Private Information and Overconfidence in Experimental Asset Markets

Noah Smith

June 24, 2012

Abstract

I study the use and overuse of private signals of fundamental value in a common type of experimental asset market. I also assess the impact of two kinds of psychological overconfidence - miscalibration, and the "better than average effect" - on the use of private signals, and on market outcomes. Traders use their private signals to make buying and selling decisions. Private signals also affect traders' reservation prices. The use of forecasts does not vary depending on measured overconfidence, with the exception that miscalibrated buyers may pay higher prices for assets. Overconfident traders do not trade more, and may actually perform better than other traders. Although the sample size is small, there is a correlation between market volume and the average "better than average effect" of the market. However, there is no observed relationship between average overconfidence and market price or volatility. These results provide strong support for models of financial markets in which traders overweight private signals; however, they offer only mixed and tepid support for psychological overconfidence as the root cause of this overweighting.

1 Introduction

A foundational result in finance theory is the No Trade Theorem (Milgrom and Stokey 1980). If all participants in a market have rational expectations (and if this is common knowledge), and if the initial allocation is Pareto optimal and traders are weakly risk-averse, then the arrival of new private information does not result in trade. The No Trade Theorem is essentially an adverse selection result, with simple intuition behind it - if someone as smart and well-informed as you is willing to sell you an asset for a price of x, why on Earth would you buy it? Of course, this intuition breaks down if you and she value the asset differently - if one of you needs to trade to for reasons of liquidity or timing - or if one of you loves to gamble. But many think it unlikely that the extremely large volumes of trade that we observe in real financial markets are a result of portfolio rebalancing or pure gambling. It could be that much real-world trading represents a failure of rational expectations, and thus a breakdown in the efficiency of the market.

A growing body of empirical research seems to support this view. Odean (1999), for example, finds that excessive trading lowers the returns that individual investors receive on their portfolios, even when controlling for liquidity and timing demands. Odean's proposed explanation for this over-trading is *overconfidence* - the tendency for investors to believe, incorrectly, that their information is of superior quality to that of their counteraparty. That sort of belief would lead traders to ignore the adverse selection problem - if you and I both believe that our information is superior to the other's, then we may trade even though we both know we disagree. For this reason, economists have constructed a number of models in which trading and asset pricing are driven by overweighting of private signals about asset fundamentals - i.e., an incorrect but dogmatic belief of each agent that her own private signals are more precise than those signals really are (and more precise than the signals of others).

This modeling assumption is typically motivated by a large psychology literature on individual overconfidence. Psychologists have shown that people tend both to overestimate the precision of their own information ("miscalibration") and to believe that they are above average in ability (the "better than average effect"). Both of these effects could cause the overweighting of private information in financial markets.

So behavioral theories of overtrading posit a two-step chain of causality: Psychological overconfidence –> Overweighting of private information –> Market outcomes. Experimental and empirical tests of these theories, however, have focused only on the direct link between overconfidence and market outcomes. In this paper, I use an experiment to disentangle the two steps in the causal chain; in other words, I ask two questions:

1. Do traders trade based on ex ante symmetric private signals of asset fundamentals?
2. Do more overconfident traders rely more on private signals to make trading decisions?

To answer these questions, I use a variant of the most common laboratory financial market setup (Smith, Suchanek, and Williams 1988), where groups of subjects trade a single risky asset in three successive fixed-horizon markets. In my experimental markets, traders receive noisy private signals of asset fundamentals. They also make predictions about the future path of asset prices. Before trading begins, I use a Bayesian inference task to gauge two types of individual overconfidence commonly cited in the behavioral finance literature - miscalibration, and the "better than average effect." This task is the same signal-extraction problem by which dividends can be inferred from forecasts during the asset market. On average, subjects are overconfident according to both of these measures; also, the two measures are somewhat correlated at the individual level.

I find that traders do trade based on private information; traders who get better-than-average forecasts tend to be buyers, while traders who get worse-than-average forecasts tend to be sellers. Adding private signals causes trade, in violation of the No Trade Theorem. Also, traders' inferred reservation prices depend on their signals to a greater than optimal degree. These findings support the modeling approach used in overconfidence-based models.

However, I find only weak evidence that the overuse of private information is caused by *psychological* overconfidence. Traders who are miscalibrated, or who consider themselves "better than average" guessers of the random processes underlying their private signals, are no more likely than others to use their signals to make trading decisions. Miscalibrated traders may set their reservation prices more in accordance with their signals than do other traders, but the evidence is not extremely strong.

Since I have data on overconfidence and on market outcomes, I revisit the direct link between psychological overconfidence and market outcomes. I find that overconfident traders do not trade more, and actually may earn higher profits than other traders. When evaluating aggregate outcomes across market trading groups, my sample size is very small; however, I observe an apparent correlation between trading volume and the average "better than average effect" of a trading group. However, I find no such relationship between aggregate confidence and either asset overpricing or price volatility.

Thus, I find only mixed evidence to support the link between psychological overconfidence and the kind of market phenomena that overconfidence is generally invoked to explain. Since my overconfidence measures are highly specialized, it may be that other measures work better to explain aggregate outcomes. However, it is clear that psychological overconfidence on signal-extraction tasks does *not* explain the overweighting of private signals that I observe in this experiment.

My findings also have an interesting implication for the "bubble experiment" literature that follows Smith, Suchanek and Williams (1988). Overweighting of private information does not disappear, and actually strengthens, over successive repetitions of the asset market. This is in stark contrast to the behavior of asset bubbles, which typically disappear by the third repetition. By focusing on bubbles, asset pricing experiments may be ignoring a more enduring source of market inefficiency.

Section II reviews the relevant experimental literature, the literature on overconfidence-based theories of asset markets, and some psychological literature on overconfidence. Section III describes the experimental procedures for the market and for the overconfidence tests. Section IV motivates and states the hypotheses linking the experiment to existing models of overconfidence. Section V presents and discusses the results of the experiment. Section VI concludes.

2 Related Literature

2.1 Overconfidence models

Overconfidence models typically attempt to explain one or more of three phenomena: overpricing of financial assets, overtrading (or high volume during periods of overpricing), and excess price volatility.

The first theories of asset markets that could reasonably be called "overconfidence models" are models of heterogeneous beliefs. In the seminal paper of Harrison and Kreps (1978), traders have different beliefs about future asset fundamentals, which do not disappear over time; furthermore, these traders "agree to disagree," disregarding the beliefs of other traders who disagree with them, even though these differing beliefs are known to all. In this model, in which short selling is not possible, asset values exceed even the most optimistic traders' beliefs about the fundamental, since assets also have an option value; as beliefs fluctuate, assets can be sold and resold at a profit. This option value turns out to be a common result in overconfidence models.

Some later models, such as that of Morris (1996), relax the assumption that agents are dogmatic in their differences of opinion, allowing traders to learn over time. Because learning takes time, results similar to that of Harrison and Kreps (1978) occur in the interim.

Another strand of model posited the existence of "noise traders" who do not trade based on fundamentals; the classic model of this type is DeLong, Shleifer, Summers, and Waldmann (1990). These models typically leave noise traders unmodeled, but the modelers conjecture that noise traders might believe in spurious information about fundamentals.

Beginning in the late 1990s, a number of models attempted to unify the idea of heterogeneous beliefs with the psychological literature on overconfidence. These papers modeled overconfidence as a trader's incorrect, persistent belief that the precision of his own private signals about asset fundamentals is higher than it really is. One such model is Odean (1998), in which price-taking traders trade a single risky over a limited time horizon, while receiving private signals about the future payoff of the asset. Although traders in this model can infer the *average* private signal from the market price, they place excessive weight on their own signal, so that they value the asset differently than the market. This generates excessive trading relative to a rational-expectations equilibrium, and results in prices that exceed fundamentals. However, because short selling is allowed, this model does not generate overpricing via an option value; overconfidence causes prices to deviate from fundamental value, but the direction of the deviation is not determined.

A second model along these lines is Scheinkman and Xiong (2003). In this model, traders trade over an infinite horizon, receiving signals in continuous time. Although the signals are public knowledge, each trader thinks of one of the signals is his "own" signal. *All* signals are actually worthless, but each trader believes that his own signal (and no other) is correlated with innovations to the asset's payoff. This model is similar in spirit to Harrison & Kreps (1978), except that instead of stubbornly disagreeing about the process governing asset fundamentals, traders disagree only "agree to disagree" about the informativeness of the various signals. Again, with short selling impossible, option value creates a "bubble,"[1] in which prices exceed any trader's expectation of fundamentals. Also, heterogeneous beliefs lead to excess volume and excess volatility.

Other, similar "behavioral" models of overconfidence proceed along similar lines. Daniel, Hirshleifer, and Subrahmanyam (2001) obtain similar results for a portfolio of many securities. Daniel, Hirshleifer, and Subrahmanyam (1998) model overreaction and underreaction to news about fundamentals as functions of slowly changing overconfidence and self-attribution bias. Gervais and Odean (2001) also model over-

[1]Note that this model does not generate a runup-and-crash in asset prices; thus, it does not generate a "bubble" in the sense of the experimental literature reviewed here.

confidence that changes due to self-attribution bias. Hirshleifer and Luo (2001) model how overconfident traders can survive in a market filled with traders with differing levels of confidence, without being competed out of the market.

2.2 Psychology literature on overconfidence

Several types of overconfidence have been found in psychology experiments. Two receive special attention in the behavioral economics literature. The first is *miscalibration*, in which people are found to overestimate the precision of their knowledge (Lichtenstein, Fischhoff, and Phillips 1982, Yates 1990). This is typically measured with a confidence interval test: subjects are asked a number of questions, and along with their answers they are asked to give a self-predicted confidence interval (usually a 90% confidence interval). If fewer than 90% of their answers lie within the predicted intervals, subjects are said to be miscalibrated. The second type of overconfidence is the *better-than-average effect*, which documents that most people see themselves as better than average in ability (Taylor and Brown 1988, Svenson 1981). This is usually measured by asking people if they are better than the average at some task, e.g. driving a car. If more than 50% claim to be better than average, it signals that overconfidence exists in the population.[2] Many studies find that overconfidence is prevalent among the population; however, on tasks that are easy or repetitive, people may be well-calibrated or even underconfident (Pulford and Colman 1997, Kahneman and Riepe 1998).

2.3 Empirical studies

Odean (1999) examines data on individual investor portfolios provided by a discount brokerage. They find that stocks sold by individual traders tend to subsequently outperform stocks bought by these traders; in other words, individual investors lose money by trading too much. Odean attempts to control for liquidity demands, tax-loss selling, rebalancing, and changes in risk aversion, and finds that once these factors are controlled for, over-trading lowers individual investor returns by an even greater amount. Barber and Odean (2001) find that this over-trading is correlated with gender; men trade more and thus perform worse than women. Since men are often found in psychology experiments to be more overconfident about their abilities than women, this may be evidence in support of the overconfidence hypothesis.

[2] Of course, the fact that an individual considers himself to be better than average does not mean that individual is overconfident; she may simply have good information that she is, in fact, better than average.

Glaser and Weber (2003) test the relationship between individual investor confidence and trading activity. Using a questionnaire distributed to a large number of investors, they measure miscalibration (using self-reported confidence intervals) and the better-than-average effect (using self-reported evaluations of investing skills). They find that miscalibration is not correlated with trading volume, but that the better-than-average effect is positiviely correlated with trading volume.

2.4 Asset market experiments

The classic asset market experiment is the "bubble experiment" of Smith, Suchanek, and Williams (1988) (henceforth "SSW"). In that experiment, small groups of subjects traded a single short-lived risky asset that paid a stream of dividends. Even though the dividend process was made public to all traders, the outcome was a large bubble, in which the price of the asset diverged strongly from the fundamental value and then crashed at the end of the market. This bubble was accompanied by high trading volume at the peak. When the same group of subjects repeated the market several times, the bubble progressively shrank and eventually disappeared; in other words, the outcome approached a rational expectations equilibrium as subjects became more experienced. This result has become the most widely replicated result in experimental finance; in the next two decades, bubbles were shown to be robust to many changes in market institutions and asset fundamentals (King et al. 1993). Because this market setup is so well-studied, I use it in the present experiment.

The experiment most similar to the one in this paper is Kirchler and Maciejovsky (2002). These authors use a setup similar to SSW, but with several modifications. First, traders are asked to predict the price of the risky asset in the upcoming period. Second, traders receive signals about the *other* traders' average predictions. Third, traders are asked to give *confidence intervals* about their price predictions. The authors use the confidence intervals to measure overconfidence (miscalibration). They find that traders are generally well-calibrated at the beginning of the experiment, but become more overconfident as they make successive predictions.[3] Interestingly, Kirchler and Maciejovsky find no correlation between overconfidence and trading volume, although they do find that overconfident traders earn less on average. The authors do not test the relationship between overconfidence and use of private information.

Other experiments use market setups quite different from the SSW setup. Deaves, Luders, and Luo (2003) use twelve single-period markets, where endowments are reset

[3] However, using a second measure of confidence relying on traders' subjective certainty about their predictions, the authors find that subjects are often well-calibrated or even underconfident.

at the beginning of each market, and payoffs occur at the end of the market. The authors give traders i.i.d. signals of fundamental value, but deceive their subjects by telling them that the signals' precision is correlated with subjects' performance on a general knowledge questionnaire. They measure overconfidence by asking subjects to give subjective confidence intervals on the questionnaire. The authors find that overconfidence (positive miscalibration) is positively correlated with trading volume. They also find a positive relationship between miscalibration and the "better than average" effect, as measured by subjects' self-predicted rankings of total profits.

Biais et al. (2005) use a trading setup with asymmetric information about asset fundamentals, measuring overconfidence (miscalibration) with a confidence interval task. They find no significant relationship between miscalibration and trading volume, although they find that miscalibrated subjects earn lower profits.

Finally, a very interesting experiment by Hales (2009) investigates the closely related question of whether traders "agree to disagree" about the value of an asset. In his experiment, subjects trade in pairs and receive private signals about asset value. Hales finds that whether traders are prompted to consider the adverse selection problem has a strong effect on whether trade occurs. When traders are asked to guess the difference between their own signal and the signal of the other trader, trade tends not to occur; however, without such prompting, trade does tend to occur. This result suggests that over-reliance on private information is not due to traders "agreeing to disagree," but simply to their failure to consider the fact that others have information that may disagree with their own.

2.5 Differences from previous literature

The present experiment differs from the existing literature in several ways. First, it directly tests agents' use of private information about fundamentals. Traders in this study receive private signals about asset payouts, of a form very similar to the type of signals in models like Odean (1999). Unlike in previous experiments on overconfidence, this information is both ex ante symmetric, and contains only information about fundamentals.[4]

Second, this experiment is the first to measure overconfidence using a test that is identical to the signal-extraction problem faced by traders in the actual asset market.

[4]Some financial market experiments have established the result that markets aggregate information that is distributed among different traders (Plott and Sunder 1988). However, these experiments have typically not given subjects noisy information, and have typically not examined whether subjects' individual behavior reflects an underestimation of the amount of noise in their own information relative to the noise in others' information.

In the psychology literature, "overconfidence" is measured by many different types of tasks, and there are often large discrepancies between these measures. The present study attempts to control for these possible sources of systematic error by using the same task for the confidence test and the asset market.

Third, this experiment uses a market setup - the "bubble experiment" or SSW setup - that is known to almost always converge to a rational expectations equilibrium by the third market repetition.

Fourth, this study measures confidence using a task that is very unlikely to be subject to overconfidence. The signal-extraction task used here is repetitive and simple, and feedback is swift - all characteristics that are found to be less conducive to overconfidence (Kirchler and Maciejovsky 2002). Thus, any observed overconfidence on this task probably represents a lower bound on the degree to which a subject overestimates her own predictive accuracy and relative performance.

3 Experimental Setup

3.1 Subjects and compensation

The experiment was conducted on four days, between June 3 and June 7, 2012, at the Institute for Social Research at the University of Michigan. Subjects were de-identified, so no individual data is known except that all were over 18 years old. The majority of subjects were probably undergraduate students at the University of Michigan. On each day, there were 12 experimental subjects in the lab, divided into two trading groups of 6 subjects each. Data from two trading groups are excluded from the analysis due to procedural errors by the experimenters, so that there are a total of 36 subjects in the sample.[5] The entire experiment lasted approximately two and a half hours. Ex ante average compensation[6] was $38, including a $5 show-up fee. Realized average compensation over the five days was $35.79. This is in line with the typical per-hour compensation for this type of experiment. Compensation was divided among the preliminary tasks and the financial market itself, as will be explained in the following subsections. Payment for all parts of the experiment was given in cash at the end of the Asset Market portion of the experiment.

[5]However, one of these two excluded groups completed the confidence test with no procedural errors; this group's confidence test data is valid, though its asset market data is not.

[6]Since some compensation depended on the accuracy of subjects' price predictions during the financial market, the $38 estimate was to some degree a guess on the part of the experimenter, based on trial runs of the experiment.

3.2 Experimental procedures

The instructions for the experiment are presented in Appendix A. The experiment consisted of four parts:

1. A financial literacy test
2. A time-series prediction task
3. A confidence test
4. The asset market itself

The purpose of the financial literacy test and time-series prediction task was twofold: A) to measure subjects' basic competence (since nothing was known about the subjects' experience or ability), and B) to induce exogenous variation in group confidence levels. For this reason, the test and the time-series prediction task were administered in two different versions, an "easy" version and a "hard" version (subjects were not told that two versions existed). The differences between these versions will be explained later. Unfortunately, however, the different versions failed to induce measurable differences in confidence, as will be explained later. The version of the preliminary tasks received by each trading group is displayed in Table 1.

All portions of the experiment were carried out using the z-Tree software package (Fishbacher 2007).

3.3 Financial literacy test

The first portion of the experiment was a five-question financial literacy test. This test was based on the literacy test in Rooij, Lusardi, and Lessie (2011). The "easy" version of the test consisted of the five "basic literacy questions" created by Rooij et al., while the "hard" version of the test consisted of five of the eleven "advanced literacy questions." The text of these questions is given in Appendix C. Subjects were given one minute to answer each question, and were told the correct answer after answering each question.

Each subject received $2 for taking the financial literacy test, regardless of score.

3.4 Time-series prediction task

The second portion of the experiment was a "prediction task". Subjects were asked to predict the future value of a time series (labeled "GDP") from past values of the time series, using no other information other than the series' past values. The time series was generated before the experiment by an AR(1) process with Gaussian white noise innovations and a linear time trend. At the beginning of the task, subjects were shown a pre-existing ten-period "history" of the time series, which they used to make their first one-period-ahead prediction. After each prediction, the realized value of "GDP" for the next period was revealed, along with the subject's prediction error; subjects would then make another one-period-ahead prediction, while seeing the initial ten-period history and all subsequent realizations of the time series. Each subject made ten predictions in all. At the end of the task, the subject was told their total prediction error.

The "easy" version of the time series had higher persistence and a smaller variance of the innovation term than the "hard" version; the time trends were the same. The values of the "easy" and "hard" time series are listed in Appendix D.

Each subject received $2 for completing the prediction task, regardless of score.

3.5 Confidence test

The confidence test was presented to the subjects as "Prediction Task 2". In this task, subjects were asked to guess whether a pair of "balls" was drawn from "Jar A" or "Jar B". Each ball was identified by its color; these colors could take on five values, red, orange, yellow, green, and blue. Subjects were told the number of balls of each color in Jar A and Jar B, respectively. Each subject made 15 guesses; the number of balls of each color in each jar did not change between guesses, indicating that the draws were i.i.d. The number of balls of each color in each jar are listed in Table 2.

Subjects were not told whether their guesses were right or wrong. After making 15 guesses, subjects were asked to predict how many correct guesses they had made (out of a possible 15). They were also asked to predict their rank among the 12 subjects in the lab on that day, so that a rank of 1 would indicate that a subject made as many or more correct guesses than any other subject in the lab. Even after making these guesses, subjects were not told how many correct guesses they had made, or their true rank. These two self-predictions were used to generate the two confidence measures used in the experiment. The difference between a subject's guess of his or her number of correct answers and the true number he or she got correct is defined as that subject's "miscalibration," while a subject's predicted ranking defines his or her "better than average effect" (or "BOA").

This task was a simple Bayesian inference task. It was essentially identical to the problem of guessing the true per-period dividend from a private forecast during the asset market (see next subsection), though the probabilities were different. This is important, because individuals may have different levels of confidence for different types of tasks. An individual's confidence on a Bayesian prediction task may be a good or a poor proxy for her confidence on, say, a market-timing task. To my knowledge, this experiment is the first to measure confidence using exactly the same kind of signal-extraction problem that is required of traders in the asset market.

3.6 Asset market

The asset market portion of the experiment closely follows the setup of Smith, Suchanek, and Williams (1988). Each market consisted of a trading group of 6 subjects. The assets traded included one riskless asset ("cash," measured in "yen"), and one risky asset (called simply "the asset," measured in "shares").

Each group of subjects traded over three market *rounds*. The subjects in each group were the same from round to round. At the beginning of each market round, each subject was given an initial endowment of cash and shares. There were three types of endowments, each with identical ex-ante value; in each market, two traders in each group received Type 1 endowments, two received Type 2 endowments, and two received Type 3 endowments; the allocation of endowments was randomized given this constraint.

Each market round included ten two-minute trading *periods*. During a period, subjects could buy and sell the asset from each other. At the end of a period, each share of the asset paid a *dividend* to its holder; this dividend was either 0 yen or 10 yen, with a 50% probability of each. Dividends were the same across all shares within a period, but different from period to period and different across groups. Dividends were i.i.d. At the end of each period, traders were told the realized dividend for that period, and the appropriate dividend income was added to traders' cash accounts[7]. After the final dividend was realized at the end of Period 10 in each round, the asset vanished (no buyout value).

At the end of a market round, subjects received the promise of an amount of cash equal to the amount of cash in their account, to be paid after the end of the experiment. The exchange rate was 1 "yen" = 1¢. Thus, there was no "carry-over" between rounds; whatever a subject earned or did during Round 1 could not affect his or her expected payoff in Round 2 or Round 3.

Trading was done via continuous double auction with an open order book. Each subject was able to enter bids or asks, or to accept the best available bid or ask. In addition, to facilitate rapid trading, any subject could press a button that submitted

[7] Thus, the cash/asset value ratio of the market as a whole increased each period.

a bid (ask) that was 1 yen higher (lower) than the highest (lowest) available bid (ask). Shares could only be bought and sold one at a time. During trading, subjects were told the "buy-and-hold value" of their current portfolio of risky asset shares; this was equal to 5 yen multiplied by the number of periods remaining, multiplied by a subject's number of shares. In other words, this "buy-and-hold value" was equal to the ex ante risk-neutral fundamental value of the asset. However, this value did *not* include the information in subject's forecasts (see following subsection).

3.6.1 Forecasts

At the beginning of each period (before the start of trading), each subject received a forecast (noisy signal) of that period's dividend value. There were five possible forecasts: "extremely bullish," "moderately bullish," "neutral," "moderately bearish," and "extremely bearish." Forecasts were i.i.d. across subjects. The probabilities of receiving the different forecasts were different depending on the different upcoming (but not yet observed) dividend realization; for example, if the dividend was to be 10 yen, the probability of a subject receiving an "extremely bullish" forecast was 35%. Subjects were given a hard copy of a chart telling them these conditional probabilities; guessing the dividend from the forecast was thus a matter of Bayesian inference, identical in form to the guessing task in the earlier Confidence Test. The forecast probability chart can be seen in Appendix B.

Henceforth, to avoid confusion, these dividend forecasts will be referred to as "forecasts," and traders' own predictions of future asset prices (described in the next subsection) will be referred to as "predictions."

3.6.2 Price predictions

Before trading in each period, subjects were asked to predict the future average trading price of the risky asset. Specifically, they were asked to predict:

1. the price in the upcoming trading period, and
2. the price in the final period of the round (Period 10).

Subjects were paid for making accurate predictions, according to an absolute deviation formula with a minimum of zero. Specifically, for each of the two price predictions in each period, a subject received a payment equal to max(10 yen −

$|predictedvalue - actualvalue|, 0)$. Thus, perfect predictions in all periods in all three rounds would yield a total payment of 600 yen, or \$6[8].

3.6.3 Control group

One group (Group C) was not shown any dividend forecasts. This group will be referred to as the "control group."

4 Theory and Hypotheses

4.1 Overconfidence measures

In the behavioral finance literature on overconfidence, two psychological phenomena - miscalibration, and the better-than-average effect - are commonly used to motivate the agents' overweighting of private signals. However, it is not clear that either one of these phenomena, *by itself*, is a necessary or sufficient condition for overweighting of private signals. To see why, let's examine the model of Odean (1998), which is most similar to the setup used in the present experiment. In Odean's model, each agent is aware of three pieces of noisy information: 1) the prior distribution of asset payouts, 2) the trader's own signal of asset payouts, and 3) other traders' signals of asset payouts. The agent's belief about the precision of each of these signals differs from the true value. The behavioral multiple that an agent assigns to his belief about the precision of his own signal is called κ and is assumed to be greater than or equal to 1, meaning that agents believe that their own signals are more precise than they really are. The multiples that agents assign to the prior distribution of payouts and to the signals of others are called η and γ, respectively, and are both assumed to be less than or equal to 1, meaning that agents believe that the prior distribution and other agents' signals are *less* precise than they really are.

However, in this experiment, it cannot be assumed a priori that these sign restrictions hold. Suppose a subject demonstrates miscalibration on a test. It may be that this person overbelieves in the precision not just of her *own* signals, but of *all* signals; in other words, that $\kappa = \eta = \gamma > 1$. In the Odean (1998) model, this would cause traders to have common beliefs, which would preclude trade and reverse the results

[8] It is theoretically possible for subjects to intentionally make predictions that are not unbiased functions of their actual beliefs, even given this symmetric loss function. This would occur if subjects attempt to make market-timing bets, e.g. betting that the price will go up, but predict a lower price for the purpose of hedging. However, previous studies, such as Haruvy, Lahav, and Noussair (2007) indicate that such behavior is negligible in these experiments.

of the model. Similarly, suppose a subject demonstrates that he or she believes him or herself to be better than average. It may be that this person *under*believes in the precision of her own signals, and simply underbelieves even more in the precision of other signals; in other words, $\gamma = \eta < \kappa < 1$. This is a general problem with these overconfidence measures, and not specific to the Odean (1998) model, which is used simply to illustrate the point. It is *possible* that either or both of these observable confidence measures is a good proxy for the overweighting of private signals, but this is not *guaranteed*.

Therefore, when testing hypotheses about the effects of overconfidence, I use three measures. The first are miscalibration and the better-than-average (henceforth "BTA") effect, as measured by the two parts of the confidence test described above:

$$MISC_i \equiv \text{self-predicted \# of correct guesses - \# of correct guesses} \tag{1}$$

$$BTA_i \equiv \text{6.5 - self-predicted rank} \tag{2}$$

The third is a dummy variable, D_CONF_i, equal to 1 if a subject both is positively miscalibrated and predicts that she is better than average (and zero otherwise). In other words, this dummy is 1 when a subject is overconfident according to both common psychological measures:

$$D\_CONF_i \equiv \begin{cases} 1 & MISC_i > 0, BTA_i > 0 \\ 0 & otherwise \end{cases} \tag{3}$$

Note that if $MISC_i$ and BTA_i are increasing functions of $\kappa - 1$ and $\kappa - \gamma$ in the Odean (1998) model, respectively, and if $\eta = 1$, then $D\_CONF_i = 1$ is sufficient to preserve all relevant results of that model (proof omitted); again, this is for illustrative purposes only, since the experiment is not an exact test of this or any other existing model. In most of the analysis, I examine only miscalibration and BTA; however, in one or two cases, D_CONF_i yields slightly different results.

4.2 Use of private signals

The first question this paper seeks to answer is: Do traders overweight private signals when they make their trading decisions? This is the mechanism by which models of heterogeneous beliefs generate market inefficiencies. It is also the mechanism by which behavioral models like Odean (1998) and Scheinkman and Xiong (2003) assume psychological overconfidence to operate.

A necessary condition for the interaction of individual overconfidence with private signals is that traders make use of private signals at all. In other words, the dividend forecasts traders receive in the asset market must cause trade where trade would not otherwise have occurred, thus violating the No Trade Theorem. Within a trading group within a period, the distribution of forecasts is i.i.d. Thus, if traders do not use their forecasts to decide whether to buy or to sell, the probability of the buyer getting a forecast of 2 (the best possible) and the seller getting a forecast of -2 (the worst) is the same as the probability of the buyer getting a forecast of -2 and the seller getting a forecast of 2. Thus, theory, in the form of the No Trade Theorem, tells us that the expected value of the difference between the forecasts of buyers and sellers should be zero:

Hypothesis 1: The average spread between the forecasts of the buyer and seller in concluded transactions will be zero.

4.3 Private signals and overconfidence

The second question this paper seeks to answer is whether psychological overconfidence is associated with greater overuse of private information. There are two observables that represent trading decisions: Quantity and price. Therefore I examine how overconfidence and private signals interact to affect buying/selling decisions and the prices at which assets are exchanged.

4.3.1 Trading decisions

Overconfidence theories predict that overconfident traders will tend to buy when they receive higher forecasts of asset payouts, and sell when they receive worse forecasts. Therefore, a null hypothesis of no influence of overconfidence on overuse of private signals is the following:

Hypothesis 2: Overconfident traders are no more likely than other traders to buy more of the risky asset when their forecasts are better.

4.3.2 Pricing

Behavioral theory also predicts that overconfidence leads to greater use of private signals when determining a trader's *reservation price*. In this experiment, there is no short-selling; hence, each contract price theoretically represents the reservation price of the buyer, who is the most optimistic trader in the market (Scheinkman & Xiong 2003).[9] In this market, dividends and forecasts are i.i.d.; thus, in any period, the

[9] More realistically, some error on the part of the experimental subjects is probably involved; still, it is reasonable to assume that the contract price is a random variable with the upper limit of

option value of the asset should be the same across traders, and reservation prices should reflect only differences in opinion about fundamentals. Thus, theory predicts that forecasts should have a greater effect on the prices of concluded contracts when the buyers of the contracts are more overconfident. Stated as a null hypothesis, this is:

Hypothesis 3: The correlation between a buyer's forecast and the price paid by that buyer does not depend on the buyer's level of overconfidence.

4.4 Overconfidence and market outcomes

In addition to studying the role of private information and its interaction with overconfidence, this paper also revisits the predictions of behavioral theory about the relationship of overconfidence and market outcomes at the individual and aggregate level.

Behavioral theories predict that overconfident traders trade more frequently than other traders. Thus, the null hypothesis is:

Hypothesis 4: The number of trades concluded by a trader over the course of the experiment is uncorrelated with that trader's measured overconfidence.

A number of experimental and empirical studies have found that some measures of overconfidence are correlated with lower earnings in financial markets. This is predicted by some theories, although others predict that overconfident traders earn equal or higher returns and lower expected utility. In any case, the null hypothesis I test is:

Hypothesis 5: Overconfidence is not related to a subject's earnings in the financial market portion of the experiment.

At the aggregate level, overconfidence-based theories generally predict that markets in which a large number of traders are overconfident will exhibit one or more of the following outcomes: higher volume (due to over-trading), increased price volatility (due to fluctuating heterogeneous beliefs), and prices above fundamentals (due to the option value associated with fluctuating heterogeneous beliefs). Thus, using only the small sample of trading groups available, I attempt to test the following null hypotheses:

Hypothesis 6: Trading groups that are more overconfident on average will exhibit higher trading volumes.

Hypothesis 7: Trading groups that are more overconfident on average will exhibit greater overpricing of the asset relative to fundamentals.

---

support at the buyer's reservation price, and is hence is correlated with the unobservable reservation price.

Hypothesis 8: Trading groups that are more overconfident on average will exhibit greater within-period price volatility.

The experimental setup used in this experiment is believed to converge as close as possible to a rational expectations equilibrium by the third market round. Therefore, when possible, I test these hypotheses both for the market as a whole and for the third market round only.

5 Results and Discussion

5.1 Description of the data

5.1.1 Overconfidence

On the type of Bayesian inference task used for the confidence test, answers are either correct or incorrect; thus, subjective confidence intervals cannot be given for each answer. However, an equivalent measure of miscalibration is a subject's own estimation of how many guesses he or she got right. If the subject's subjective confidence intervals for his or her guesses are too small on average, she will tend to predict that she got more questions correct than she actually did.[10]

On average, subjects are miscalibrated. Define $CORRECT_i$ as the number of times (out of a possible 15) that subject i guessed the correct "jar" in the confidence test. Define $SELFPRED_i$ as that subject's guess for his or her total number of correct guesses. Then define $MISC_i$ as the difference of the two:

$$MISC_i \equiv SELFPRED_i - CORRECT_i \tag{4}$$

The average value of $MISC_i$ is 1.72, with a standard deviation of 2.68. A total of 26 out of 36 subjects are positively miscalibrated by this measure, while 6 are negatively miscalibrated and the remaining 4 are well-calibrated. A Wilcoxon sign-rank test rejects the null of $MISC_i = 0$ with a p-value of 0.0003. Hence, on average, subjects are positively miscalibrated.

On average, subjects also predict that they are better than average guessers. Define BTA_i ("better than average") as the difference between 6.5, the correct ex ante average rank of all subjects (1=best, 12=worst), and the rank that a subject predicts for his or her own guessing accuracy (out of all subjects present in the lab).

[10] An alternative measure of miscalibration would be to compare a subject's predicted number of correct answers with the ex ante expected number of correct guesses made by a guesser who made the proper inference every time. However, this would not take into account the subject's knowledge of his or her own ability to make correct inferences.

The average value of BTA_i is 1.39 (standard deviation of 1.74). A Wilcoxon sign-rank test rejects the null of $BTA_i = 0$ with a p-value of 0.0001.[11] Thus, on average, subjects suffer from the "better than average effect."

There is some evidence that the two measures of overconfidence are correlated across individuals. An OLS regression of $MISC_i$ on BTA_i with a sample size of 36 fails to reject the null of no relationship (p-value 0.184). However, there are an additional 6 subjects who successfully completed the confidence test even though their asset market data is excluded due to procedural error. When these additional subjects are included in the sample, a regression of $MISC_i$ on BTA_i yields a p-value of 0.085, showing that the two measures are probably weakly correlated (the correlation coefficient is 0.269).

22 of the 36 subjects are overconfident by both measures ($D\_CONF_i = 1$), and 24 out of 42 in the expanded sample. A two-sided test cannot reject the null that $D\_CONF_i = 0.5$ (p-value 0.182); in other words, there is not sufficient evidence to conclude that subjects, on average, suffer from *both* miscalibration and the better-than-average effect.

Because this task is repetitive and simple, the psychology literature predicts that subjects will be less overconfident on this task than on other tasks. Hence, it is reasonable to assume that subjects will tend to be even *more* overconfident on other tasks - for example, market timing, assessing the rationality of other traders, and understanding the experimental setup itself.

5.1.2 Bubbles

The prices in these markets clearly display the classic "bubble" result. Figure 1 displays price histories for the three market rounds, along with the risk-neutral fundamental value (unconditional on forecasts), and the maximum possible dividend value. Prices are measured as the difference between the realized contract price per share and the risk-neutral fundamental value, which in turn is defined as the expected total dividend payment per share.[12] The general pattern of decreasing overpricing over successive market rounds is clearly visible. Stockl et al. (2010) suggest a mea-

[11] It is possible that subjects incorrectly imagined the ex ante average rank as 6 rather than 6.5. Define BTA_i' as 6 minus a subject's predicted rank. A test of the null of $BTA_i' = 0$ rejects the null with a p-value of 0.002.

[12] When forecasts are visible, i.e. for all groups except C, the expected dividend payment includes the true dividend value (0 or 10) for the current period rather than the expected value (5). This assumes that forecasts, in aggregate, are perfectly revealing of the upcoming dividend. In fact, this is an approximation, but only a slight one.

sure of overpricing called "Relative Deviation," defined by the following:

$$RD_m \equiv \frac{10\sum_t P_{m,t} - FV_{m,t}}{\sum_t FV_{m,t}} \tag{5}$$

Here, $P_{m,t}$ is the average realized contract price in period t of market m, and $FV_{m,t}$ is the risk-neutral fundamental value. The Relative Deviation for all trading groups in this experiment decreases from Round 1 to Round 2, and again from Round 2 to Round 3. The Relative Deviations for each market round can be seen in Table 3. Furthermore, in Round 3, the Relative Deviation is very small or negative for all groups except one, indicating that bubbles generally disappear by the third trading round. These facts are all in agreement with the rest of the bubble experiment literature.

For this reason, Round 3 of these markets probably presents the closest approximation to a Rational Expectations Equilibrium achievable in an experimental setting.

It is clear from theory that the bubbles observed in many of these markets, especially Round 1 markets, cannot possibly be explained by overconfidence about the dividend forecasts - or, in fact, about any signals of fundamental value that are actually present in this experiment. This is because in many of these markets, price exceeds the maximum possible dividend payout. Overconfidence about signals of fundamental value can never lead to this outcome.

Proof. Suppose that all traders believe their private information to be perfectly revealing of the current-period dividend in every period. Also, suppose each trader expects her private information to be fully revealing of dividend value in all future periods. Furthermore, suppose that in every period, each trader believes that some other trader will receive an incorrect signal of fundamental value, and will believe this signal with certainty. In the final period, this situation will lead a trader to expect to be able to complete an arbitrage trade with a value equal to the difference between the maximum and minimum possible dividend realizations - i.e., 10 yen. By backwards induction, this adds an option value of 10 yen to the asset's value to any trader in the previous period. But in the previous period, a similar arbitrage trade will also be available. Thus, the maximum that any trader is willing to pay for the asset in a given period is given by the maximum possible dividend (10 yen) plus 10 x the number of subsequent periods. This is equal to 10 x the total number of periods remaining, which is also equal to the maximum possible dividend payout. ■

Thus, the observed bubbles require the presence of some other wedge between price and fundamental value, in addition to overconfidence. This could be provided by subjects' initial misunderstanding of the parameters of the experiment, as suggested by Lei, Noussair, and Plott (2001), or by speculative behavior, as suggested by Smith, Suchanek, and Williams (1988).

5.2 Use of private signals

Define F_{it} as the forecast received by trader i in period t. Recall that there are five possible forecasts, ranging from -2 to 2. Now define ΔF_{ijt} as the difference in forecasts between buyer i and seller j at time t, so that $\Delta F_{ijt} = F_{it} - F_{jt}$. To test whether forecasts are used in trading decisions, I use a Wilcoxon sign-rank test of the null that ΔF_{ijt} across all completed contracts in groups that received forecasts (#obs = 1266). The test rejects the null in the positive direction with a p-value of 0.000. Restricting the sample only to Round 3 - where learning is expected to be complete - the test again rejects the null, with a p-value of 0.0002. The average difference between buyer and seller forecasts in Round 3 is 0.357 forecast categories, which is equal to the average ΔF_{ijt} for all rounds combined. Thus, even in the third round, traders are using the dividend forecasts to decide whether to buy or sell the asset; learning, which eliminates bubbles, does not eliminate this over-reliance on private signals. Note that this result does not show that *all* subjects use forecasts to make their decisions, only that some do.

For comparison, a test of the same hypothesis for the control group, for which forecasts were invisible[13], fails to reject the null of $\Delta F_{ijt} = 0$ (p-value of 0.509).

A small robustness check in order. First, since the trading system used in the experiment only allows traders to trade asset shares one at a time, it is possible to view multiple transactions between a buyer and a seller during a single period as one single transaction. Repeating the above tests using within-period averages of ΔF_{ijt} across buyer-seller pairs, treating each within-period buyer-seller pair as a single observation, yields identical results.

This result shows, in a very stark setting, that injecting private information - "news" - into an asset market induces trade, even when the signals are ex ante symmetric. What's more, the use of private information persists even after bubbles are gone. The result supports theories in which heterogeneous beliefs are induced by excessive belief in the precision of private signals. However, it does not necessarily support the link between psychological overconfidence traits and overestimation of private signal precision. The following subsections deal with that question.

[13] Forecasts for the control group were still *generated* by the same random process, behind the scenes; they were simply not *shown* to the traders.

5.3 Overconfidence and private signals

5.3.1 Trading decisions

To test Hypothesis 2, I use data on subjects' trading activity. The dependent variable is $NETBUY_{it}$, the net asset purchase of trader i in period t. The key regressor is the interaction between the buyer's confidence and the buyer's *differential* forecast $\widetilde{F}_{it}$. $\widetilde{F}_{it}$ is the difference between the forecast of trader i in period t and the average of the forecasts of all other traders in the same group-period:

$$\widetilde{F}_{it} = F_{it} - \frac{1}{5}\sum_{j \neq i} F_{jt} \tag{6}$$

I use the differential forecast instead of the raw forecast for two reasons. The first is that the group-period's average forecast contains information about the dividend itself; dividends may be bunched at the beginning or end of a market round. The second reason is that using the differential forecast more clearly illustrates that adverse selection is being ignored. Although forecasts are discretized, the average forecast is close to the best available forecast of the upcoming dividend; therefore a trader who trades based on her differential forecast is behaving suboptimally. Another way of saying this is: I want to ask the question of whether those traders who get the better forecasts within a group within a period tend to be the traders who buy assets in that period.

Here are the regression equations:

$$NETBUY_{it} = \alpha + \beta_1 MISC_i + \beta_2 \widetilde{F}_{it} + \beta_3 (MISC_i \times F_{it}) + \epsilon_{it} \tag{7}$$

$$NETBUY_{it} = \alpha + \beta_4 BTA_i + \beta_5 \widetilde{F}_{it} + \beta_6 (BTA_i \times F_{it}) + \epsilon_{it} \tag{8}$$

Results for these regressions are in Table 4. Standard errors are clustered at the group-period level.[14] Forecasts have a strongly significant effect on trading decisions, as was shown in the previous subsection. This is true even when the sample is restricted to Round 3. However, neither β_3 nor β_6 is significant. Although traders in general buy more when their forecasts are good and sell more when their forecasts are bad, there is no evidence that overconfident traders do this more so.

This result is quite robust to the inclusion (as control variables) of the cash and asset shares in the subject's portfolio at the beginning of the period. It is also robust

[14]This clustering is due to the fact that the decision of a trader to buy is correlated with the decision of another trader in the same group to sell.

to the subtraction of individual fixed-effects in net buying, and to the scaling of net buying by total trading activity within a period (i.e. using the fraction of a period's net asset purchases rather than the level of purchases). These results are not shown.

The result does not support the theory that overconfidence is responsible for overweighting of private signals.

5.3.2 Pricing

To test Hypothesis 3, I regress the price paid by a buyer with the interaction of the buyer's confidence and her forecast. $BMISC_c$ and $BBTA_c$ are the miscalibration and BTA effect, respectively, of the buyer of contract c.

For this regression I again use the differential rather than the raw forecast. Since all forecasts within a group-period are correlated with the upcoming dividend (and, hence, with each other), forecasts should affect prices even in a rational expectations equilibrium. However, if the traders within a group-period who get better forecasts pay higher prices, it indicates the presence of heterogeneous beliefs that persist even after trade.

Because prices vary from period to period depending on any number of factors (e.g. bubbles), for this regression I use the differential price of a contract $\widetilde{P}_{ct}$, where c is an index specifying a contract and t is the group-period in which the contract occurs. This is defined analogously to $\widetilde{F}_{it}$ above:

$$\widetilde{P}_{ct} = P_{ct} - \frac{1}{NC_t - 1} \sum_{c' \neq c} P_{c't} \tag{9}$$

NC_t is the number of contracts in the group-period.

The two regression equations are:

$$\widetilde{P}_{ct} = \alpha + \beta_1 BMISC_c + \beta_2 \widetilde{F}_{ct} + \beta_3 \left(BMISC_c \times \widetilde{F}_{ct} \right) + \epsilon_{ct} \tag{10}$$

$$\widetilde{P}_{ct} = \alpha + \beta_4 BBTA_c + \beta_5 \widetilde{F}_{ct} + \beta_6 \left(BBTA_c \times \widetilde{F}_{ct} \right) + \epsilon_{ct} \tag{11}$$

Again, one contract is used as one observation.[15] Standard errors are now clustered at the subject level. The results of these regressions are in Table 5. All results are robust to the inclusion of subjects' beginning-of-period portfolios as controls, to

[15] Again, the results are robust to the alternative specification of one observation as the average across all contracts between a specific buyer-seller pair within a group-period.

the subtraction of individual-level fixed effects, and to the averaging of contracts over buyer-periods (results not shown).

First, the results for miscalibration. The coefficient β_3 is significant at the 10% level. The point estimate is 0.317, meaning that a subject who is miscalibrated by 1 unit is expected to pay a price premium of 0.317 yen relative to the other buyers in the group when her forecast is 1 level better (one forecast level corresponds to a difference of either 2.5 yen or 1.25 yen in expected dividend payout). Given that 1 unit of miscalibration is less than one standard deviation of miscalibration (which is 1.72 units), this is a large point estimate for β_3. Note that β_2 is not significant here; well-calibrated buyers do not tend to pay higher prices when their forecasts are higher, at least when averaged over all three market rounds.

However, when the sample is restricted to Round 3 only, β_3 becomes insignificant, and β_2 becomes significant (p-value 0.009). In the final round, when learning is assumed to have occurred, it is no longer only miscalibrated buyers who pay higher prices when they get higher forecasts than the average; it is now *all* buyers.

What this means is that instead of miscalibrated subjects learning to ignore forecasts that are different from the average forecast, well-calibrated subjects are "learning" to pay attention to forecasts when deciding how much to pay for the asset. This is a very interesting result, because it represents *anti-learning.* Recall that when forecasts are different from the group average, they are more likely than not to be *incorrect* forecasts. That well-calibrated traders make this mistake in increasing amounts as the market progresses is very interesting, because it runs exactly counter to the notion that subjects and markets become more rational as the experiment progresses. It shows that while bubbles are an ephemeral phenomenon in laboratory asset markets, over-reliance on private signals can persist and even increase over time.

The estimate for β_6 is not significant; BTA buyers do not pay more than other buyers when they get better forecasts. Oddly, β_4 is negative and significant; buyers who suffer from the BTA effect pay *lower* prices, on average, than other traders. Why? One hypothesis is that BTA subjects are more sensitive to signals of a bubble. However, when I regress price on the interaction of the buyer's BTA level and the true overpricing of the asset, the interaction is not significant, and the main effect remains. Thus, the tendency of BTA subjects to pay lower prices, regardless of dividend forecasts or bubbles, remains a puzzle.

To sum up, I find strong evidence that traders use private signals of fundamentals when they make their trading decisions, exactly in accordance with overconfidence-based behavioral theories like Odean (1998) and Scheinkman and Xiong (2003). However, I find only statistically weak evidence of a link between reliance on private signals and measured overconfidence. I now examine whether the predictions of the behavioral literature hold true with regards to the effect of overconfidence on individual-level and market-level outcomes.

5.4 Effects of individual overconfidence

5.4.1 Trading frequency

Define $NUMTRADES_i$ as the number of trades made by subject i. To determine if overconfident subjects trade more, I regress $NUMTRADES_i$ on $MISC_i$ and BTA_i. The results are in Table 6. There is no observable correlation. This is also the case when the sample is restricted to Round 3. In other words, there is no evidence that overconfident subjects trade more than other subjects. This is consistent with the findings of Kirchler & Maciejovsky (2002) and Biais et al. (2005), though it contradicts the findings of Deaves, Luders, and Luo (2003). The lack of correlation between miscalibration and trading frequency agrees with the empirical study of Glaser and Weber (2003), though those authors do find a relationship between the BTA effect and trading frequency, which I do not.

5.4.2 Performance

Some behavioral theories predict (and some empirical papers argue, e.g. Barber & Odean 2001) that overconfident traders perform worse than other traders. Regressing a subject's total profit Π_i on miscalibration and BTA, I find no relationship between confidence and total profit. However, I find a *positive* relationship between D_CONF_i and profit. Looking at profit earned in Round 3 only, I find no relationship between profit and miscalibration, a positive relationship between profit and the BTA effect, and no relationship between D_CONF_i and profit. These results are in Table 6.

Some models of overconfidence predict that overconfident traders have expected returns equal to or greater than other traders, but lower expected utility, driven by a greater variance in profit. To test this, I define $\Pi SQ_i \equiv (\Pi_i - \overline{\Pi})^2$, where $\overline{\Pi}$ is the average profit across all subjects, and regress this on confidence. However, I find no relationship; overconfident traders do not have a detectably higher variance in their earnings that would lead a risk-averse individual to have lower expected utility from trading (results not shown).

Other experiments, including Kirchler & Maciejovsky (2002) and Biais et al. (2005), found a negative correlation between miscalibration and earnings. However, those studies measured miscalibration on general tasks, while the present experiment measured miscalibration only on a very specific and simple task related to the type of private signals provided in the experiment. Hence, the lack of correlation between miscalibration and profits should not be construed as a contradiction of those studies' findings.

5.5 Effects of average overconfidence

There are only six trading groups in this experiment. This is a very small sample. Hence, the ability of this data set to evaluate the effects of overconfidence on market outcomes is very limited; we should not expect too much action at the aggregate level. However, there is one interesting pattern in the aggregate data.

5.5.1 Volume

Regressing average per-period trading volume on average overconfidence at the trading group level, I find no relationship between average miscalibration and trading volume, but I find a significant positive relationship (p-value of 0.058) between group average BTA effect and volume. The results are in Table 7. The point estimate is large; an average BTA effect of 1 rank is associated with an increase in volume of 8.04 trades per period, equal to almost two standard deviations. The R-squared of the regression is 0.634. The relationship can be seen visually in Figure 2. If the sample were larger, this would be clear evidence that average overconfidence increases trading volume; as it is, the sample is very small and this result should be taken with a grain of salt. Still, if true, this would provide support for the link between trading volume and psychological overconfidence as measured by the BTA effect, found by Glaser and Weber (2003).

Why would there be a relationship between overconfidence and volume at the aggregate level, but not between overconfidence and trading frequency at the individual level? The answer is not clear, but it might be a function of higher-order beliefs. Non-overconfident traders may not assume (as in most models) that other traders are more confident than themselves, and hence may be more sensitive to adverse selection when dealing with overconfident traders, stifling trade.

Alternatively, this result may be a spurious one, given the small sample. When only volume in Round 3 is considered, the relationship between volume and group BTA becomes insignificant (p-value of 0.120). Therefore, this result is merely suggestive, not conclusive, evidence in favor of overconfidence-based theories.

5.5.2 Overpricing and Volatility

I define group overpricing B_g to be the difference between contract price and the asset's risk-neutral fundamental value, conditional on perfect knowledge of the current period's dividend (as justified in the earlier subsection on bubble measurement), averaged across all contracts concluded within a trading group. I define price volatility SDP_g as the within-period standard deviation of contract prices, averaged across all contacts concluded within a group. I find no relationship between these aggregate quantities and overconfidence of either type. These results can be seen in Table 7. Since the sample is very small, this negative result probably does not reveal much.

6 Conclusion

The results of this experiment provide solid support for models of financial markets in which traders rely excessively on their private signals, relative to the information implicit in other traders' actions. Adding ex ante symmetric signals about fundamentals to a common asset market experiment resulted in trade, in which those who received positive signals bought more and those who received negative signals sold more. Signals also affected the reservation prices of traders in the way predicted by behavioral theory.

However, the link between overconfidence and over-use of private signals was less clear. There was some evidence suggesting a link. Miscalibrated subjects had reservation prices that were more sensitive to private signals than those of their well-calibrated counterparts; also, groups where the average trader subscribed to the "better than average effect" tended to have higher trading volume. However, the former result was only weakly statistically significant, while the latter result needs a larger sample size of trading groups before it can be regarded as definitive.

Meanwhile, the rest of the hypotheses of overconfidence-based models were not supported. Overconfident traders did not rely more on their forecasts to decide whether to buy or sell. They did not trade more frequently than other traders. They did not earn lower profits, and may in fact have made *higher* profits.

These findings do not debunk the notion that overconfidence leads to over-trading, overpricing, or excess volatility. Overconfidence can be measured in many ways, and there is no clear consensus as to what sort of psychological test best relates to financial market behavior. But what these findings do show is that over-reliance on private signals is not driven by overconfidence on the very task that agents must complete in order to extract information from their signals. In other words, subjects in this experiment acted as if the noise in their signals represented real information, but the degrees to which they did this was not related to how good they thought they were at separating signal from noise.

There are other possible reasons for overweighting of private information besides psychological overconfidence. One of these is myopia, as conjectured by Hales (2009). Overconfidence tasks deal only with an agent extracting information from her *own* signals. But even if she does this perfectly, bounded rationality or higher-order uncertainty may stop her from properly extracting the information implicit in the actions of *others.* She may not know if other agents are rational, or how overconfident they are. She may find it cognitively difficult to make inferences from their actions. Or she may simply be attentionally limited, concentrating only on what she is doing and not monitoring the offers and trades being concluded by other agents.

The findings of these experiments indicate that economists and psychologists should continue the search for psychological measures that predict financial market behavior. But they also indicate that modelers of financial markets need not rely too heavily on the psychology literature to motivate the idea that traders overuse private information. This modeling assumption looks like a good one.

The results presented here also have important implications for the asset pricing experiment literature. Experimental results like that of Plott and Sunder (1988) have found that markets aggregate information; however, the finding that traders treat noise as information means that experimental markets may aggregate *mis*information in fairly predictable ways. Also, the "bubble experiment" literature begun by Smith, Suchanek, and Williams (1988) has found that simple asset markets of the type used here tend to converge to a rational expectations equilibrium after three market rounds, and some have concluded that once agents learn how the market works, they behave rationally. However, in this experiment, traders' over-reliance on private signals not only persisted but increased as the experiment progressed. This indicates that noise trading may be a much more persistent source of inefficiency in experimental asset markets than the dramatic bubbles that usually characterize the beginning of these experiments.

References

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: 773-806.

Biais, Bruno, Denis Hilton, Karine Mazurier, and Sébastien Pouget. 2005. Judgemental overconfidence, self-monitoring, and trading performance in an experimental financial market. *Review of Economic Studies* 72(2): 287-312.

Daniel, Kent D., David Hirshleifer, and Avanidhar Subrahmanyam. 1998. Investor psychology and security market under- and over-reactions. *Journal of Finance* 56(3): 1839-1885.

Daniel, Kent D., David Hirshleifer, and Avanidhar Subrahmanyam. 2001. Overconfidence, arbitrage, and equilibrium asset pricing. *Journal of Finance* 56(3): 921:965.

Deaves, Richard, Erik Lüders and Guo Ying Luo. 2009. An experimental test of the impact of overconfidence and gender on trading activity. *Review of Finance* 13(3): 555-575.

DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, Robert J. Waldmann. 1990. Noise trader risk in financial markets. *Journal of Political Economy* 98(4): 703-738.

Fischbacher, Urs. 2007. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10(2) 171–178.

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

Glaser, Markus and Martin Weber. 2003. Overconfidence and trading volume. AFA 2004 San Diego Meetings.

Hales, Jeffrey. 2009. Are investors really willing to agree to disagree? An experimental investigation of how disagreement and attention to disagreement affect trading behavior. *Organizational Behavior and Human Decision Processes* 108(2): 230-241.

Harrison, J.Michael and David M. Kreps. 1978. Speculative investor behavior in a model with heterogeneous expectations. *The Quarterly Journal of Economics* 92(2): 322-336.

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

Kahneman, Daniel and Mark W. Riepe. 1998. Aspects of investor psychology. *Journal of Portfolio Management* 24: 52–65.

King, Ronald R., Vernon L. Smith, Arlington Williams, Mark V. van Boening, Mark V. The robustness of bubbles and crashes in experimental stock markets. In R. H. Day and P. Chen, *Nonlinear Dynamics and Evolutionary Economics*. Oxford, England: Oxford University Press, 1993.

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

Lichtenstein, Sarah, Baruch Fischhoff, and Lawrence Phillips. Calibration of probabilities: the state of the art to 1980. In Daniel Kahneman, Paul Slovic, and Amos Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press, 1982.

Milgrom, Paul and Nancy Stokey. 1982. Information, Trade, and common knowledge. *Journal of Economic Theory* 26(1): 17-27.

Morris, Stephen. 1996. Speculative investor behavior and learning. *The Quarterly Journal of Economics* 111(4): 1111-1133.

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

Odean, Terrance. 1999. Do investors trade too much? *American Economic Review* 89(5): 1279–1298.

Pulford, Briony D. and Andrew M. Colman. 1997. Overconfidence: feedback and item difficulty effects. *Personality and Individual Differences* 23: 125–133.

Scheinkman, Jose A., Wei Xiong. 2003. Overconfidence and speculative bubbles. *Journal of Political Economy* 111(6): 1183-1219.

Smith, Vernon, Gerry Suchanek, and Arlington Williams. 1988. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, 56(5): 1119–1151.

Svenson, O. 1981. Are we all less risky and more skillful than our fellow drivers? *Acta. Psychologica* 47: 143–148.

Taylor, Shelley E. and Jonathan, D. Brown. 1988. Illusion andwell-being: a social psychological perspective on mental health. *Psychological Bulletin* 103: 193–210.

Van Rooij, Marten, Annamaria Lusardi, and Rob Alessie. 2011. Financial literacy and stock market participation. *Journal of Financial Economics* 101(2): 449-472.

Yates, J. Frank. *Judgment and Decision Making*. Englewood Cliffs: Prentice-Hall, 1990.

Appendix A: Experimental Instructions

(Note: Instructions that appear in normal type were given in all treatments. Instructions that appear in highlighted font were omitted from the control group.)

General Instructions

This is an experiment about decision-making in financial markets. In this experiment, you will participate in a computerized market, in which you will buy and sell a financial asset. The better investment decisions you make, the more money you take home, so invest wisely!

Before the asset market begins, you will take three tests – a written financial literacy test, and two prediction tests.

After the experiment has begun, we ask that you do not talk to the other test subjects until the entire experiment is over. You may direct any questions to the experimenters.

Payment will be given at the end of the experiment.

Literacy Test

This is a test of your general financial literacy. There are 5 questions. You have 1 minute to answer each question.

The computer screen shows one question at a time. All questions are multiple choice. When you have chosen your answer to a question, click the "**OK**" button. You can't go back once you have gone to the next question.

At the end of the test you will be told your score. You will be paid $2 for taking this test, regardless of your score.

Prediction Test 1

In this test, you will be trying to predict the value of something called "GDP". This "GDP" has absolutely nothing to do with the financial assets that you will be trading later on. Each period's "GDP" is generated by a mathematical formula, and depends somewhat on past GDP values and somewhat on randomness (we don't show you this formula).

On the right hand side of the screen will be a graph and a table, both of which show the history of "GDP." When the test begins, you will see a history of 10 past periods. Looking at this history will help you make your predictions.

On the left of the screen is a box where you enter your prediction for the next period's GDP. When you have entered your prediction, hit "OK". You have 30 seconds to enter each prediction. Try to focus on making the best predictions you can.

After you enter each prediction, you get to see the actual GDP for the next periods, as well as your **prediction error** and a message telling you how accurate your prediction was.

After you make 10 predictions, this test ends. You will get to see your total prediction error – that is, the total difference between your predictions and the real GDP value in those periods.

You will be paid $2 for taking this test, regardless of how well you do.

Prediction Test 2

In this test, you will be trying to guess whether some "balls" are being drawn out of "Jar A" or "Jar B".

Jar A and Jar B each contain different numbers of balls of various colors. The five colors are red, orange, yellow, green, and blue. The numbers of each color ball in each jar are listed on the side of the screen.

Each period, we flip a (computerized) coin to choose which jar to draw balls from. Without telling you which jar we're drawing from, we draw two balls and show you the colors. Based on the color of the balls, you guess which jar (Jar A or Jar B) the two balls came from. Enter your guess and press "**Submit Guess**". There are 15 periods, so you will make 15 guesses. You will get 30 seconds to make each guess.

This time, we won't tell you if you guessed right or wrong each time. At the end of the test, you will be asked to guess how many you got right. Try to guess this as accurately as you can.

You will also be asked to guess your *rank* (out of the total number of subjects in the room) – if you think you got more right than any other subject, enter "1", if you think three people guessed more right than you, enter "4", etc. Again, try to guess as accurately as you can.

You will be paid $2 for taking this test, regardless of how well you do.

The Asset Market

What You Will Be Doing

You will participate in three “Rounds” of trading. In each Round, you will get the chance to buy and sell a simulated financial asset (called “Asset”), using cash that we (the experimenters) give you. You can think of this asset as a mortgage or a bond. Each share of the asset pays some cash each period – this cash is called a “dividend”.

After each Round you will get to keep all the cash that is in your account at the end of that Round. We will record this amount and add it to your payment at the end of the experiment.

You will trade in a market of up to six people (including yourself). We, the experimenters, will pick your market group for you. If there are more than six people in the experiment, you will be divided into two groups.

The unit of currency is called “yen”. During each Round, you will have your “yen” in an account, and the shares of the asset will be priced in “yen”. At the end of each Round, you will be able to exchange all the “yen” in your account for real dollars, at an exchange rate of 1 yen = $0.01 (1 yen = 1 cent). So remember, these “yen” are real money!

How Trading Works

At the beginning of a Round, you will be given some cash and some shares of the asset. Different people will get different starting amounts of each. For example, you might start out with 400 yen and 8 shares of the asset, or 700 yen and 4 shares of the asset.

A Round consists of 10 trading periods. Each period lasts for 2 minutes.

At the end of each period, each share of the asset pays a **dividend**. This dividend is an amount of yen that you get for owning the share. The dividend is random, and is determined each period by a computerized coin-flip. In each period, there is a 50% chance that each share will pay a dividend of 10 yen, and a 50% chance that each share will pay a dividend of 0 yen. The dividends are the same for all shares, but different from period to period.

So the cash in your account after each period is given by:

Your cash after Period T =
Your cash at the beginning of Period T
– Purchase price of all asset shares you purchased in Period T
+ Sale price of all asset shares you sold in Period T
+ Dividend paid in Period T x # of asset shares you hold at the end of Period T

The amount of money you take home from the Round will just be the amount of cash in your account at the end of the Round (that is, at the end of Period 10).

Predictions

In addition to trading during each period, you will get a chance to make two predictions before the period. You will be asked to predict the **Average Asset Price in the Upcoming Period**, and the **Average Asset Price in Period 10**. These are the average prices that you think a share of the asset will trade for in those periods. (Note: In Period 10, the final period, these two will be the same thing.)

In addition to your earnings from the market, <u>you will be paid for how accurate these predictions are</u>. For each prediction, you will get 10 yen – |your prediction – the actual value| (with a minimum of zero). So the better your predictions turn out to be, the more you get paid.

Remember, the <u>other traders will never see your predictions</u>. So just try to make the most accurate predictions you can.

Analyst Forecasts

To help you make your decisions about buying and selling the asset, before each period (but after you make Predictions), each of you will receive an **Analyst Forecast**. This is a forecast regarding the dividend that the asset will pay in the next period. The forecasts may be different for each person, and no one else will be able to see your forecast.

The Analyst Forecasts are not perfect. The forecasts can be one of five types: “extremely bullish,” “moderately bullish,” “neutral,” “moderately bearish,” or “extremely bearish”. (Note: in financial slang, “bullish” means “optimistic” and “bearish” means “pessimistic”!)

Which forecast you see depends somewhat on random chance (since each person’s forecast is different), but also depends on the actual dividend that is going to be paid in the next period. We have provided you with a sheet, labeled “Forecast Chart”, that tells you how the probabilities of each type of forecast depend on the probabilities of the next period’s dividend.

The Trading Software

The trading software you will be using is called z-Tree. We will show you how it works.

You will always see the following information at the top of the screen:

Period ← This is the current trading period. If it is between periods, this is the number of the next period.

Time Remaining ← This is the number of seconds left in the period.

On the right hand side of the screen, you will see your **Account Information**:

Cash ←This shows how much cash you currently have in your account. Cash is denominated in “yen”.

Shares ← This shows how many shares of the asset you currently own. This is also called your “portfolio”.

Buy-and-Hold Value ← This is the expected (average) amount of money you would get for all the shares in your portfolio, if you held onto those shares until the end of the experiment. (<u>Note: This value does NOT take your most recent Analyst Forecast into account!</u> It is the buy-and-hold value with ALL dividend percentages set to 50-50, including the current period’s dividend.)

Below your account information, you will also see a **Price History box**, which tells you about asset prices in previous periods:

Average Price ← This shows the average price of a share of the asset in each of the previous periods in the Round
(empty in Period 1)

Below the Price History box, you will see the **Traded Price box**. This lists the sale prices of every share sold in the <u>current</u> period.

During trading, you will have two columns: the **Sell Column** on the left, and the **Buy Column** in the center.

Sell Column

There are two ways to sell assets: 1) You can make a sell offer that other traders can choose to accept, and 2) You can accept bids (buy offers) made by other traders. In the Sell Column, you will see the following three boxes:

Make Sell Offer ← This lets you make an offer to sell one share of asset. Type in the price you are willing to accept for a share, and press "**Enter**". This price must be a whole number (no decimals!), and must be positive. When you enter a sell offer, your offered price will appear in the "Offer Price" box in the Buy Column in the lower middle of the screen, where buyers can choose to accept it.

Undersell Lowest Offer ← Pressing this button will enter a sell offer that is 1 yen lower than the lowest existing sell offer. It is a way to make an offer quickly. If there are no existing sell offers, then this button won't do anything.

Bid Price box ← This box contains a list of all outstanding bids. You can sell a share by pressing the "**SELL NOW**" button at the bottom of the box. You will automatically sell one share for the highest available bid price (as long as it's not your own bid!). Each time you sell a share, ALL of the existing bids and sell offers of both you and the buyer will disappear.

Buy Column

There are two ways to buy assets: 1) You can make a bid that other traders can choose to accept, and 2) You can accept sell offers made by other traders. In the Buy Column, you will see the following three boxes:

Make bid ← This lets you make an offer to buy one share of the asset. Type in the price you are willing to pay for a share, and press "OK". This price must be a whole number (no decimals!), and must be positive. When you enter a bid, your bid price will appear in the "Bid Price" box in the Sell Column on the lower left hand side of the screen, where sellers can choose to accept it.

Top Highest Bid ← Pressing this button will enter a bid that is 1 yen higher than the highest existing bid. It is a way to make a bid quickly. If there are no existing bids, then this button won't do anything.

Offer Price box ← This box contains a list of all outstanding sell offers. You can buy a share by pressing the "**BUY NOW**" button at the bottom of the box. You will automatically buy one share for the lowest available sell offer price (as long as it's not your own sell offer!). Each time you buy a share, ALL of the existing bids and sell offers of both you and the seller will disappear.

Note: If you try to buy a share for more cash than is in your account, or if you try to sell a share when you have zero in your portfolio, you will get an error message.

Appendix B - Forecast Chart

(Note: This chart was given to all treatment groups except for the control group.)

If the Dividend is 10 Yen

Forecast	% Chance
Extremely Bullish	35
Moderately Bullish	30
Neutral	20
Moderately Bearish	10
Extremely Bearish	5

If the Dividend is 0 Yen

Forecast	% Chance
Extremely Bullish	5
Moderately Bullish	10
Neutral	20
Moderately Bearish	30
Extremely Bearish	35

Table 1: Groups

Group	Forecasts	Preliminary Tasks	Average Miscalibration	Average BTA Effect
C	no	hard	-0.33	1.17
P	yes	hard	1.33	1.67
Q	yes	easy	3.33	1.17
R	yes	hard	4.17	0.83
S	yes	easy	1.00	1.50
T	yes	hard	0.83	2.00

Table 2: Confidence Test

Jar	Color	# of Balls
A	Red	15
	Orange	18
	Yellow	20
	Green	22
	Blue	25
B	Red	25
	Orange	22
	Yellow	20
	Green	18
	Blue	15

Table 3: Bubble Sizes

	Relative Deviation		
Group	Round 1	Round 2	Round 3
C	26.67	1.16	-1.24
P	3.2	-2.3	-2.48
Q	23.88	14.47	0.77
R	53.11	24.86	0.35
S	40.56	23.53	11.07
T	3.4	2.79	-0.12

Table 4: Forecasts and Trading Decisions

	All Rounds		Round 3	
Regressor	Coefficient	p-value	Coefficient	p-value
MISCi	0.02	0.37	0.05	0.11
Fit	0.35***	0.00	0.28*	0.06
MISCi x Fit	-0.02	0.21	0.00	0.90
BTAi	0.00	0.93	0.05	0.61
Fit	0.34***	0.00	0.31*	0.06
BTAi x Fit	-0.02	0.51	-0.01	0.90
LHS: NETBUYit	#obs: 1080		#obs: 360	

Table 5: Forecasts and Prices

	All Rounds		Round 3	
Regressor	Coefficient	p-value	Coefficient	p-value
BMISCi	-0.57	0.32	0.32	0.66
Fit	0.32	0.52	2.58***	0.01
BMISCi x Fit	0.35*	0.07	-0.56*	0.10
BBTAi	-1.38*	0.07	-1.05	0.17
Fit	0.94**	0.02	1.78**	0.03
BBTAi x Fit	-0.05	0.82	-0.19	0.56
LHS: DMPRICEit	#obs: 1262		#obs: 396	

Table 6: Individual Outcomes

		All Rounds		Round 3	
LHS	Regressor	Coefficient	p-value	Coefficient	p-value
NUMTRADESi	MISCi	-0.08	0.50	0.01	0.87
	BTAi	-0.05	0.77	-0.05	0.41
Πi	MISCi	40.67	0.53	-9.69	0.61
	BTAi	93.29	0.34	55.93**	0.06
	D_CONFi	608.37**	0.07	-13.34	0.90
		#obs: 36			

Table 7: Group Outcomes

LHS	Regressor	Coefficient	p-value
VOLg	MISCg	-0.13	0.31
	BTAg	8.04**	0.06
Bg	MISCg	5.22	0.53
	BTAg	-32.17	0.32
SDPg	MISCg	-0.09	0.98
	BTAg	-14.07	0.15
	#obs: 6		

Figure 1a: Average Market Prices in Round 1

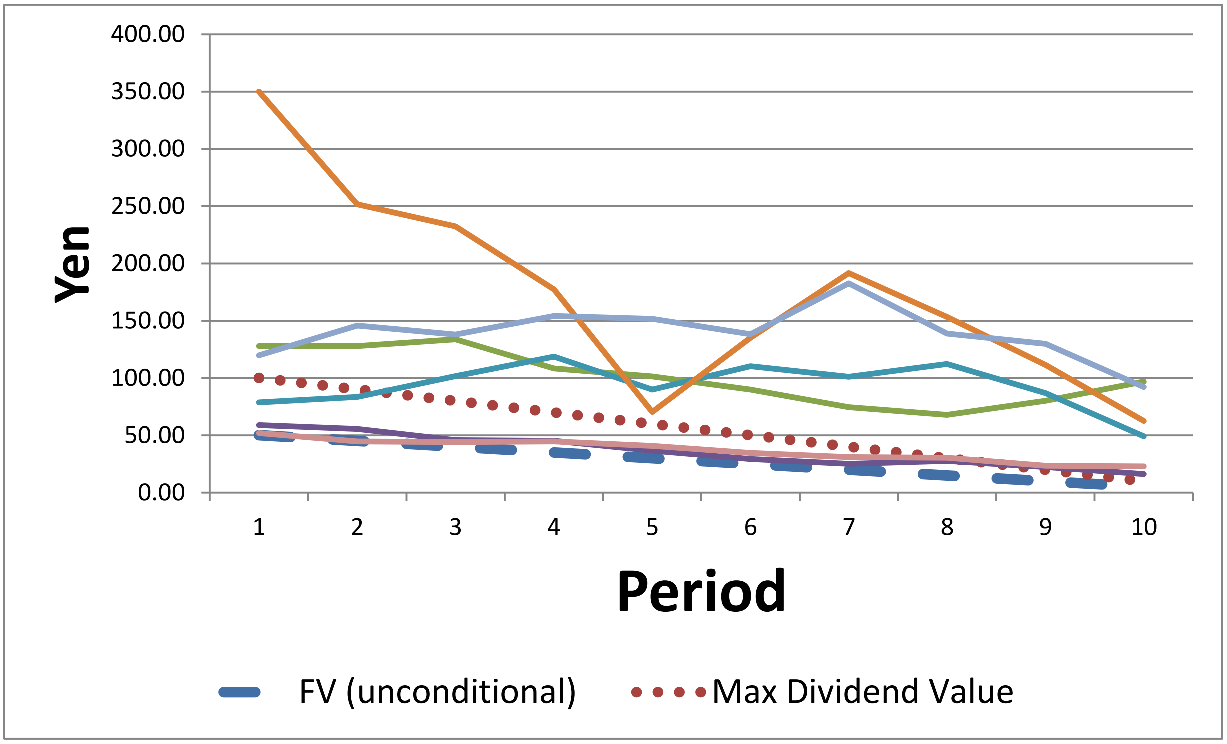

Figure 1b: Average Market Prices in Round 2

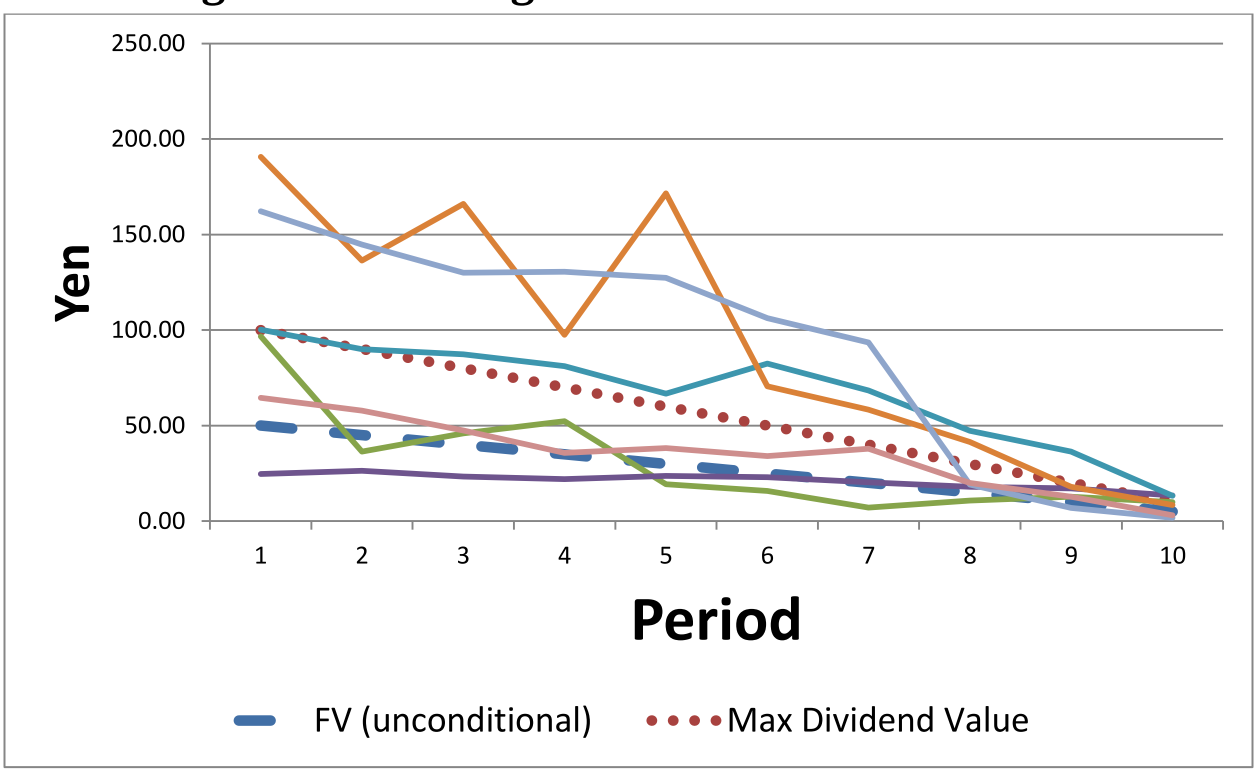

Figure 1c: Average Market Prices in Round 3

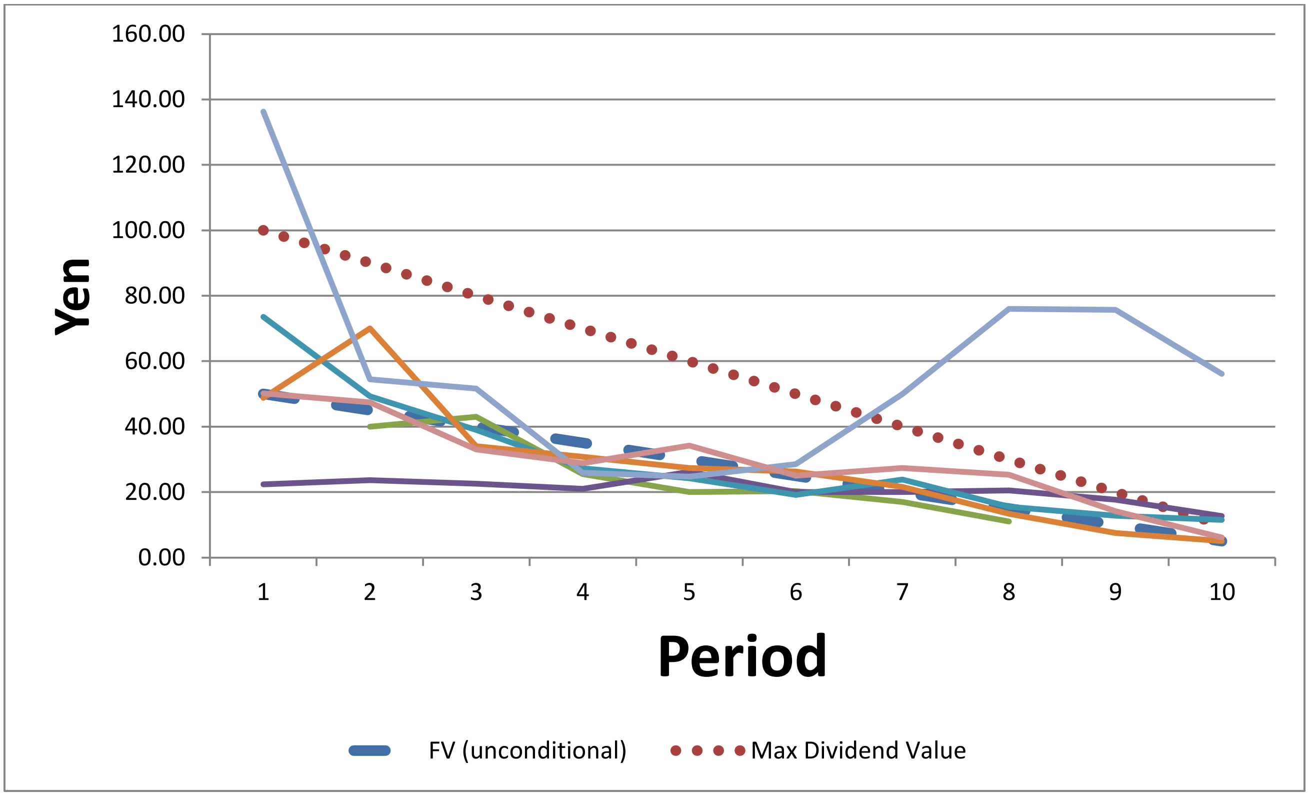

Figure 2a: Volume vs. Better Than Average Effect

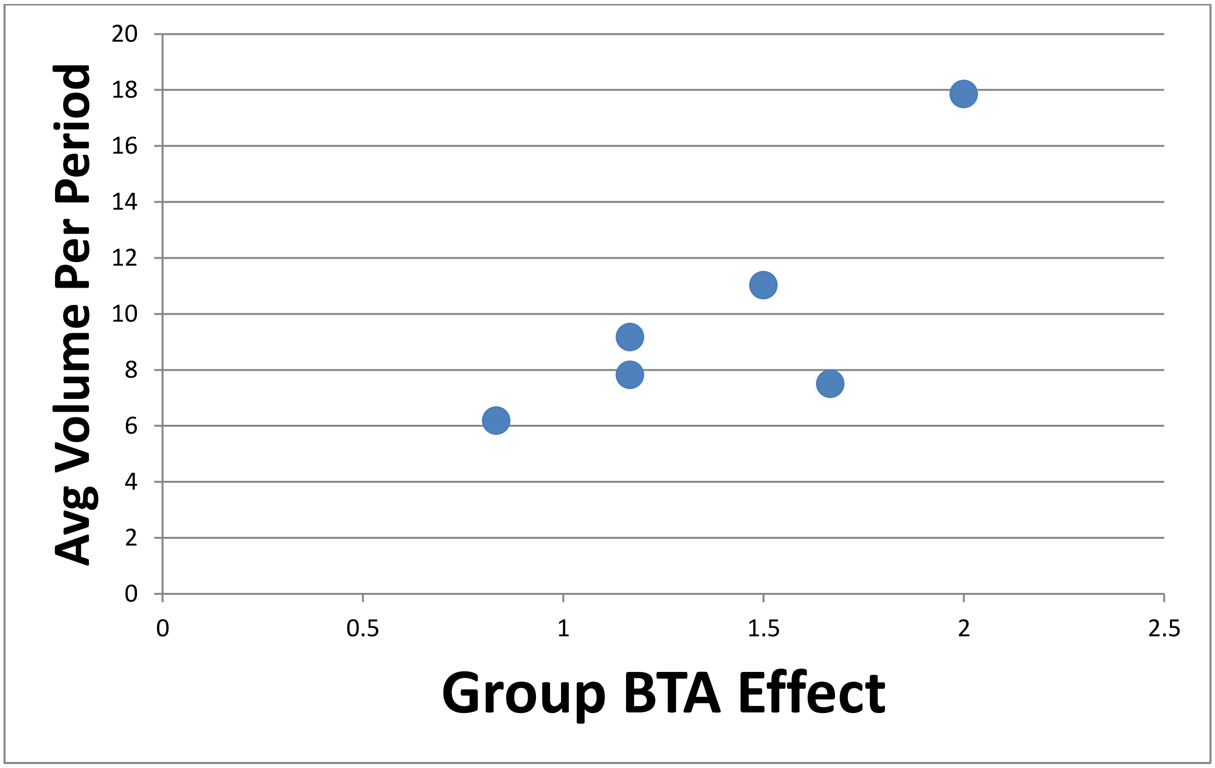

Figure 2a: Volume vs. Miscalibration