Distinguishing the Effect of Overconfidence from Rational Best-Response on Information Aggregation

Shimon Kogan
Tepper School of Business, Carnegie Mellon University

This article studies the causal effect of individuals' overconfidence and bounded rationality on information aggregation by using a new multiperiod game in which agents are rewarded for submitting accurate estimates of an unknown asset's value based on (i) their private information and (ii) others' past estimates. By carrying out laboratory sessions of this game in which subjects' overconfidence is a treatment variable, I find that overconfidence affects the information aggregation process by increasing the dispersion of estimates and decreasing the rate of estimates' convergence. However, even when subjects are not overconfident, qualitatively similar deviations from the fully rational model predictions are observed. I show that this can be explained by subjects' strategic response to errors. (*JEL* G12, C92, D83)

Overconfidence—individuals' tendency to overestimate their abilities and information—has been documented by a large body of psychology literature (see, for example, Fischhoff, Slovic, and Lichtenstein 1977; and Kahneman, Slovic, and Tversky 1982). In finance, it has been incorporated into models of asymmetric information to rationalize a set of long-standing asset pricing "anomalies." For example, Odean (1998) shows that in a variety of models overconfidence leads to excess trading volume, and in some settings can also lead to higher price volatility and lower price quality. Daniel, Hirshleifer, and Subrahmanyam (1998) find that overconfidence can cause positive returns autocorrelation in the short run (e.g., momentum) and negative returns autocorrelation in the long run. In these settings, overconfident agents exhibit inefficient information processing as they overweight a particular information set compared with what a perfectly Bayesian agent would do. For example, an overconfident

This is an updated version of a paper titled "Distinguishing Bounded Rationality from Overconfidence in Financial Markets—Theory and Experimental Results." I am indebted to my thesis committee Co-Chairs, Jonathan Berk, and Jacob Sagi, and my thesis committee members, Shachar Kariv, John Morgan, Terry Odean, and Nancy Wallace, for their insightful suggestions and continuous support, and I wish to thank the editor, Tobias Moskowitz, and an anonymous referee for their guidance. I would also like to acknowledge helpful comments by George Akerlof, Bob Anderson, Stefano DellaVigna, Vincent Glode, Rick Green, Teck Ho, Botond Koszegi, Lars Lochstoer, Rich Lyons, Ulrike Malmendier, Barbara Mellers, seminar participants at Carnegie Mellon University, Federal Reserve Bank of Boston, Harvard Business School, London Business School, MIT Sloan School of Management, Tel-Aviv University, University of British Columbia, University of California at Berkeley, University of Texas at Austin, University of Utah, Western Finance Association 2005 Meetings, and Yale School of Management for their valuable comments. The financial research support of The Center for Responsible Business, The Dean Witter Foundation, and IBER is gratefully acknowledged. Send correspondence to Shimon Kogan, Tepper School of Business, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213; telephone: (412) 268-8501; fax: (412) 268-7357. E-mail: kogan@andrew.cmu.edu.

doi:10.1093/rfs/hhn075 Advance Access publication August 27, 2008

agent may overweight her private signal relative to the information revealed by aggregate quantities (e.g., prices).

In this article, I suggest an alternative reason that may cause agents to *rationally* "overweight" their information: they respond to the errors of others. If some people make mistakes by failing to follow a fully rational behavior, then aggregate quantities, which facilitate the process of information transmission, may contain errors (see Akerlof and Yellen 1985). In such cases, a rational agent would strategically respond to these errors (see Camerer 2003) by discounting the informativeness of fellow traders' actions and overweighting their own private signal.

To examine these issues, I devise a simple, nonstrategic information aggregation game focused on how agents process private information and learn from each other. In the game, there are two players. Each receives (i) a private signal and (ii) a private-*signal precision*. The players' task is to estimate an unknown fundamental value, around which their signals are drawn. The game consists of multiple decision turns in which players first observe each other's *previously* submitted estimates, and then simultaneously submit new ones. Players' payoffs do not come from trading but rather depend on the accuracy of their *individual* estimates. Therefore, players' estimates are a solution to a Bayesian updating problem in which private information is weighted relative to the information learned from the other players' estimates. The more confidence a player has in her private information, the less she adjusts her estimates across turns. Over time (under full rationality), players are predicted to perfectly aggregate their private information, converging to the rational expectations level.

I extend the fully rational baseline model in two ways. First, I derive best responses when players are overconfident. In this case, players' confidence in their private signal exceeds their objective signal precision. Second, I derive players' best responses when interacting with others prone to have mean-zero errors. I show that in this case, players discount observed estimates of others. In both cases, I find that the level of "disagreement" among players, measured by their estimates' dispersion, exceeds the fully rational benchmark.

I conduct experimental sessions in which subjects participate in the information aggregation game described above and are rewarded in cash based on their decisions. Private signal precision (high or low) is determined by the subjects' rank on a task that takes place at the beginning of sessions. In some sessions, denoted as *baseline treatment* (BLT), participants roll a die and *privately* observe the outcome. In other sessions, denoted as *overconfidence treatment* (OCT), participants answer a short Scholastic Aptitude Test (SAT) quiz, similar to a treatment used by Camerer and Lovallo (1999). In both treatments, subjects are not told what their rank is, but are made aware of the way it is determined. While the die roll is a neutral treatment, the SAT is not; many previous studies document the tendency of individuals to perceive themselves as "better than average" (e.g., Svenson 1981). Therefore, subjects who mistakenly believe that they are better than their peers on the SAT quiz will also mistakenly believe that their signal precision is better than it really is. They will be overconfident about their private signal—not by conjecture, but rather by the experimental design.

I use the observed adjustment rates to back out subjects' implied confidence in their signal and compare those to the objective precisions. I find that, consistent with the treatment design, subjects are (on average) overconfident in the OCT but not in the BLT. In particular, overconfidence seems to be concentrated among subjects that have the lowest signal precision. Furthermore, observed overconfidence seems to have an adverse effect on information aggregation. Comparing the results across treatments, I find that in the OCT, estimates' dispersions are substantially higher and the rate of estimates' convergence is lower. However, comparing the results in the BLT with the predictions of the fully rational model, I find deviations that are qualitatively similar to the differences across treatments. To explain this, I estimate an extension of the fully rational model, which accounts for best responses to errors and finds support for it.

The information aggregation game used in this article abstracts from many aspects typically embedded in trading markets, such as risk-sharing, portfolio choice, strategic behavior, and optimal order submission. Central to the game's tractability is the fact that participants' payoffs are not a fixed sum (as is the case when trading is involved). As such, the game does not generate prices, allocations, and trading volume, and therefore cannot be used to provide direct observation on these quantities. I am able to derive a tractable model that captures how players integrate information obtained privately with information indirectly revealed to them from aggregate quantities. I focus on this aspect since it captures the channel through which overconfidence is hypothesized to affect agents' decisions in markets.

The results obtained in this study can be relevant to the examination of overconfidence in financial markets in a number of ways. First, it shows that the two sources for inefficient learning—overconfidence and errors—can have economically similar effects on information aggregation. Models of overconfidence could benefit from including the two sources of inefficiencies and providing predictions that separately identify them. Second, the mechanism studied in this article shares some important features with markets populated by asymmetrically informed agents. In both cases, the object of interest is the way private information is mapped into agents' valuation of the underlying asset. While the mechanism connecting the inputs and outputs differs, in both cases, it relies on Bayesian updating to determine how agents weight their private information with that implied from others' behavior.[1] As such, the findings shed light on the ways deviations from rationality affect this updating.

[1] Indeed, the idea that prices reflect a weighted average of traders' individual estimates about the fundamental value is very common in information-based asset-pricing models (Diamond and Verrecchia 1981), and the notion that speculative volume can be driven by traders' differing valuations is discussed by Varian (1986), Harris and Raviv (1993), and Kandel and Pearson (1995), among others.

The remainder of the article is organized as follows. Section 1 summarizes the related literature. Section 2 sets up the theoretical model and derives the unique subgame-perfect Nash equilibrium. Section 3 describes the experimental design mirroring this model. Section 4 discusses and analyzes the results. I summarize in Section 5.

1. Literature Review

In this section, I focus on the most closely related experimental and psychology work on overconfidence, at the individual and market level, and relate it to this study. In the interest of brevity, I omit a number of strands of literature including theoretical asset pricing models studying the effect of traders' overconfidence and work on social learning (see Bikhchandani, Hirshleifer, and Welch 1998 for a survey).

Kirchler and Maciejovsky (2002) measure individuals' miscalibration and how it is affected by trading. They allow subjects to trade securities that pay stochastic dividends over a number of periods in double auction markets. At the beginning of each period, subjects are asked to provide an assessment of the distribution of trading prices to be observed during the period. The authors use these predictions to construct two different measures of overconfidence. They find substantial heterogeneity in subjects' levels of over-/underconfidence but find no aggregate overconfidence.

Biais et al. (2002) correlate individual measures of overconfidence and self-monitoring, collected through surveys at the beginning of sessions, to their earnings from trading in an asset market similar to the one described in Plott and Sunder (1988). Absent private values, this setting is subject to winner's curse since participants hold imperfect private information. Thus, no-trade results apply. The authors find that subjects prone to overconfidence generate relatively low earnings and those that exhibit high self-monitoring abilities generate relatively high earnings. At the same time, they do not find that overconfidence leads to more intense trading.

Deaves, Lüders, and Luo (2005) study the link between miscalibration, gender, and trading intensity. The authors design an experiment in which subjects' signal quality depends on the accuracy of their responses on a survey. They find that overconfidence leads to increased trading activity among subjects. In contrast with other studies, they find no difference in overconfidence and trading intensity between men and women.

Glaser and Weber (2007) conduct a survey among broker investors to assess their overconfidence as expressed in their miscalibration and better than the average effect (as well as illusion of control/unrealistic optimism). Results from 215 individuals are matched with their own trading volume. The authors find that the two measures yield distinctively different results: above average effect, but not miscalibration, is related to trading volume.

There are a number of important differences between these papers and this study. First, I measure participants' over-/underconfidence that is *implicit* in their decisions by estimating a structural model of behavior. That is, instead of using direct elicitation, I estimate revealed miscalibration as it applies in the information aggregation game studied here. As a result, the approach in this study is more robust to survey methodology; it is quite possible that while individuals are not able to *communicate* probabilistic assessments well, they are able to *incorporate* them into their decisions.

Second, the game played in this study differs substantially from the canonical double-auction markets, used by virtually all experimental asset markets. The tractability of the mechanism used in this study allows one to examine a rich set of aggregate measures and to generate clear *level* predictions. In contrast, most previous work almost exclusively dealt with *comparative static* tests. For example, no study that I know of looks at the effect of overconfidence on the precision of agents' valuation of the underlying asset, while it is clearly of central interest to economists.

Third, most previous work in this area follows the approach of correlating individuals' psychological attributes and their behavior in markets. As such, it was centered on *individual level* results. The focus in this article is on making some initial steps toward understanding the aggregation process of individuals' biases.

The experimental design in this study makes use of two distinct forms of overconfidence: miscalibration and better-than-average effect. Both forms have been studied extensively in cognitive psychology. Miscalibration refers to individuals' tendency to overestimate the accuracy of their knowledge (see Kahneman, Slovic, and Tversky 1982 for a review). Studies of miscalibration typically involve the elicitation of confidence intervals: subjects are presented with a set of open-ended questions (e.g., "How much does a fully loaded Boeing 747 weigh?") and asked to provide a range of values such that they are 90% confident that the unknown quantity would fall inside the stated range. Evidence suggests that subjects report too narrow confidence intervals. That is, the fraction of answers landing inside the confidence interval is substantially lower than 90%. An alternative design presents subjects with multiple-choice questions and asks subjects to estimate the probability that the answer they pick is correct. Once again, in aggregate, the fraction of correct answers is lower than the stated probabilities. Studies of the better-than-average effect suggest that people are overconfident about their ability relative to others. For example, Svenson (1981) asked groups of subjects to compare their driving ability (skill and safety) to their peers in a group. Around 70–80% rated themselves as above median. This phenomenon has been documented in various other areas: health (Larwood 1978), managerial skills (Larwood and Whittaker 1997), and business success (Cooper, Woo, and Dunkelberg 1988). While I do not suggest that, in general, these two forms of overconfidence are related, I make use of the better-than-average effect to generate miscalibration:

since subjects overestimate their relative ability to answer correctly the SAT quiz questions, they should also overestimate the precision of their private signal.

2. Theory

2.1 General

In this game, there are two players, both trying to estimate the realization of a random variable v, referred to as "fundamental value," where v is distributed uniformly between L and H. Each player is assigned a type: $t_i \in \{h, l\}$ such that one player is of type h and the other is of the complementary type l. A player of type h receives a perfect signal while a player of type l receives an imperfect signal:[2]

- perfect signal: $s_i^h = v$;
- imperfect signal: $s_i^l = v + e_i$, where $e_i \overset{iid}{\sim} U[-Y, Y]$.

Players' type is determined as follows. Both players (i and j) independently draw quantities q_i and q_j from a uniform distribution with a support [0, 1]. The player with the highest q is assigned to be type h while the other player is assigned to be type l. For player i, the probability that she is of type h is q_i, since $\Pr(q_j < q_i | q_i) = q_i$.

For now, assume that subject's belief about the probability that she is of type h, denoted by $\widetilde{q}_i$, is correct (i.e., $\widetilde{q}_i = q_i$). Thus, the realization $\{s_i, q_i\}$ makes up subject i's private information set.

The game consists of three stages: at the beginning of each turn, t, both players simultaneously submit an action, $a_{i,t}$, which comes in the form of a numerical estimate of the realized fundamental value. At the end of each turn, both players' estimates are announced. As I show later, three turns are needed for players to arrive at the fully revealing equilibrium. The intuition is straightforward. There are two dimensions of uncertainty for each player—the other player's signal and her signal precision. Since each turn can allow for at most one new dimension to be observed, players need to observe each other's estimates for two turns, arriving at full revelation in turn three.

At the end of the game, one turn is randomly chosen (with equal probability) and players receive a payoff $\pi_i(a_{i,t}, v)$, ensuring that expected utility is maximized at the expected value of v: $E(v|I_{i,t}) \in \arg\max E[u_i(\pi_i(a_{i,t}, v))|I_{i,t}]$, where $I_{i,t}$ represents player i's information set (both private and public) in turn t. Put differently, the payoff scheme ensures that if players act myopically, they minimize the forecasting error at each turn of the game. Note that each player is paid according to the accuracy of their actions, irrespective of the actions of the other player. Hence, this is not a fixed-sum game (unlike most trading

[2] Subscripts i, j index the players.

institutions). This feature is important in neutralizing the payoff externalities that typically arise in market settings, thereby removing strategic incentives.

2.2 Optimal actions: full rationality

I start by characterizing the fully rational solution of this game for players i, j, denoted by a_i^*, a_j^*. Recall that the game starts with the players receiving their private information, $\{s_i, q_i\}$, followed by three decision turns. Since the exogenous information is fixed across the turns, subjects revise their estimates due to *endogenous* information only, which is obtained by observing others' actions. Also, since exactly one player is perfectly informed but the identity of that player is uncertain, optimal actions are a convex combination of players' signals. How far one's estimate is from her signal would depend on the confidence she has in her signal.

While I characterize the solution for the full range of signals in the Appendix, in this section I focus on the "interior case" where $s_i, s_j \in [L+Y, H-Y]$. The interior case gives rise to tractable predictions, allows one to obtain clear intuition about the dynamics of this game, and can be made arbitrarily large relative to the "exterior case."[3]

Proposition 1. *There exists a Perfect Bayesian Equilibrium (PBE) in pure strategies, where:*

- *in turn* 1 $\{a_{i,1}^* = s_i, a_{j,1}^* = s_j\}$
- *in turn* 2 $\{a_{i,2}^* = q_i s_i + (1-q_i)a_{j,1}, a_{j,2}^* = q_j s_j + (1-q_j)a_{i,1}\}$
- *in turn* 3 $\{a_{i,3}^* = \text{Ind}_{(q_i > \frac{a_{j,2}-a_{i,1}}{a_{j,1}-a_{i,1}})} s_i + \text{Ind}_{(q_i < \frac{a_{j,2}-a_{i,1}}{a_{j,1}-a_{i,1}})} a_{j,1} + \text{Ind}_{(q_i = \frac{a_{j,2}-a_{i,1}}{a_{j,1}-a_{i,1}})} \times (\frac{s_i + a_{j,1}}{2}),\quad a_{j,3}^* = \text{Ind}_{(q_j > \frac{a_{i,2}-a_{j,1}}{a_{i,1}-a_{j,1}})} s_j + \text{Ind}_{(q_j < \frac{a_{i,2}-a_{j,1}}{a_{i,1}-a_{j,1}})} a_{i,1} + \text{Ind}_{(q_j = \frac{a_{i,2}-a_{j,1}}{a_{i,1}-a_{j,1}})} \times (\frac{s_j + a_{i,1}}{2})\}$.

See the Appendix for proofs of the propositions. At the first turn, players should announce their own signals. At the second turn, players should announce a weighted average of their signal and the other player's estimate. Since players follow the equilibrium strategies, this amounts to player i weighting her signal with the *signal* of player j, revealed in player j's first turn estimate. The weight depends on their confidence in their private information. Therefore, players that adjust their estimates very little between turns one and two have high confidence in their own private signals. Likewise, players that adjust their estimates very much have low confidence in their own private signals.

Proposition 2. *The pure strategy equilibrium in Proposition 1 is unique.*

[3] Note that the probability of an interior case goes to 1 as the relative magnitude of the signal imprecision (Y) to the range of v's ($H-L$) goes to 0.

The intuition for the uniqueness result is that players' payoffs are not a fixed sum. Therefore, they cannot do better than to provide their Bayesian estimates of v at each turn of the game. They would deviate from this strategy if they could learn faster the information the other player holds, but this is not feasible.

In the game, information is aggregated sequentially. In turn one, optimal action depends on one's own signal. In turn two, optimal action depends on one's own signal, as well as the other's observed $t = 1$ action and own subjective confidence. To measure the information aggregation rate, I focus on the level of "disagreement" between players, labeled "estimates' dispersion."[4]

Definition. Estimates' dispersion is defined as $\mathrm{ED}_t = |a_{i,t} - a_{j,t}|$.

Proposition 3. *Estimates' dispersion strictly decreases from turn one to three* $(a.s.)$.

This result is intuitive: if players aggregate information, then their information sets will get closer together over time. Since estimates are functions of these sets, estimates too will get closer together. At the first turn, since players' estimates are equal to their signals, the dispersion depends on the realization of the signals. At the second turn, dispersion decreases since players' estimates are a convex combination of their own signal and the initial estimate of the other player. At the third turn, both players submit the same estimate, bringing the dispersion to zero.

2.3 Optimal actions: miscalibration

Recall that I am interested in understanding the effects of errors in beliefs and errors in actions on players' best responses. This part addresses the former by allowing for the possibility that players hold erroneous beliefs about the probability of being perfectly informed. That is, I set the individual subjective probability to be equal to the objective probability plus miscalibration: $\widetilde{q}_i = q_i + \mathrm{MC}_i$, where MC_i denotes subject i's miscalibration. Positive miscalibration represents overconfidence while negative miscalibration represents underconfidence. I allow for arbitrary subjective beliefs as long as they are admissible, that is, $0 \leq \widetilde{q}_i \leq 1\ \forall i$.[5] To simplify matters, assume that subjects are naïve in the sense that they are not aware of the other player's potential miscalibration (this assumption will be maintained throughout this article).[6]

Proposition 4. *If players' miscalibration is equal to* MC_i *and* MC_j*, then estimates' dispersion is given by:*

[4] Results not reported here also analyzed estimates' imprecision, a measure of the distance between the average estimate and v. Most of the qualitative predictions carry over to this measure too.

[5] This definition corresponds to the way miscalibration is defined in the cognitive psychology literature (see Kahneman, Slovic, and Tversky 1982 for a review).

[6] I naturally assume that players are not aware of their own miscalibration (see, for example, Odean 1998).

- *Turn 1:* $\text{ED}_1^M = |s_i - s_j| = |e|$;
- *Turn 2:* $\text{ED}_2^M = |1-(q_i+q_j)-(\text{MC}_i+\text{MC}_j)||e|$;
- *Turn 3:* $\text{ED}_3^M = 0$.

In this setting, miscalibration should not affect estimates' dispersion in the first turn. This is because first turn estimates do not reflect players' subjective confidence in their own signals. For the same reason, miscalibration would impact estimates' dispersion in the second turn. Since both players regard their subjective beliefs to be equal to the objective precisions of their signals, they would converge on the signal held by the player with the higher subjective probability. Therefore, turn three estimates' dispersion would be zero. Next, I compare the level of estimates' dispersion in the presence and absence of miscalibration.

Proposition 5. *The expected level of estimates' dispersion is (weakly) greater in the presence of average overconfidence for all turns (a.s.).*

In the first and last run of the game, overconfidence should not have an effect on estimates' dispersion. However, if players are on average overconfident, i.e., $\text{MC}_i + \text{MC}_j > 0$, turn two estimates should be further apart. Player i would adjust her estimate too little toward the estimate of player j, compared with the fully rational case.

2.4 Optimal actions: rational response to errors

I now study the second form of deviation from full rationality. I ask: what should player i's best response be if she knows that player j's actions include some error? I address this by following the spirit of the probabilistic choice model of Goeree and Holt (1999) and Mckelvey and Palfrey (1995), in which players' actions are *distributed around* their best responses; best responses are formed while taking into account that others' actions include mean-zero errors.

Specifically, I let player j's observed action be composed of her fully rational optimal action, as derived in Section 2.2, plus a mean-zero error:

$$a_{j,1} = a^*_{j,1}(s_j) + e_{j,1} = s_j + e_{j,1}. \tag{1}$$

While player i cannot observe player j's error, she at least accounts for the fact that these errors have magnitude σ_1. Therefore, player i's best response in turn two is not only a function of her information set $I_{i,2}$, but also of the error magnitude σ_1:

$$a^*_{i,2} = a^*_{i,2}(I_{i,2}, \sigma_1). \tag{2}$$

Assume that players' observed actions are normally distributed around the optimal action in that turn. This assumption automatically satisfies the condition

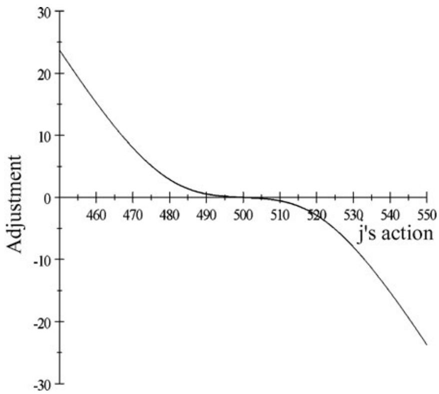

Figure 1
Turn two adjustment term effect (numerical example)
This figure depicts an example for the adjustment player i applies to player j's observed turn one estimate.

that the probability of observing deviations from best response is inversely related to their cost since the payoff function used in the experiment has the same functional form as the normal density (both are negative quadratic exponential). As such, this assumption is economically appealing.

After some calculations (see the Appendix for derivation), I obtain that:

$$a_{i,2}^{*} = q_i s_i + (1 - q_i) \times \left(a_{j,1} + \frac{2\sigma_1^2(\phi(a_{j,1}; s_i - Y, \sigma_1) - \phi(a_{j,1}; s_i + Y, \sigma_1))}{\operatorname{erf}\left(\frac{\sqrt{2}}{2\sigma_1}(Y - s_i + a_{j,1})\right) - \operatorname{erf}\left(\frac{\sqrt{2}}{2\sigma_1}(a_{j,1} - s_i - Y)\right)} \right). \quad (3)$$

To interpret this expression, I contrast it with the solution obtained in the fully rational model, $a_{i,2} = q_i s_i + (1 - q_i)a_{j,1}$. Notice that player i's response in turn two to the actions of player j in turn one is now adjusted relative to the fully rational case. This adjustment becomes larger the further player j's actions are from player i's signal. The reason is that "extreme" estimates are more likely to include a large error term relative to "moderate" estimates. Therefore, the informativeness of extreme estimates should be discounted. Figure 1 plots this adjustment term for the following parameter values: $s_i = 500, \sigma_1 = 10,$ and $Y = 30$.

One can see that for "moderate" values (480–520), the adjustment is small. Going outside that range results in a steep adjustment in the direction of the signal possessed by player i. As a result, player i's turn two estimate is closer to her signal than would be the case if players were fully rational. Note that if I did not account for the possibility of errors, player i's estimate would appear to be biased in the same direction as it would if she were overconfident.

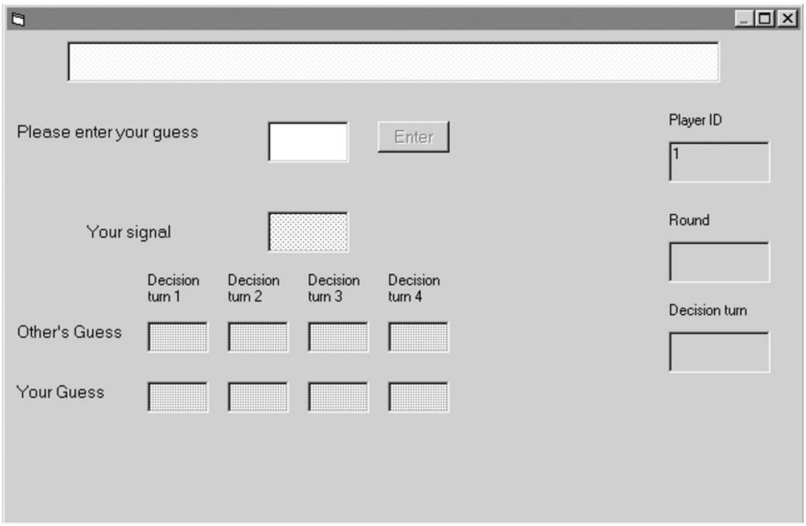

Figure 2
Interface screenshot
This is a screenshot of the interface used in this study.

3. Experimental Design

3.1 General

The experiment was run at the Haas School of Business, University of California at Berkeley: a total of 12 sessions were conducted in which 72 subjects participated; five sessions were BLT and seven sessions were OCT.[7] Subjects were recruited from undergraduate classes and had no previous experience with similar experiments. They received a show-up payment of $5 and an additional performance-based pay of $0–$10, which was paid in private and in cash at the end of the session. Sessions were about 60 minutes long and included six participants each.

At the beginning of each session, an administrator read the instructions aloud and answered questions in private.[8] Each subject entered their decision using a computerized interface, which was built for the purpose of this experiment (Figure 2), thus maintaining both isolation and anonymity.[9]

3.2 Structure

Each session started with an initial phase, followed by 10 independent and identical rounds.[10] At the beginning of each round, subjects were randomly

[7] The order of treatments was determined randomly.

[8] Instructions are available upon request.

[9] The application developed by the author for this experiment is available upon request.

[10] I report here the results from the first 10 rounds, while a few sessions were conducted with more rounds.

assigned into "markets" consisting of two players each and were presented with their private signal. Each round was composed of four decision turns, and for each turn, each subject was asked to enter a decision.[11] Throughout the turns, subjects' pairing and private information remained the same. Transition from one turn to the next occurred only after all subjects submitted their decisions. No time restriction was imposed.

The experiment was carried out along a single treatment: BLT or OCT, which differed *only* in their initial phase. In the BLT, the initial phase consisted of subjects privately throwing a die and observing its outcome.[12] Draws were recorded by the administrator and fed into the computer, which then determined the rank of the draws; three of the subjects, with the highest draws, were classified as "perfectly informed" while the other three, with the lowest draws, were classified as "imperfectly informed" (ties were resolved randomly). Subjects observed their own draw but not the draws of other participants.

In the OCT, subjects were asked to answer 20 multiple-choice SAT questions (taken from sample tests that were posted on the College Board website). Scores were recorded by the computer, which then ranked subjects according to the number of correct answers, as a primary key, and by the length of time required to complete the quiz, as a secondary key.

In both treatments, subjects were not told what their rank was but were made aware of the way it was determined. Altering the procedure by which precision of signals was determined was designed to make subjects overconfident (on average) about their private signal in the OCT. While the die throw is a neutral treatment, the SAT is not; many previous studies document the tendency of individuals to perceive themselves as "better than average" (e.g., Svenson 1981). Therefore, subjects who mistakenly believe that they are better than their peers on the SAT quiz will also mistakenly believe that their signal precision is better than it really is. They will be overconfident about their private signal—not by conjecture, but rather by the experimental design.

The choice of using SAT questions was deliberate and intended to bias the results in favor of the null of no treatment effect by facing subjects with a task with which they are familiar. All subjects had taken the SAT test before and were aware of their performance relative to their peer group.

3.3 Information

At the beginning of each round, a quantity v was drawn by the computer, where $v \sim U[50, 950]$. Then, subjects received independent signals. Subjects classified as "perfectly informed" received a signal equal to the realization of

[11] While the fourth turn is redundant under the fully rational model, it need not be redundant in practice. I also run two sessions (not reported here) with six decision turns, but behavior during the last two turns seemed very close to the one exhibited in turn four.

[12] At the beginning of each experiment, one subject was publicly asked to examine the die and confirm that it appeared normal.

v (i.e., $e_i = 0$). Subjects classified as "imperfectly informed" received a signal equals to $v + e_i$, where $e_i \sim U[-30, 30]$.

All information was continuously displayed on subjects' interfaces for them to observe. Note that aside from the information specified above, no additional feedback was given. In particular, the realization of the unknown quantity, v, was not revealed at any stage of the experiment (not even at the end of the round) and subjects did not see their earnings until the end of the session. This may be likened to an environment where traders never get to observe the liquidating value; subjects can only learn from their interaction with other players, not from exogenous cues.[13]

3.4 Assignment

Pairing into markets was randomly determined while ensuring that exactly one subject was perfectly informed and the other one was imperfectly informed. This is important for three reasons: (i) it makes ex ante distribution of information equal across all market instances, (ii) it disables subjects from easily unveiling their type, and (iii) it allows posterior probability updating to take on a particularly simple and intuitive form.[14,15]

3.5 Decisions and payoffs

At the beginning of each turn, t, subjects simultaneously submit their estimates $a_{i,t}$ by entering a number on their screen. No restrictions are imposed on the value the estimate could take. Upon receiving submissions from both subjects, the turn comes to an end and no changes are accepted. At that point, subjects are informed of each other's estimate and are allowed a short transition time into the next turn.

At the end of the session, one turn from each round was randomly drawn and earnings (for subject i in round r) were calculated as follows:

$$\pi_i = \sum_r c_1 * \exp\left(-\frac{(v_r - a_{i,r})^2}{c_2}\right), \tag{4}$$

where $c_1 = 100$ and $c_2 = 50$. At the end of the experiment, the number of points earned was converted into dollars using an exchange rate of 100–1 and subjects were paid in private and in cash. Average earnings (including the show-up fee) are \$12 with a standard deviation of \$3.50.

[13] In a few sessions, I have extended the number of rounds to include a full feedback round: subjects' payoff and the realization of v was revealed at the end of the round. Subjects discovered whether they were the perfectly or imperfectly informed type almost immediately.

[14] If two subjects submit the same estimate in turn one, most likely they are both perfectly informed and thus from the next round on both players know their type with certainty.

[15] I have conducted a few sessions (not reported in this article) with different rules of market assignment. The problem discussed here does appear: when two perfectly informed subjects are paired together, they tend to find out their type. Nonetheless, the qualitative features of the experiment and the results are similar.

The payoff function in Equation (4) is chosen for a number of reasons. First, its convexity ensures that payoffs were nonnegative everywhere. This is desirable because of the bankruptcy possibility arising from subjects submitting estimates that are distant from the fundamental value (due to errors). Generally, bankruptcy is nonenforceable in the laboratory and once encountered may influence subjects' decisions in a substantial manner and may result in loss of experimental control (see Friedman and Sunder 1994). Second, the symmetry of payoffs around the fundamental value suggests to subjects that they should submit estimates that minimize estimation error. Indeed, the instructions reinforce this idea by stating that "the more precise your guesses are, the more money you will earn at the end of the experiment."

4. Results

4.1 Overconfidence

The game enables the study of how players incorporate their signal prevision beliefs into their estimates. While I cannot directly observe subjects' beliefs, I can estimate them from the sequence of their observed decisions. Specifically, since subject i's turn two optimal action, under miscalibration, is given by $\widetilde{q}_i s_i + (1 - \widetilde{q}_i)a_{j,1}$, I estimate the following regression model:

$$a_{i,n,2} = \alpha + \beta_1 s_{i,n} + \beta_2 a_{j,n,1} + \epsilon_{i,n}. \tag{5}$$

Treating the data as a panel and clustering errors by subjects, I estimate Equation (5) for each treatment. Columns 1 and 2 of Table 1 report the results. Under the null, subjects should put equal weight on their signal and on the other subjects' first turn estimates, i.e., $\beta_1 = \beta_2 = 0.5$. The results from the BLT are in line with this prediction; the average weights assigned to own signals and to the others' estimates are 0.55 and 0.45, respectively. Thus, I cannot reject the null that in the BLT, subjects are well calibrated. In contrast, the results from the OCT suggest that subjects significantly overweight their own signal: $\beta_1 = 0.72$ and $\beta_2 = 0.28$. I can reject the null that subjects are well calibrated in the OCT in favor of the alternative that they are overconfident.

To further explore the attributes of overconfidence, I back out subjects' signal precision beliefs from the rate at which estimates are adjusted across turns. Rearranging the expression for player i's best response in turn two, $a_{i,2} = \widetilde{q}_i s_i + (1 + \widetilde{q}_i)a_{j,1}$, I get that $\widetilde{q}_i = a_{i,2} - a_{j,1}/s_{i-a_{j,1}}$. To see how I interpret this quantity, consider two extreme cases. In the first case, subject i does not adjust her turn two estimate (i.e., $a_{i,2} = s_1$). This behavior implies that i is sure to hold the perfect signal; indeed, $\widetilde{q}_i$ would equal 1 in this case. In the second case, subject i adjusts her turn two estimate all the way toward subject j's estimate (i.e., $a_{i,2} = a_{j,2}$). This reflects that subject i believes that subject j holds the perfect signal; $\widetilde{q}_i$ would equal 0 in this case.

Table 1
Overconfidence and rational best response

	BLT (1)	OCT (2)	BLT (3)	OCT (4)
Own signal	0.545	0.7167	0.392	0.6361
	[0.0717]	[0.0597]	[0.0870]	[0.0754]
Other's estimate	0.4519	0.2839	0.6049	0.3644
	[0.0722]	[0.0592]	[0.0872]	[0.0749]
(Own signal −			0.0002	0.0001
Other's estimate)^3			[0.0001]	[0.0000]
Constant	1.8514	0.1268	1.6863	0.0514
	[1.2481]	[1.0618]	[1.1431]	[1.0270]
Observations	300	408	300	408

This table reports panel-regression results of subjects' turn two estimates. In Columns 1 and 2, I estimate the model $a_{i,n,2} = \alpha + \beta_1 s_{i,n} + \beta_2 a_{j,n,1} + \epsilon_{i,n}$, for each treatment separately. For Rows 3 and 4, I estimate the model $a_{i,n,2} = \alpha + \beta_1 s_{i,n} + \beta_2 a_{j,n,1} + \beta_3 (s_i - a_{j,1})^3 + \epsilon_{i,n}$, for each treatment separately. $a_{i,n,2}$ corresponds to subject i's turn two estimate in round n. Likewise, $a_{j,n,1}$ corresponds to subject j's turn one estimate in round n (subjects i and j are matched together in round n). $s_{i,n}$ corresponds to player i's signal in round n. "Own signal," "Other's estimate," and "(Own signal − Other's estimate)^3" correspond to β_1, β_2, and β_3, respectively. Standard errors, reported in square brackets, are pulled by subjects.

Figure 3 depicts the average implied confidence across treatments (BLT and OCT) while grouping subjects by their objective precision.[16] Objective precision is computed by bootstrapping the performance across subjects from the initial task (i.e., SAT). A subject who obtained X correct answers in the quiz was assigned an objective precision that corresponds to the probability that another subject, randomly drawn, would have obtained less than X correct answers on the quiz.

First, subjects' objective precisions are related to their implied confidence levels. This is true for both treatments. For example, in the BLT, subjects with low objective precision (0 to 1/3) have an implied confidence of 13% compared with 51% for subjects with average objective precision (1/3 to 2/3) and 89% for subjects with high objective precision (2/3 to 1). The same qualitative relation between objective precision and implied confidence is observed in the OCT.

Second, overconfidence is concentrated among subjects that are the least informed in the OCT. For each objective precision group, I test the null that the implied confidence is equal to the objective precision. For each subject, I average the implied confidence across rounds and compare it with her objective precision. Therefore, each subject is treated as a single observation. With this conservative use of data, I can reject the null that the poorly informed subjects in the OCT are well calibrated at the 1% significance level. To get a sense for the magnitude of these subjects' overconfidence, I compare them to their peers

[16] I excluded observations for which the implied confidence was outside the range 0–1; by doing so, I have taken out about 15% of the observations, across both treatments.

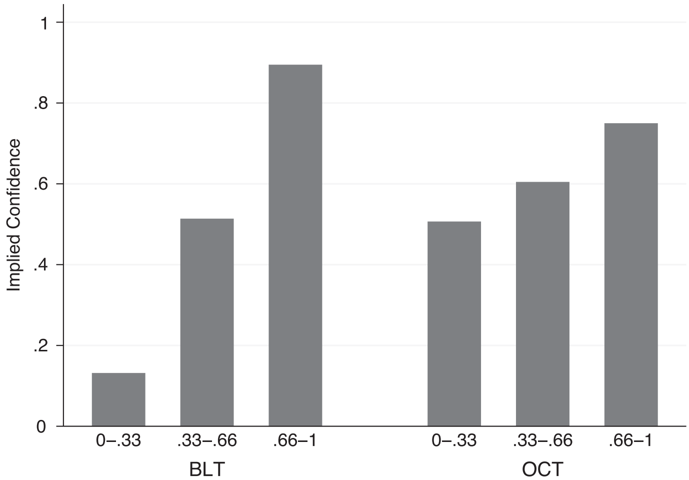

	Objective Precision Group		
	0 − .33	.33 − .66	.66 − 1
(p-values)			
BLT	0.1614	0.3882	0.7982
OCT	0.0010	0.5754	0.2273

Figure 3
Implied confidence across treatments
This figure reports average implied confidence across subjects. Implied confidence is computed as the ratio of own estimate change from turn one to turn two over the absolute difference between subjects' turn one estimates. The left panel depicts the results for the BLT and the right panel depicts the results for the OCT. Within each treatment, subjects are assigned into one of three objective precisions groups: $[0-1/3)$, $[1/3-2/3]$, and $(2/3-1]$. The table reports p-values associated with testing the null that the implied confidence is equal to the objective precision (using a nonparametric test). I exclude observations in which the implied confidence is not an admissible probability (outside the range 0–1).

in the BLT. I find that the former have almost four times higher confidence in their private information than the latter.

4.2 Estimates' dispersion

4.2.1 Aggregation of information. Theory predicts that players aggregate and disseminate private information across turns by observing each other's estimates. Therefore, estimates' dispersion, $\mathrm{ED}_{r,t}$, which is the absolute value of the difference between players' estimates in a given turn, $|a_{i,t} - a_{j,t}|$, decreases across turns. That is, players' estimates get closer together over turns.

To test this prediction, I calculate the average level of estimates' dispersion across turns and treatments (reported in Table 2). The results suggest that subjects aggregate information in both treatments. I see that the estimates' dispersion level decreases across turns in both treatments. In the BLT, the estimates' dispersion level start at around 17.6, decreasing to 5.0 by the fourth

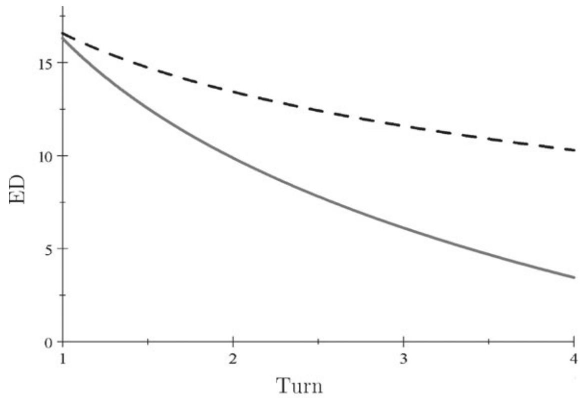

$ED_{i,t} = \alpha_1 + \beta_1 Dummy_i + \beta_2 \ln t + \beta_3 Dummy_i \ln t + \epsilon_{i,t}$

	$\widehat{\alpha}_1$	$\widehat{\beta}_1$	$\widehat{\beta}_2$	$\widehat{\beta}_3$
	16.295	0.269	−9.274	4.737
	[0.8623]	[1.1825]	[0.7215]	[1.0291]
Observations	1,416			
Groups	354			

Figure 4
Convergence of estimates' dispersion across turns
The figure depicts the convergence rate of estimates' dispersion across turns in the BLT (solid line) and OCT (dashed line). The rates are obtained from estimating the model $\text{ED}_{i,t} = \alpha_1 + \beta_1 \text{Dummy}_i + \beta_2 \ln t + \beta_3 \text{Dummy}_i \ln t + \epsilon_{i,t}$, treating the data as a panel (market instances and turns). $\text{ED}_{i,t}$ is the estimates' dispersion in turn t of observation i, Dummy_i takes the value 1 if observation i is obtained in the OCT and 0 otherwise, and t is the turn number. Robust standard errors are reported in square brackets.

turn. In the OCT, estimates' dispersion level starts at around 17.0 and decreases to about 9.7 by the fourth turn. To obtain a field for the magnitude of the results, recall that in each instance of the game, there is exactly one subject who receives a perfect signal and another subject who receives an imperfect signal. Since the imperfect signal is uniformly distributed around the liquidating value with bounds of ± 30, if subjects did not aggregate information at all, estimates' dispersion would have been constant at 15.

I test the statistical significance of the null that the level of estimates' dispersion is constant across subsequent turns using a nonparametric Wilcoxon matched-pairs signed-ranks test. The p-values from these tests are reported in rows 4–6 of Table 2 for the BLT and rows 10–12 of that table for the OCT. In the BLT, the level of estimates' dispersion decrease from turn one to turn three but not after that, in line with the rational model. In the OCT, I find a similar

Table 2
Estimates' dispersion across turns and treatments

Turn		1	2	3	4
		BLT			
			(average levels)		
Observed ED		17.57	7.51	5.62	5.00
Fully rational ED		14.95	4.03	0.00	0.00
			(p-values)		
$ED_1 = ED_2$	[0.000]				
$ED_2 = ED_3$	[0.000]				
$ED_3 = ED_4$	[0.280]				
$ED^{Observed} = ED^{Rational}$		[0.000]	[0.000]	[0.000]	[0.000]
		OCT			
			(average levels)		
Observed ED		16.97	11.95	13.19	9.74
Fully rational ED		15.34	3.80	0.00	0.00
			(p-values)		
$ED_1 = ED_2$	[0.000]				
$ED_2 = ED_3$	[0.358]				
$ED_3 = ED_4$	[0.000]				
$ED^{Observed} = ED^{Rational}$		[0.009]	[0.000]	[0.000]	[0.000]
$ED_t^{BLT} = ED_t^{OCT}$		[0.470]	[0.000]	[0.000]	[0.000]

This table reports the average level of the estimates' dispersion (ED) and the p-values (in square brackets) from nonparametric tests comparing the observed levels across turns, across treatments, and with the predicted levels derived from the fully rational model. The top/bottom panel reports the results for the BLT/OCT. The first two rows in each panel report the average ED level, for each turn separately, observed in the data and predicted by the fully rational model. Rows 3–5 of each panel report the p-values when testing the null that the observed ED level is constant across subsequent turns. The last row of each panel reports p-values associated with testing the null that observed and predicted levels are equal. The last row of the table reports the p-values resulting from the equality test across treatments.

pattern, delayed by a turn: there is a significant drop going from turn one to turn two and from turn three to turn four but not in between.

To quantify the convergence rate and measure how it is affected by the treatment, I estimate the following regressions model (treating the data as a panel), which interacts a treatment dummy with the log of the turn number:

$$ED_{i,t} = \alpha_1 + \beta_1 \, Dummy_i + \beta_2 \ln t_i + \beta_3 \, Dummy_i \ln t_i + \epsilon_{i,t}, \qquad (6)$$

where $Dummy_i$ takes the value of 1 if the observation was obtained under the OCT and 0 otherwise, and t is the turn number. This specification allows me to test parametrically the following hypotheses.

- Estimates' dispersion does not decrease across turns, i.e., $\beta_2 = 0$ and $\beta_2 + \beta_3 = 0$.
- There is no difference, across treatments, in the initial level of estimates' dispersion, i.e., $\beta_1 = 0$.
- There is no difference across treatments in the rate of estimates' dispersion convergence, i.e., $\beta_3 = 0$.

The estimation results are reported and depicted in Figure 4. I find that in both treatments, the initial estimates' dispersion is virtually identical; $\widehat{\beta}_1$ is not statistically different from zero. This is consistent with the theoretical predictions in that beliefs about signal precision should not enter players' estimates in the first turn. Convergence in estimates occurs in both treatments; $\widehat{\beta}_2$ and $\widehat{\beta}_2 + \widehat{\beta}_3$ are statistically different from zero. However, the rate of convergence is substantially lower in the OCT compared with the BLT, as $\widehat{\beta}_3$ is positive and statistically different from zero.

Consistent with these results, I find that the *level* of estimates' dispersion is different across treatments. Rows 1 and 7 of Table 2 show that estimates' dispersion is higher in the OCT than in the BLT, with the exception of the first turn. As predicted by the model (see Proposition 5), subjects' confidence does not enter their first turn estimates, and thus, no difference is observed. In turns two to four, estimates' dispersion in the OCT is roughly twice as high as in the BLT. Nonparametric tests confirm these observations. The last row in Table 2 reports the p-values of testing for treatment effects.

4.2.2 Excess level of estimates' dispersion. The analysis so far suggests that subjects form estimates that are further apart in the OCT than in the BLT, and that while these estimates get closer together as a result of interaction, the rate of convergence is lower in the OCT. These findings are consistent with the results discussed in Section 4.1, which suggest that subjects in the OCT are overconfident while those in the BLT are not.

But, are the results in the BLT consistent with the level predictions of the fully rational model? To answer that, I use the information held by subjects in each round to compute turn-by-turn predicted estimates and estimates' dispersion. Rows 2 and 8 of Table 2 report the level of estimates' dispersion derived from the fully rational model discussed in Section 2.2. This model serves as a baseline, and thus does not allow for the possibility of either overconfidence or errors. Rows 6 and 12 in Table 2 report the p-values from testing equality between fully rational predictions and observed levels for the BLT and OCT, respectively. It may not be surprising to find that in the OCT, estimates' dispersion levels are in excess of those predicted by the fully rational model. After all, one would expect these results from subjects' overconfidence. It is more surprising to find that estimates' dispersion levels are higher than predicted in the BLT too, while I find no evidence of overconfidence among subjects in that treatment.

I argue that these deviations from the fully rational predictions are consistent with subjects' best response to errors in others' estimates. To explore this possibility, I approximate the best response function, discussed in Section 2.4, which suggested that

$$a^*_{i,2} = q_i s_i + (1 - q_i)(a_{j,1} + \Theta(s_i, a_{j,1}, \cdot)), \tag{7}$$

where $\Theta(s_i, a_{j,1}, \cdot)$ is a term that adjusts the observed estimate, $a_{j,1}$, based on player i's signal, player j's estimate, and player j's expected estimate error (see Figure 1 for an illustration). Since the functional form of Θ interacts observed variables with unobserved variables in a way that does not allow for linear estimation, I approximate this function using a cubic function around $s_i - a_{j,1}$. This allows estimation of the following linear model:

$$a_{i,n,2} = \alpha + \beta_1 s_{i,n} + \beta_2 a_{j,n,1} + \beta_3 (s_i - a_{j,1})^3 + \epsilon_{i,n}. \tag{8}$$

The model in Equation (5), which allows for the possibility that subjects are overconfident, I also extended to account for the possibility that subjects "discount" others' estimates. The coefficient β_3 captures the steepness of this "discount." Under the null, β_3 is equal to zero.

Columns 3 and 4 in Table 1 report the result for the BLT and OCT. In both treatments, I find evidence consistent with rational response to errors. The coefficient associated with the term $s_i - a^3_{j,1}$ is different from zero at the 1% level. Further, this coefficient is positive as predicted by the rational response to errors alternative. At the same time, this channel does not eliminate the overconfidence observed in the OCT, as $\widehat{\beta}_1$ is substantially higher in the OCT, compared with the BLT, and exceeds 0.5, indicating that subjects assign too high a weight to their own signal in the OCT. In other words, I am able to separate overconfidence from the rational response to errors and show that both affect subjects' estimates.

5. Summary

In this article, I suggest a simple information aggregation game through which I study—theoretically and empirically—how participants aggregate multidimensional private information together with information learned indirectly from observing others' actions. In order to distinguish two widespread behavioral biases, erroneous actions, and mistaken beliefs, I combine an experimental design, which controls for the presence of overconfidence, and a model that nests both biases. By doing so, I am able to measure subjects' confidence, in both variants of the experiment, along with their responses to errors.

I am able to estimate subjects' implicit overconfidence. I find that the treatment induces overconfidence in one set of sessions but not in the other. Taking advantage of this difference, I establish the effect of overconfidence by showing that it gives rise to more dispersed estimates and to a lower rate of estimates' convergence. At the same time, I find support for the second form of deviation from full rationality errors in actions. In both treatments, subjects appear to respond to others' mistakes in forming their own estimates.

This article contributes to the existing literature in a number of ways. First, it suggests a theoretically tractable mechanism that can be used to study how individual biases affect information aggregation. This can be pursued further by changing the structure of information: one can introduce public announcements alongside private signals and replace exogenous endowment of private signals with an endogenous acquisition of information.[17] Second, it provides a laboratory implementation in which subjects are endowed with private signals *and* private signal precision, and shows how to turn the latter into a treatment variable. Third, it shows how the fully rational model can be extended to incorporate different deviations from rationality and how it can be used for empirical identification.

Appendix

A.1 Proofs

Proof of Proposition 1. Assume that player j follows $a^*_{j,t}$. I will show that player i's dominant strategy is also characterized by $a^*_{i,t}$. In turn three, $I_{i,3} = \{s_i, q_i, a^*_{j,1}, a^*_{j,2}\} = \{s_i, q_i, s_j, q_j\}$ since $a^*_{j,1} = s_j$, and since $\{a^*_{j,1}, a^*_{j,2}, a_{i,1}\}$ can be used by player i to back out q_j. Hence, player i has full information and she can do no better to submit $a^*_{i,3} = E[v|I_{i,3}] = \text{Ind}_{(q_i > \frac{a_{j,2}-a_{i,1}}{a_{j,1}-a_{i,1}})} s_i + \text{Ind}_{(q_i < \frac{a_{j,2}-a_{i,1}}{a_{j,1}-a_{i,1}})} a_{j,1} + \text{Ind}_{(q_i = \frac{a_{j,2}-a_{i,1}}{a_{j,1}-a_{i,1}})} (\frac{s_i + a_{j,1}}{2}) = \text{Ind}_{(q_i > q_j)} s_i + \text{Ind}_{(q_i < q_j)} s_j + \text{Ind}_{(q_i = q_j)} (\frac{s_i + s_j}{2}) = v \;\; a.s.$ In turn two, $I_{i,2} = \{s_i, q_i, s_j\}$. Given that in turn three, player i knows the fundamental value ($a.s.$), he can do no better than act myopically in turn two. Therefore, $a^*_{i,2} = E[v|I_{i,2}] = \Pr(h|I_{i,2})E(v|I_{i,2}, h) + (1 - \Pr(h|I_{i,2}))E(v|I_{i,2}, l) = q_i s_i + (1 - q_i)s_j$. In turn one, player i will act myopically and submit $a^*_{i,1} = E[v|I_{i,1}] = \Pr(h|I_{i,1})E(v|I_{i,1}, h) + (1 - \Pr(h|I_{i,1}))E(v|I_{i,1}, l) = q_i s_i + (1 - q_i)s_i = s_i$ since by deviating she would reduce own payoffs without changing the information that can be extracted from player j's actions.

Proof of Proposition 2. I use backward induction for this proof. Since the myopic best response equilibrium maximized expected payoffs at each turn of the game separately, the player would deviate from it only if they could increase future payoffs. Since turn three is the last turn, it follows trivially that $a_{i,3} = a^*_{i,3}$. In turn two, assume that $a_{i,2} \neq a^*_{i,2}$. Since $E(u_{i,2}(a_{i,2})) < E(u_{i,2}(a^*_{i,2}))$, it must be the case that $E(u_{i,3}(a^*_{i,3}(I(a_{j,2}(a_{i,2}))))) > E(u_{i,3}(a^*_{i,3}(I(a^*_{j,2}(a_{i,2})))))$ but since actions are submitted simultaneously, this cannot hold. Thus, in turn two, $a^*_{i,2} = E(v|I_{i,2})$. In turn one, assume that $a_{i,1} \neq a^*_{i,1} \Rightarrow E(u_{i,2}(a_{i,1})) < E(u_{i,2}(a^*_{i,1}))$, so it must be the case that $E(u_{i,3}(a^*_{i,3}(I(a_{j,2}(a_{i,1}))))) > E(u_{i,3}(a^*_{i,3}(I(a_{j,2}(a^*_{i,1})))))$ but since turn three actions arrive at full information revelation (a.s.), this cannot hold.

Proof of Proposition 3.

$\text{ED}_1 = |s_i - s_j| = |e_j| > 0 \; a.s.$
$\text{ED}_2 = |q_i s_i + (1 - q_i)s_j - (1 - q_j)s_i - q_j s_j| = |(1 - q_i - q_j)s_j - (1 - q_i - q_j)s_i|$
$= |(1 - q_i - q_j)(s_j - s_i)|.$
Since $-1 \leq (1 - q_i - q_j) \leq 1$ and $-e_j \leq (s_j - s_i) \leq e_j$, we get
$|(1 - q_i - q_j)e_j| < |e_j| \; a.s.$ (notice that $q_i + q_j$ need not equal 1).

Also, since $a^*_{i,3} = a^*_{j,3}$, $\text{ED}_3 = 0$.
Thus, $\text{ED}_1 > \text{ED}_2 > \text{ED}_3 = 0$.

[17] This topic has been of interest to a large body of empirical literature that documents underreaction to public announcements, such as earning releases, and the theoretical literature that suggest a role for overconfidence in this (e.g., Bondt and Thaler 1985 and Jegadeesh and Titman 1993).

Proof of Proposition 4. Replacing q_i with $\widetilde{q}_i$ and q_j with $\widetilde{q}_j$, and following the Proposition 3 proof, I get that

- $\mathrm{ED}_1^M = |s_i - s_j| = |e_j|$;
- $\mathrm{ED}_2 = |1 - \widetilde{q}_i s_i - \widetilde{q}_j)||s_j - s_i| = |1 - (q_i + q_j) - (\mathrm{MC}_i + \mathrm{MC}_j)||e|$;
- $\mathrm{ED}_3^M = 0\ (a.s.)$.

Proof of Proposition 5. For turns one and three, the proof is trivial.

For turn two, recall that

$\mathrm{ED}_2^M = |1 - (q_i + q_j) - (\mathrm{MC}_i + \mathrm{MC}_j)||e| \propto |1 - (q_i + q_j) - (\mathrm{MC}_i + \mathrm{MC}_j)|$, since $|e| \geq 0$.

Denoting $(q_i + q_j) \equiv q_{ij}$ and $(\mathrm{MC}_i + \mathrm{MC}_j) \equiv \mathrm{MC}_{ij}$, and squaring both sides of the expression, I get

$$\left(\mathrm{ED}_2^M\right)^2 = (1 - q_{ij} - \mathrm{MC}_{ij})^2 (s_j - s_i)^2 = (1 - q_{ij} - \mathrm{MC}_{ij})^2 (e_j)^2.$$

To find the parameter value ranges for which the dispersion of estimates is increasing, take a derivative with respect to

$$\mathrm{MC}_{ij} : \frac{d\left(\mathrm{ED}_2^M\right)^2}{d\mathrm{MC}_{ij}} = -2(e_j)^2 (1 - q_{ij} - \mathrm{MC}_{ij}),$$

which is increasing if $q_{ij} + \mathrm{MC}_{ij} - 1 > 0$. Since in expectations $q_{ij} = 1$, $E(\mathrm{ED}_2^M)$ is increasing in MC_{ij} if players are on average overconfident.

A.2. Derivation of optimal actions for the general case

In this section, I derive players' optimal actions in each turn of the game under the general case. Since players' strategies are symmetric, I abuse notation and drop players' identification subscript when appropriate.

In turn one, optimal actions are

$$E[v|s_i, q_i] = \Pr(h|s_i, q_i) E(v|s_i, q_i, h) + (1 - \Pr(h|s_i, q_i)) E(v|s_i, q_i, l).$$

Computing each of these terms separately, I get

$$\Pr(h|q_i, s_i) = q \Big/ \left(\frac{2Y}{\min(s_i + Y, H) - \max(s_i - Y, L)}\right) = \frac{q_i(\min(s_i + Y, H) - \max(s_i - Y, L))}{2Y}.$$

$$E(v|s_i, q_i, h) = s_i . E[v|s_i, q_i, l] = \int_{\max(s_i - Y, L)}^{\min(s_i + Y, H)} \frac{1}{\min(s_i + Y, H) - \max(s_i - Y, L)} x dx$$
$$= \frac{1}{2}(\max(s_i - Y, L) + \min(Y + s_i, H)).$$

That is, conditional on being a type l player, v is distributed uniformly around the signal s_i. The intuition is that in the interior case ($s_i \in [L + Y, H - Y]$), player i is uniformly distributed around v *and* v is drawn uniformly. Around v's distribution boundaries, the distribution mass is being reallocated uniformly.

To summarize,

$$a_1^* = \frac{q(\min(s + Y, H) - \max(s - Y, L))}{2Y} s + \left(1 - \frac{q(\min(s + Y, H) - \max(s - Y, L))}{2Y}\right)$$
$$\times \left(\frac{1}{2}(\max(s - Y, L) + \min(Y + s, H))\right),$$

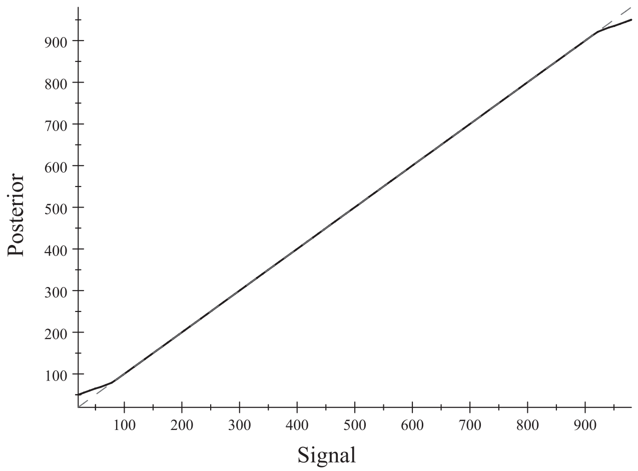

Figure A.2
Turn one posterior
The figure depicts an example for the turn one posterior a player forms given the signal they receive.

which in the interior case ($s_i, s_j \in [L+Y, H-Y]$) simplifies to

$$a_1^* = s.$$

Figure A.2 depicts the posterior as a function of the signal, where v ranges from 50 to 950, signal dispersion is set to 30 (i.e., $Y = 30$), and $q = 2/3$. For most of the domain, the posterior is equal to the signal. If the signal is sufficiently close to the "edges," the posterior tapers off. It is important to note that for all possible signals, the posterior is a monotonic function of the signal.

At the end of turn one, $a_{i,1}$ and $a_{j,1}$ are announced. Therefore, turn two optimal actions are

$$\begin{aligned} a_{i,2}^* &= E(v|I_{i,2}) = E(v|\{s_i, q_i, a_{j,1}^*\}) = \Pr(h|s_i, q_i, a_{j,1}^*)s_i + (1 - \Pr(h|s_i, q_i, a_{j,1}^*)) \\ &\quad \times E[v|l, s_i, q_i, a_{j,1}^*], \end{aligned}$$

where

$$\Pr(h|s_i, q_i, a_{j,1}^*) = \int_{\max(s_i - Y, L)}^{\min(s_i + Y, H)} \Pr(h|s_i, q_i, s_j) \Pr(s_j|a_{j,1}^*(s_j)^{-1}) ds_j$$

and $a_{j,1}^*(s_j)^{-1}$ is the inverse function of player j's turn one actions with respect to her q_j.

Therefore,

$$\begin{aligned} a_{i,2}^* &= E[v|l, s_i, q_i, a_{j,1}^*] \\ &= \int_{\max(s_i - Y, L)}^{\min(s_i + Y, H)} \left(\int_{\min(s_i, s_j)}^{\max(s_i, s_j)} (\Pr(v|l, s_i, q_i, a_{j,1}^*)v) dv \right) ds_j \\ &= \int_{\max(s_i - Y, L)}^{\min(s_i + Y, H)} \left(\int_{\min(s_i, s_j)}^{\max(s_i, s_j)} (\Pr(v|l, s_i, q_i, s_j) \Pr(s_j|a_{j,1}^*(s_j)^{-1}) v) dv \right) ds_j. \end{aligned}$$

Since $a_{j,1}^*(s_j)^{-1}$ cannot be explicitly derived, the close-form solution to $a_{i,2}^*$ is difficult to obtain.

By the end of turn two, all information is revealed ($a.s.$). Therefore, optimal action in turn three is

$$a^*_{i,3} = \text{Ind}_{(q_i > \frac{a_{j,2}-a_{i,1}}{a_{j,1}-a_{i,1}})} s_i + \text{Ind}_{(q_i < \frac{a_{j,2}-a_{i,1}}{a_{j,1}-a_{i,1}})} a_{j,1} + \text{Ind}_{(q_i = \frac{a_{j,2}-a_{i,1}}{a_{j,1}-a_{i,1}})} \left(\frac{s_i + a_{j,1}}{2}\right), a^*_{j,3}.$$

A.3. Derivation of turn two optimal actions

Generally, it is easy to show that if x, y, z are r.v., then

$$\Pr(x|y,z) = \frac{\Pr(y|x,z)\Pr(x|z)}{\Pr(y|z)}.$$

Therefore,

$$\Pr(s_2|a_{21},s_1) = \frac{\Pr(a_{21}|s_2,s_1)\Pr(s_2|s_1)}{\Pr(a_{21}|s_1)}.$$

Calculating the elements of this expression, I get

- $\Pr(a_{2,1}|s_2,s_1) = \Pr(a_{2,1}|s_2) = \phi(a_{21}-s_2;0,\sigma_1)$;
- $\Pr(s_2|s_1) = \frac{1}{2Y}$.

However,

$$\Pr(a_{21}|s_1) = \int_{s_1-Y}^{s_1+Y} \frac{1}{2Y}\phi(a_{21}-s_2;0,\sigma_1)ds_2$$

$$= \frac{\sqrt{4\pi}\,\text{erf}\left(\frac{1}{\sigma_1}\left(\frac{1}{2}Y\sqrt{2}-\frac{1}{2}s_1\sqrt{2}+\frac{1}{2}a_{21}\sqrt{2}\right)\right)}{\frac{1}{8\sqrt{\pi}Y}} - \frac{\sqrt{4\pi}\,\text{erf}\left(\frac{1}{\sigma_1}\left(\frac{1}{2}a_{21}\sqrt{2}-\frac{1}{2}s_1\sqrt{2}-\frac{1}{2}Y\sqrt{2}\right)\right)}{\frac{1}{8\sqrt{\pi}Y}}$$

$$= \frac{1}{4Y}\left(\text{erf}\left(\frac{\sqrt{2}}{2\sigma_1}(Y-s_1+a_{21})\right) - \text{erf}\left(\frac{\sqrt{2}}{2\sigma_1}(a_{21}-s_1-Y)\right)\right).$$

Collecting these terms, I obtain

$$E(s_2|a_{21},s_1) = \int_{s_1-Y}^{s_1+Y} s_2 \Pr(s_2|a_{21},s_1)ds_2$$

$$= \int_{s_1-Y}^{s_1+Y} s_2 \frac{\phi(a_{21}-s_2;0,\sigma_1)\frac{1}{2Y}}{\frac{1}{4Y}\left(\text{erf}\left(\frac{\sqrt{2}}{2\sigma_1}(Y-s_1+a_{21})\right)-\text{erf}\left(\frac{\sqrt{2}}{2\sigma_1}(a_{21}-s_1-Y)\right)\right)}ds_2$$

$$= \frac{-\sqrt{2}\left(\frac{1}{2}\sigma_1\sqrt{2}e^{-\frac{1}{2}\frac{a_{21}^2}{\sigma_1^2}}\left(-\sigma_1\sqrt{2}\exp\left(\frac{a_{21}}{\sigma_1^2}(Y+s_1)-\frac{1}{2\sigma_1^2}(Y+s_1)^2\right)-\sqrt{\pi}a_{21}e^{\frac{1}{2}\frac{a_{21}^2}{\sigma_1^2}}\,\text{erf}\left(\frac{1}{2}\frac{a_{21}}{\sigma_1}\sqrt{2}-\frac{1}{2\sigma_1}(Y+s_1)\sqrt{2}\right)\right)}{\sqrt{\pi}\sigma_1\,\text{erf}\left(\frac{1}{\sigma_1}\left(\frac{1}{2}a_{21}\sqrt{2}-\frac{1}{2}s_1\sqrt{2}-\frac{1}{2}Y\sqrt{2}\right)\right)-\sqrt{\pi}\sigma_1\,\text{erf}\left(\frac{1}{\sigma_1}\left(\frac{1}{2}Y\sqrt{2}-\frac{1}{2}s_1\sqrt{2}+\frac{1}{2}a_{21}\sqrt{2}\right)\right)}$$

$$+\frac{\sqrt{2}\left(\frac{1}{2}\sigma_1\sqrt{2}e^{-\frac{1}{2}\frac{a_{21}^2}{\sigma_1^2}}\left(-\sigma_1\sqrt{2}\exp\left(\frac{a_{21}}{\sigma_1^2}(s_1-Y)-\frac{1}{2\sigma_1^2}(s_1-Y)^2\right)-\sqrt{\pi}a_{21}e^{\frac{1}{2}\frac{a_{21}^2}{\sigma_1^2}}\,\text{erf}\left(\frac{1}{2}\frac{a_{21}}{\sigma_1}\sqrt{2}-\frac{1}{2\sigma_1}\sqrt{2}(s_1-Y)\right)\right)\right)}{\sqrt{\pi}\sigma_1\,\text{erf}\left(\frac{1}{\sigma_1}\left(\frac{1}{2}a_{21}\sqrt{2}-\frac{1}{2}s_1\sqrt{2}-\frac{1}{2}Y\sqrt{2}\right)\right)-\sqrt{\pi}\sigma_1\,\text{erf}\left(\frac{1}{\sigma_1}\left(\frac{1}{2}Y\sqrt{2}-\frac{1}{2}s_1\sqrt{2}+\frac{1}{2}a_{21}\sqrt{2}\right)\right)}$$

$$= a_{21} + \frac{e^{-\frac{1}{2}\frac{a_{21}^2}{\sigma_1^2}}\left(\sigma_1\sqrt{2}\left(\exp\left(\frac{a_{21}}{\sigma_1^2}(s_1-Y)-\frac{1}{2\sigma_1^2}(s_1-Y)^2\right)-\exp\left(\frac{a_{21}}{\sigma_1^2}(Y+s_1)-\frac{1}{2\sigma_1^2}(Y+s_1)^2\right)\right)\right)}{-\sqrt{\pi}\,\text{erf}\left(\frac{\sqrt{2}}{2\sigma_1}(a_{21}-s_1-Y)\right)+\sqrt{\pi}\,\text{erf}\left(\frac{\sqrt{2}}{2\sigma_1}(Y-s_1+a_{21})\right)}$$

$$= a_{21} + \frac{2\sigma_1^2(\phi(a_{21};s_1-Y,\sigma_1)-\phi(a_{21};s_1+Y,\sigma_1))}{\text{erf}\left(\frac{\sqrt{2}}{2\sigma_1}(Y-s_1+a_{21})\right)-\text{erf}\left(\frac{\sqrt{2}}{2\sigma_1}(a_{21}-s_1-Y)\right)}.$$

References

Akerlof, G. A., and J. L. Yellen. 1985. Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria? *American Economic Review* 75(4):708–20.

Biais, B., D. Hilton, K. Mazurier, and S. Pouget. 2002. Psychological Dispositions and Trading Behaviour. Working paper, CEPR.

Bikhchandani, S., D. Hirshleifer, and I. Welch. 1998. Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades. *Journal of Economic Perspectives* 12(3):151–70.

Bondt, W. F. M. D., and R. Thaler. 1985. Does the Stock Market Overreact? *Journal of Finance* 40(3):793–805.

Camerer, C. F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.

Camerer, C. F., and D. Lovallo. 1999. Overconfidence and Excess Entry: An Experimental Approach. *American Economic Review* 89(1):306–18.

Cooper, A. C., C. Y. Woo, and W. C. Dunkelberg. 1988. Entrepreneurs' Perceived Chances for Success. *Journal of Business Venturing* 3(2):97–108.

Daniel, K., D. Hirshleifer, and A. Subrahmanyam. 1998. Investor Psychology and Security Market Under- and Overreactions. *Journal of Finance* 53(6):1839–85.

Deaves, R., E. Lüders, and G. Y. Luo. 2005. An Experimental Test of the Impact of Overconfidence and Gender on Trading Activity. CoFE Discussion Paper 05-07, Center of Finance and Econometrics, University of Konstanz.

Diamond, D. W., and R. E. Verrecchia. 1981. Information Aggregation in a Noisy Rational Expectations Economy. *Journal of Financial Economics* 9(3):221–35.

Fischhoff, B., P. Slovic, and S. Lichtenstein. 1977. Knowing with Certainty: The Appropriateness of Extreme Confidence. *Journal of Experimental Psychology* 3(4):552–64.

Friedman, D., and S. Sunder. 1994. *Experimental Methods: A Primer for Economists*. Cambridge, MA: Cambridge University Press.

Glaser, M., and M. Weber. 2007. Overconfidence and Trading Volume. *The Geneva Risk and Insurance Review* 32:1–36.

Goeree, J. K., and C. A. Holt. 1999. Stochastic Game Theory: For Playing Games, Not Just for Doing Theory. *Proceedings of the National Academy of Sciences of the United States* 96:10,564–567.

Harris, M., and A. Raviv. 1993. Differences of Opinion Make a Horse Race. *Review of Financial Studies* 6(3):473–506.

Jegadeesh, N., and S. Titman. 1993. Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *Journal of Finance* 48(1):65–91.

Kahneman, D., P. Slovic, and A. Tversky (eds.). 1982. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge, MA: Cambridge University Press.

Kandel, E., and N. D. Pearson. 1995. Differential Interpretation of Public Signals and Trade in Speculative Markets. *Journal of Political Economy* 103(4):831–72.

Kirchler, E., and B. Maciejovsky. 2002. Simultaneous Over- and Underconfidence: Evidence from Experimental Asset Markets. *Journal of Risk and Uncertainty* 25(1):65–85.

Larwood, L. 1978. Swine Flu: A Field Study of Self-Serving Biases. *Journal of Applied Social Psychology* 8(3):283–89.

Larwood, L., and W. Whittaker. 1977. Managerial Myopia: Self-Serving Biases in Organizational Planning. *Journal of Applied Psychology* 62:194–98.

Mckelvey, R. D., and T. R. Palfrey. 1995. Quantal Response Equilibria in Normal Form Games. *Games and Economic Behavior* 10:1–5.

Odean, T. 1998. Volume, Volatility, Price, and Profit When All Traders Are Above Average. *Journal of Finance* 53(6):1887–934.

Svenson, O. 1981. Are We All Less Risky and More Skillful Than Our Fellow Drivers? *Acta Psychologica* 47(2):143–48.

Varian, H. R. 1986. Differences of Opinion in Financial Markets. *Proceeding of the Federal Reserve Bank of St. Louis* 3–40.