated	Miscalibrated		
		Once	Twice	More than twice
1	8	1	-	-
2	5	0	0	-
3	7	1	1	0
4	3	1	0	0
5	1	0	0	0
6	4	2	0	1
7	3	0	1	0
8	0	1	0	1
9	2	1	0	1
Total	33	7	2	3

Notes: The cell entries in Panel A show the number of times that informed traders' predictions per period (over periods 4-12) are calibrated and not calibrated. The percentage is shown parenthetically. A prediction is calibrated as long as it falls within the interval [forecast – 200, forecast + 200], else it is not calibrated. The cell entries in Panel B show the frequency of forecast acquisitions (over periods 4-12) by various categories regarding the calibration of the prediction. Dashes indicate that the cell entry is not applicable for a particular row (e.g., if the forecast is acquired once, the trader's prediction cannot be miscalibrated more than one period).

Table 2
Biased forecasts and informed traders' calibration

Panel A: Frequency of calibrated predictions by number of informed traders per market

Number of informed traders per market	Calibrated	Not Calibrated
One	23 (0.85)	4 (0.15)
Two	40 (0.74)	14 (0.26)
Four	87 (0.81)	21 (0.19)

Panel B: Frequency of forecast acquisitions and calibrated predictions

Frequency of forecast acquisition	Calibrated	Miscalibrated		
		Once	Twice	More than twice
1	7	2	-	-
2	2	4	1	-
3	3	0	0	0
4	3	0	1	0
5	0	3	0	1
6	6	0	1	0
7	1	0	1	0
8	0	0	0	0
9	2	2	0	3
Total	25	11	4	4

Notes: The cell entries in Panel A show the number of times that informed traders' predictions per period (over periods 4-12) are calibrated and not calibrated. The percentage is shown parenthetically. A prediction is calibrated as long as it falls within the interval [forecast – 200, forecast + 200], else it is not calibrated. The cell entries in Panel B show the frequency of forecast acquisitions (over periods 4-12) by various categories regarding the calibration of the prediction. Dashes indicate that the cell entry is not applicable for a particular row (e.g., if the forecast is acquired once, the trader's prediction cannot be miscalibrated more than one period).

Table 3
Experimental parameters: Unbiased estimates of asset value and period-end dividend

Period	Unbiased Forecast	Estimate with Bayesian Adjustment	Period-End Dividend (Asset Value)
1	1,900	1,854	1700
2	1,600	1,588	1600
3	800	878	1000
4	2,400	2,304	2400
5	1,100	1,146	1200
6	1,600	1,588	1500
7	1,000	1,057	1000
8	1,900	1,854	1800
9	1,400	1,412	1300
10	800	878	1000
11	1,900	1,854	1900
12	1,200	1,235	1300

Notes: The unbiased estimate of asset value can be computed assuming that the forecast distribution (i.e., process generating the forecast) is unknown or known. If the forecast distribution is unknown, the unbiased estimate is the unbiased forecast. If the forecast distribution is known (i.e., learned over the course of the experiment), a Bayesian adjustment can be applied to refine the unbiased estimate. The period-end dividend (asset value) is drawn from a known distribution (included in the experimental materials).

Table 4
Accuracy of informed traders' prediction of asset value with unbiased forecasts

Panel A: Normalized prediction errors

Number of Informed Traders	Mean NPE (Std Dev)	t-statistic (p-value)
One	−0.029 (0.077)	−1.962 (0.061)
Two	0.012 (0.148)	0.597 (0.553)
Four	−0.001 (0.222)	−0.046 (0.963)
Total	−0.001 (0.187)	−0.094 (0.925)

Panel B: Signed prediction errors

Number of Informed Traders	PE < 0	PE = 0	PE > 0
One	9	15	3
Two	9	37	8
Four	43	33	32
Total	61	85	43

Notes: The normalized prediction error (NPE) is computed as the difference between informed traders' prediction of asset value and the unbiased estimated (the unbiased forecast of asset value), normalized by the unbiased estimate. The data include informed traders predictions over periods 4-12. Panel A reports the results of t-tests to determine whether the mean NPE is significantly different from zero. Panel B reports the number of times that informed traders' prediction (PE) is negative, zero, and positive. Prediction error is computed as the difference between informed traders' prediction of asset value and the unbiased estimate.

Table 5
Accuracy of informed traders' prediction of asset value with biased forecasts

Panel A: Normalized prediction errors

Number of Informed Traders	Mean NPE (Std Dev)	t-statistic (p-value)
One	0.112 (0.125)	4.687 (0.000)
Two	0.070 (0.185)	2.779 (0.008)
Four	0.066 (0.194)	3.517 (0.001)
Total	0.074 (0.183)	5.531 (0.000)

Panel B: Signed prediction errors

Number of Informed Traders	PE < 0	PE = 0	PE > 0
One	3	1	23
Two	14	4	36
Four	18	32	57
Total	35	37	116

Panel C: Normalized versus unadjusted (biased) prediction errors

Number of Informed Traders	Mean Difference (Std Dev)	t-statistic (p-value)
One	−0.039 (0.086)	−2.325 (0.028)
Two	−0.081 (0.167)	−3.557 (0.001)
Four	−0.085 (0.171)	−5.121 (0.000)
Total	−0.077 (0.161)	−6.572 (0.000)

Notes: The normalized prediction error (NPE) is computed as the difference between informed traders' prediction of asset value and the unbiased estimated (the unbiased forecast of asset value), normalized by the unbiased estimate. The data include informed traders predictions over periods 4-12. Panel A reports the results of t-tests to determine whether the mean NPE is significantly different from zero. Panel B reports the number of times that informed traders' prediction (PE) is negative, zero, and positive. Prediction error is computed as the difference between informed traders' prediction of asset value and the unbiased estimate. Panel C reports the results of paired t-tests to determine whether the NPE differs significantly from NPE that does not adjust for bias at all (i.e., the prediction error equals the systematic bias).

Table 6
Profit of informed and uninformed traders with unbiased forecasts

Number of Informed Traders	Profit	Average Profit (Std Dev)	t-statistic (p-value	z-statistic (p-value)
One	Informed GP	3,364 (1,410)	3.476 (0.002)	−3.035 (0.002)
	Uninformed	2,659 (831)		
	Informed NP	2,312 (1,573)	−1.454 (0.158)	−1.874 (0.061)
Two	Informed GP	2,786 (825)	0.973 (0.340)	−0.721 (0.471)
	Uninformed	2,637 (874)		
	Informed NP	1,591 (814)	−6.505 (0.000)	−4.252 (0.000)
Four	Informed GP	2,772 (834)	0.615 (0.544)	−0.805 (0.421)
	Uninformed	2,728 (877)		
	Informed NP	2,522 (825)	−2.640 (0.014)	−2.427 (0.015)

Notes: Informed GP (gross profit) is the average informed trader profit per period excluding the cost of information (periods 4-12). Uninformed is the average uninformed trader profit per period (periods 4-12). Informed NP (net profit) is the average informed trader profit per period subtracting the cost of information (periods 4-12). The last two columns report test statistics and p-values (two tailed) for parametric paired t-tests and nonparametric Wilcoxon signed ranks tests.

Table 7
Profit of informed and uninformed traders with biased forecasts

Number of Informed Traders	Profit	Average Profit (Std Dev)	t-statistic (p-value	z-statistic (p-value)
One	Informed GP	3,362 (1,582)	2.702 (0.012)	−2.931 (0.003)
	Uninformed	2,646 (850)		
	Informed NP	2,166 (1,614)	−1.801 (0.083)	−2.984 (0.003)
Two	Informed GP	2,910 (1,059)	1.940 (0.063)	−1.610 (0.107)
	Uninformed	2,626 (911)		
	Informed NP	1,545 (1,302)	−5.305 (0.000)	−3.822 (0.000)
Four	Informed GP	2,766 (715)	0.693 (0.495)	−0.094 (0.925)
	Uninformed	2,692 (773)		
	Informed NP	2,572 (727)	−1.107 (0.278)	−1.309 (0.190)

Notes: Informed GP (gross profit) is the average informed trader profit per period excluding the cost of information (periods 4-12). Uninformed is the average uninformed trader profit per period (periods 4-12). Informed NP (net profit) is the average informed trader profit per period subtracting the cost of information (periods 4-12). The last two columns report test statistics and p-values (two tailed) for parametric paired t-tests and nonparametric Wilcoxon signed ranks tests.

Table 8
Trading activity with unbiased forecasts

Panel A: Markets with one informed traders

Activity	Trader	Average (Std Dev)	t-statistic (p-value)	z-statistic (p-value)
Trades	Informed	2.07 (1.57)	3.201 (0.004)	−3.506 (0.005)
	Uninformed	1.20 (0.67)		
Quotes	Informed	6.11 (2.69)	4.449 (0.000)	−2.797 (0.000)
	Uninformed	3.64 (1.46)		

Panel B: Markets with two informed traders

Activity	Trader	Average (Std Dev)	t-statistic (p-value)	z-statistic (p-value)
Trades	Informed	2.48 (0.92)	3.825 (0.001)	−3.273 (0.001)
	Uninformed	1.77 (0.63)		
Quotes	Informed	5.26 (2.47)	1.519 (0.141)	−1.069 (0.285)
	Uninformed	4.40 (0.98)		

Panel C: Markets with four informed traders

Activity	Trader	Average (Std Dev)	t-statistic (p-value)	z-statistic (p-value)
Trades	Informed	1.27 (0.57)	0.567 (0.576)	−0.524 (0.600)
	Uninformed	1.19 (0.58)		
Quotes	Informed	4.31 (1.24)	1.840 (0.077)	−1.665 (0.096)
	Uninformed	3.76 (1.37)		

Notes: For each type of trader (informed and uninformed), we compute the average activity (trades and quotes) per trader per period over periods 4-12. The last two columns report test statistics and p-values (two tailed) for parametric paired t-tests and nonparametric Wilcoxon signed ranks tests.

Table 9
Trading activity with biased forecasts

Panel A: Markets with one informed traders

Activity	Trader	Average (Std Dev)	t-statistic (p-value)	z-statistic (p-value)
Trades	Informed	1.85 (1.43)	3.201 (0.029)	−2.021 (0.043)
	Uninformed	1.26 (0.82)		
Quotes	Informed	4.74 (2.55)	2.314 (0.002)	−2.772 (0.006)
	Uninformed	2.77 (0.94)		

Panel B: Markets with two informed traders

Activity	Trader	Average (Std Dev)	t-statistic (p-value)	z-statistic (p-value)
Trades	Informed	1.76 (1.16)	3.304 (0.003)	−2.818 (0.005)
	Uninformed	1.06 (0.38)		
Quotes	Informed	3.91 (2.41)	1.904 (0.068)	−1.406 (0.160)
	Uninformed	2.94 (0.74)		

Panel C: Markets with four informed traders

Activity	Trader	Average (Std Dev)	t-statistic (p-value)	z-statistic (p-value)
Trades	Informed	1.44 (0.70)	2.180 (0.038)	−2.187 (0.029)
	Uninformed	1.21 (0.56)		
Quotes	Informed	3.67 (1.94)	3.185 (0.004)	−2.783 (0.005)
	Uninformed	2.60 (1.09)		

Notes: For each type of trader (informed and uninformed), we compute the average activity (trades and quotes) per trader per period over periods 4-12. The last two columns report test statistics and p-values (two tailed) for parametric paired t-tests and nonparametric Wilcoxon signed ranks tests.

References

Ackert, Lucy F., Bryan K. Church, and Mohamed Shehata. 1997. Market behavior in the presence of costly, imperfect information: Experimental evidence. *Journal of Economic Behavior and Organization* 33 (May), 61-74.

Allen, W. David, and Dorla A. Evans. 2005. Bidding and overconfidence in experimental financial markets. *Journal of Behavioral Finance* 6(3), 108-120.

Aronson, Elliot. 1969. A theory of cognitive dissonance: A current perspective. In L. Berkowitz (ed.) *Advances in Experimental Social Psychology* 4, (New York: Academic Press), 1-34.

Barber, Brad M., and Terrance Odean. 2000. Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance* 55(2), 773-806.

Barber, Brad M., and Terrance Odean. 2001. Boys will be boys: Gender, overconfidence, and common stock investment. *Quarterly Journal of Economics* 116(1), 261-292.

Barefield, Russell M., and Eugene E. Comiskey. 1975. The accuracy of analysts' forecasts of earnings per share. *Journal of Business Research* 3(3), 241-252.

Baumeister, Roy F. 1998. The self. In D.T. Gilbert, S.T. Fiske, and G. Lindzey (eds.) *Handbook of Social Psychology* (New York: McGraw-Hill), 680-740.

Bazerman, Max H., and William F. Samuelson. 1983. I won the auction but don't want the prize. *Journal of Conflict Resolution* 27(4), 618-634.

Benos, Alexandros V. 1998. Aggressiveness and survival of overconfident traders. *Journal of Financial Markets* 1, 353-383.

Biais, Bruno, Denis Hilton, Karine Mazurier, and Sebastian Pouget. 2005. Judgemental overconfident, self-monitoring, and trading performance in an experimental financial market. *Review of Economic Studies* 72, 287-312.

Cason, Timothy N., and Daniel Friedman. 1996. Price formation in double auction markets. *Journal of Economic Dynamics and Control* 20 (August), 1307-1337.

Copeland, T.E., and D. Friedman. 1992. The market value of information: Some experimental results. *Journal of Business* 65 (2), 241-266.

Daniel, Kent D., David Hirshleifer, and Avanidhar Subrahmanyam. 1998. Investor psychology and security market under- and overreactions. *Journal of Finance* 53(6), 1839-1885.

Festinger, Leon. 1957. *A Theory of Cognitive Dissonance*, Stanford University Press (Stanford, CA).

Fischhoff, Baruch, Sarah Lichtenstein, and Paul Slovic. 1977. Knowing with certainty: The appropriateness of extreme confidence. *Journal of Experimental Psychology* 3(4), 552-564.

Galinsky, Adam D., and Thomas Mussweiler. 2001. First offers as anchors: The role of perspective-taking and negotiator focus. *Journal of Personality and Social Psychology* 81(4), pp. 657-669.

Gervais, Simon, and Terrance Odean. 2001. Learning to be overconfident. *Review of Financial Studies* 14(1), 1-27.

Glaser, Markus, and Martin Weber. 2007. Overconfidence and trading volume, *Geneva Risk and Insurance Review* 32, 1-36.

Greene, William H. 1997. *Econometric Analysis*, Third Edition, Prentice-Hall (Upper Saddle River, New Jersey).

Grinblatt, Mark, and Matti Keloharju. 2009. Sensation seeking, overconfidence, and trading activity. *Journal of Finance* 64(2), 549-578.

Grosskopf, Brit, Yoella Bereby-Meyer, and Max Bazerman. 2007. On the robustness of the winner's curse. *Theory and Decision* 63(4), 389-418.

Hastie, Reid, David A. Schkade, and John W. Payne. 1999. Juror judgments in civil cases: Effects of plaintiff's requests and plaintiff's identity on punitive damage awards. *Law and Human Behavior* 23 (August), 445-470.

Hellwig, Martin F. 1980. On the aggregation of information in a competitive market. *Journal of Economic Theory* 22, 477-498.

Hirshleifer, David, and Guo Ying Lou. 2001. On the survival of overconfident traders in competitive securities market. *Journal of Financial Markets* 4(1), 73-84.

Hong, Han, and Matthew Shum. 2002. Increasing competition and the winner's curse: Evidence from procurement. *Review of Economic Studies* 69(4), 871-898.

Kagel, John H. 1995. Auctions: A survey of experimental research. In J.J. Kagel and A.E. Roth (eds.) *The Handbook of Experimental Economics* (Princeton, NJ: Princeton University Press), 501-585.

Kirchler, Erich, and Boris Maciejovsky. 2002. Simultaneous over- and underconfidence: Evidence form experimental asset markets. *Journal of Risk and Uncertainty* 25(1), 65-85.

Ku, Gillian, Deepak Malhotra, and J. Keith Murhnighan. 2005. Towards a competitive arousal model of decision-making: A study of auction fever in live and Internet auctions. *Organizational Behavior and Human Decision Processes* 96(2), 89-103.

Kyle, Albert S., 1984. Market structure, information, futures markets, and price formation. In G.G. Storey, A. Schmitz, and A.H. Harris (eds.) *International Agricultural Trade: Advanced Readings in Price Formation, Market Structure, and Price Stability* (Boulder, CO: Westview Press), 45-64.

Kyle, Albert S. 1989. Informed speculation with imperfect competition. *Review of Economic Studies* 56(3), 317-355.

Kyle, Albert S., and F. Albert Wang. 1997. Speculation duopoly with agreement to disagree: Can overconfidence survive the market test? *Journal of Finance* 52(5), 2073-2090.

Lewinsohn, Shai, and Haim Mano. 1993. Multi-attribute choice and affect: The influence of naturally occurring and manipulated moods on choice processes. *Journal of Behavioral Decision Making* 6(1), 33-51.

Mano, Haim. 1994. Risk taking, framing effects, and affect. *Organizational Behavior and Human Decision Processes* 57(1), 38-58.

Milgrom, Paul, and Nancy L. Stokey. 1982. Information, trade, and common knowledge. *Journal of Economic Theory* 26, 17-27.

O'Brien, Patricia C. (1988). Analysts' forecasts as earnings expectations. *Journal of Accounting and Economics* 10(1), 53-83.

Odean, Terrance. 1998. Volume, volatility, price, and profit when all traders are above average. *Journal of Finance* 53(6), 1887-1934.

Oskamp, Stuart. 1965. Overconfidence in case-study judgments. *Journal of Consulting Psychology* 29(3); 261-265.

Richardson, Scott, Siew Hong Teoh, and Peter Wysocki. 2004. The walkdown to beatable analyst forecasts: The roles of equity issuance and insider trading incentives. *Contemporary Accounting Research* 21(4), 885-924.

Schnitzlein, Charles R. 2002. Price formation and market quality when the number and presence of insiders is unknown. *Review of Financial Studies* 15(4), 1077-1109.

Sunder, Shyam. 1992. Market for information: Experimental evidence. *Econometrica* 60 (May), 667-69.

Svenson, Ola. 1981. Are we all less risky and more skillful than our fellow drivers? *Acta Psychologica* 47(2), 143-148.

Taylor, Shelley, E., and Jonathan D. Brown. 1988. Illusion and well-being: A social psychological perspective on mental health. *Psychological Bulletin* 103(2), 193-210.

Thaler, Richard H. 1992. *The Winner's Curse: Paradoxes and Anomalies of Economic Life* (Princeton University Press, Princeton, New Jersey).

Appendix

Experimental Instructions

The following instructions are for the treatment with 4 informed traders. Changes with two or one informed trader(s) are notes in brackets ([]) and ({}), respectively.

General Instructions
This experiment is concerned with the economics of market decision-making. We are going to create a market in which you will buy and sell certificates in a sequence of market periods. Based on your trading decisions, you are able to generate profits, which will be paid to you in cash at the conclusion of this experiment.

Attached to these materials you will find a set of Information and Record Sheets, one for each period, one Dividend Prediction Sheet, and one Profit Sheet. Please refer to these sheets while going through the instructions.

Specific Instructions
Trading

Your trading profits come from two sources--from collecting dividends on all certificates you hold at the end of each period and from buying and selling certificates. During each market period you are free to purchase or sell as many certificates as you wish, provided you follow the rules below. For each certificate you hold at the end of a period you will receive a dividend. At the beginning of each period, **four** [two] {one}participant(s) will be allowed to purchase a forecast of the current period's dividend. The process by which forecasts are acquired is explained later in these instructions. How the dividend and corresponding forecast are determined each period are also explained later.

Your total dividends for a period are computed by multiplying the dividend per certificate by the number of certificates held. Suppose, for example, that you hold two certificates at the end of period 1. If for that period the dividend is $1,500 per certificate, then your total dividends in the period are 2 x $1,500 = $3,000.

Sales from your certificate holdings increase your cash on hand by the amount of the sale price. Similarly, purchases reduce your cash on hand by the amount of the purchase price. Thus you can gain or lose money on the purchase and resale of certificates.

At the beginning of each period you are provided with 2 certificates and $50,000. You are free to hold, buy, or sell certificates according to the rules discussed subsequently. If you hold a certificate, then you receive a dividend at the end of the period.

You are required to pay a personal tax on your dividend earnings. Personal tax rates differ across traders and between periods. Each period, your personal tax rate is either 0% or 20%, as reflected on row 17 of your Information and Record Sheet: **note that your rate can vary across periods**. Your net profit at period end is determined by the

dividends you receive, the cost of acquiring a forecast of the dividend, gains and losses from purchases and sales of certificates, and your personal tax rate. At the end of the experiment, you will be paid 0.1% of this after-tax profit in cash. Thus, for every $1,000 dollars you accrue in profit, you earn $1.00.

At the end of each period, you are required to return the $50,000 cash endowment to the experimenter and certificates do not carry forward to the next period.

Determination of Dividends

The dividend you receive for each certificate held at period end is determined by drawing from a distribution. The distribution, which depicts all possible dividends per certificate and the corresponding probabilities of occurrence, is shown below.

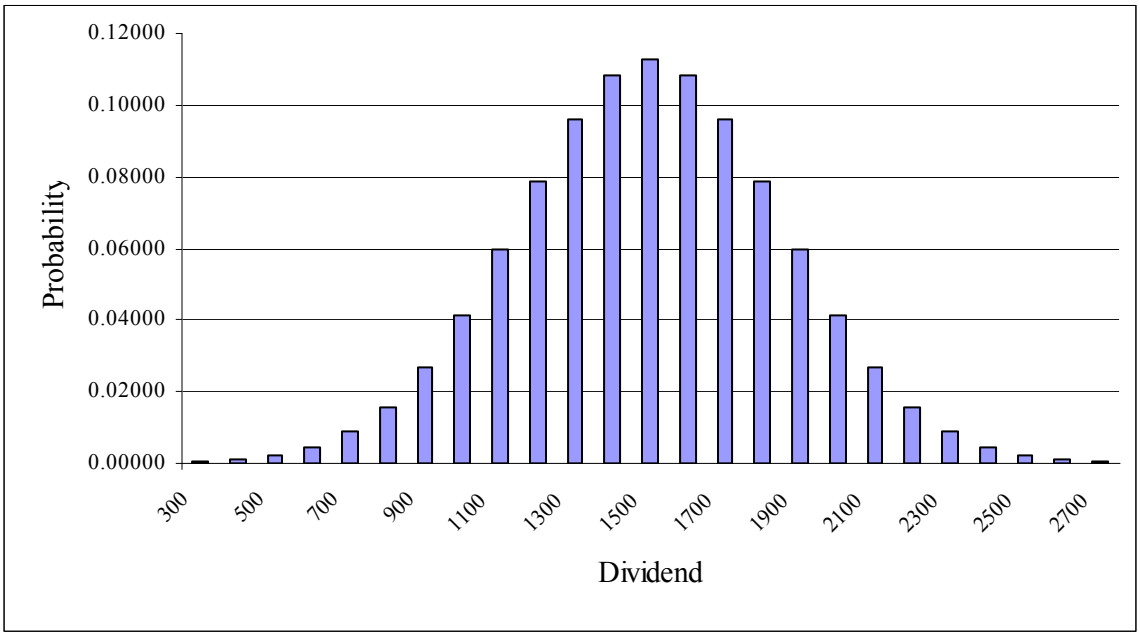

Note that the minimum dividend per certificate is $300 and the maximum is $2,700. Dividends take on values in $100 increments and the average dividend per certificate is $1,500.

Prior to conducting the experiment, dividend values were determined using a computer. The experimenters will announce the dividend per certificate at period end.

Forecast of Dividends

At the beginning of each period, you will participate in an auction to determine who acquires a forecast of the dividend. You have been provided with an index card. One side of the card includes your trader number and the other side includes numbers from 1 – 20. These numbers correspond to periods. The experiment will consist of 20 periods at most and may consist of fewer periods. At the beginning of each period, record the amount that you are willing to pay for the forecast. Record the amount on the numbered

line that corresponds to the period. After doing this, place the index card face down in front of you: that is, with your trader number facing up.

The experimenters will collect the index cards and determine who acquires the forecast and how much it costs. **The four [The two] {The} participant(s) with the four [the two] {the} highest bid(s) will acquire the forecast. They {this participant} will pay the fifth [the third] {the second} highest bid.** Any ties will be resolved randomly. **The experimenters will announce the cost of the forecast after it has been determined.**

If you are one of the four [one of the two] {the} highest bidder(s), the experimenters will record the forecast on your index card on the numbered line corresponding to the period. If you did not acquire the forecast, "XX" will be recorded on your index card. All index cards will then be returned. **It is very important that you do not let other participants see your index card.** Showing your card to others can adversely affect the amount of money that you make for participating in the experiment. **Also, please do not tell others if you acquired the forecast or what your bid was, even when the period is over. Remember that "XX" will be recorded on your card if you were NOT one of the four [one of the two] {the} highest bidder(s) so you did not acquire the forecast.**

The process of generating the forecast is unique and constant across periods in the experiment. To allow you to assess the usefulness of the forecast, a forecast history, collected over 10 practice periods, follows.

Forecast History			
Period	Forecast	Dividend	Forecast – Dividend
1	1,700	1,800	-100
2	1,000	800	+200
3	2,300	2,300	0
4	1,900	1,800	+100
5	1,500	1,600	-100
6	1,200	1,200	0
7	1,100	1,000	+100
8	1,700	1,900	-200
9	900	900	0
10	1,600	1,700	-100

You will not be provided with any other information regarding the forecast for the remainder of the experiment unless you acquire it.

Prediction of Dividend

At the conclusion of each period, just prior to the announcement of the dividend, you will predict the dividend for the period. You will receive $0.25 per period for recording your prediction. Your prediction should be recorded on the Dividend Prediction Sheet.

Market Organization

The market for certificates is organized as follows. The market will be conducted in a series of periods. **Each market period lasts four minutes**.

Anyone wishing to purchase a certificate is free to raise his or her hand and make a verbal bid to buy one certificate at a specified price, and anyone with certificates to sell is free to accept or not accept the bid. Likewise, anyone wishing to sell a certificate is free to raise his or her hand and make a verbal offer to sell one certificate at a specified price. **Please wait until the experimenter calls on you to make a bid or offer.** When you are called on, please announce your trader number followed by your bid or offer. If you wish to accept an outstanding bid or offer, shout out accept to buy or accept to sell. In this case, you do not need to raise your hand and wait to be called on. Any ties in acceptances will be resolved by random choice.

If a bid or offer is accepted, a binding contract has been closed for a single certificate, and the contracting parties will record the transaction on their Information and Record Sheets.

Except for the bids, offers, and their acceptances, you are not to speak to other participants. There are likely to be many bids and offers that are not accepted, but you are free to keep trying. **You are free to make as much profit as you can**.

Recording Rules

Trading

(1) During each period, all transactions are for one certificate at a time. After each of your sales or purchases you must record the **TRANSACTION PRICE** in the appropriate column of the Information and Record Sheet depending on the nature of the transaction. The first transaction is recorded on row 1, and succeeding transactions are recorded on subsequent rows.

(2) After each transaction you must calculate and record your new holdings of certificates and your new cash balance. **Your holdings of certificates may never go below zero. Your cash balance may never go below zero.** The last two columns of the information and record sheet are used to keep a running balance of certificates and cash.

(3) At the end of the period, record your dividend earnings on row 16.

(4) Compute your after-tax dividends as follows. Multiply one minus your tax rate by your dividend earnings. Record the amount on row 17.

(5) Compute your end-of-period trading profit as follows. Add your final cash balance from row 15 to your after-tax dividend earnings from row 17. Record the sum on row 18. Then subtract $50,000 and record the difference on row 20. At period end, transfer this amount to your Profit Sheet.

(6) To compute your net profit for the period, you must subtract the cost of the forecast if it was acquired. The Profit Sheet includes columns for trading profit, forecast cost, and net profit.

Predicting

(1) Prior to the announcement of the dividend each period, record your prediction of the dividend on the Dividend Prediction Sheet.

Throughout the course of the experiment, it is extremely important that you record transactions accurately and compute holdings of certificates and cash without error. If at any time you have a question, please ask the experimenters for assistance.

Trader No. _____

Profit Sheet

	(1)	(2)	(3)
Period	*End-of-Period* **After-Tax Profit**	*Cost of Forecast*	*Net Profit* **(1) – (2)**
1			
2			
3			
4			
5			
6			
7			
8			
9			
10			
11			
12			
13			
14			
15			
16			
17			
18			
19			
20			
Sum of net profit over periods			
Multiply the total in the row above by 0.001 to determine your take-home cash from trading			

Dividend Prediction Sheet

Period	**Dividend Prediction**	*Record $0.25 for each prediction recorded*
1		
2		
3		
4		
5		
6		
7		
8		
9		
10		
11		
12		
13		
14		
15		
16		
17		
18		
19		
20		
Sum up to determine your total dividend prediction earnings		

Information and Record Sheet

	Transaction Number	**Transaction Price**		Certificates On Hand	Cash Balance
		Sale	**Purchase**		
Begininng of Period Holdings	Endowment			2	$50,000
	1				
	2				
	3				
	4				
	5				
	6				
	7				
	8				
	9				
	10				
	11				
	12				
	13				
	14				
	15	Final Balance (bring down balances from above)			
Dividend $ _____	16	**Total Dividends:*Dividend per per certificate multiplied by # of certificate on hand at period end***			
Tax Rate _____ %	17	After-tax dividends (multiply (1-tax rate) by dividends)			
	18	Cash balance + after-tax dividends			
	19	Less Endowment			<$50,000>
	20	End of Period After-Tax Profit			

Transfer this amount to the profit sheet