Analysts, Incentives and Exaggeration

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Timothy Wayne Shields

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

October, 2007

Abstract

Sell-side analysts are compensated, at least in part, by brokerage commissions. These commissions create an incentive to bias forecasts to generate trade. Thus, analysts have clear economic incentives to deceive and traders have economic incentives to detect deception. Prior analytical theories of information transmission games starkly predict that there will always be some deception (with trade) at best and uninformative messages (and no trade) at worst when the sender's and receiver's incentives are not aligned. Prior experimental evidence of information transmission games shows senders do elect to deceive, although they send messages more informative than theory predicts. Likewise, receivers rely more upon messages than theory predicts.

Can behavior that deviates from prediction be explained by normative social behavior, such as trust and honesty? Alternatively, are subjects boundedly rational, failing to sufficiently consider other players' incentives when predicting their decisions? To answer these questions, I design and conduct an experiment to investigate whether forecasting and trading behaviors are best explained by analytical theory, limited strategic sophistication, or social norms. The experimental results confirm a minority of subjects adopt honest forecasting strategies, but at the same time, a majority of subjects adopts trusting trading strategies. Additionally, subjects do not appear to revise trading behavior despite evidence of deceptive forecasts. The results suggest subjects' behavior within the setting is better explained by a framework of hierarchical reasoning than by social normative behavior or analytic theory.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

"Honesty is the best policy - when there is money in it." - Mark Twain

"Though I am not naturally honest, I am so sometimes by chance." - William Shakespeare

Wall Street security analysts garnered attention after the burst of the technology-stock bubble. Analysts were alleged to inflate recommendations about the future prospects of firm in order to secure or maintain investment-banking relationships (Brennan (2004); Michaely and Womack (1999)). During this timeframe, the rules of the NYSE and NASD required analysts, in some circumstances, to disclose certain conflicts of interest when recommending the purchase or sale of a specific security or issuing a forecast.

On May 10, 2002, the SEC approved proposed changes to these rules, increasing the disclosures that analysts and brokerage firms must make. These efforts sought to make transparent conflicts of interest. For example, the rule changes prohibit research analysts from being supervised by the investment-banking department and also bar securities firms from tying an analyst's compensation to specific investment banking transactions. Furthermore, if an analyst's compensation is based on the firm's general investment banking revenues, that fact must be disclosed in the firm's research reports.

Sell-side analysts are compensated, at least in part, by brokerage commissions. The SEC cites brokerage commissions as a potential conflict of interest, stating "(An) analyst report can help firms make money indirectly by generating more purchases and sales of covered securitie–which, in turn, result in additional brokerage commissions (U.S. Securities and Exchange Commission (2005))." This suggests forecasts that move the market are possibly attractive to analysts.

I construct an analytical model of reporting and trading where there are potential gains from trading. In certain states of nature it may be advantageous for two shareholders to trade with each other, one selling a portion of his shares to other. I construct the model such that these gains are increasing as the state deviates from the expected. If the state deviates enough, then the potential gains may exceed the dead-weight costs of trading. Before the state is realized and trading occurs an analyst receives a signal on the state and releases a public forecast.

For example, imagine there are four states A, B, C and D having equal probability. In states A and D the gains to trading exceed the costs of trading, and in states B and C the costs exceed the gains. The analyst receives a perfect signal upon the state (I relax this assumption later). The model is similar to Crawford and Sobel (1982), in that the interest of the analyst (sender) is not always aligned with the shareholder (receiver).

This analytical theory predicts a stark result: in a finite game every forecast will induce trade, or no forecast will induce trade. The intuition is straight-forward. Using the example above, imagine the analyst's reporting decision in the last period of the finite game where the shareholders believe the analyst will release a forecast equal to her signal. In order to generate trade and thus commissions, the analyst will always release a forecast that induces trade, and will never release a forecast of B or C. Sequentially-rational shareholders conjecture this behavior and trade, or not, dependent upon trading costs. If trading costs are small enough, then the analyst's reporting strategy always induces trade, otherwise, there is no trade. The analyst anticipates the shareholders' behavior and releases a forecast in the next to last period that induces trade. Using backwards induction we find the analyst will never release a forecast that does not induce trade.

However, prior experimental studies show that some subjects overcome pareto-suboptimal predictions similar to my analytical prediction of no trade (King-Casas et al. (2005); Berg et al. (1995); Hoffman et al. (1998); Engle-Warnick and Slonim (2004)). When subjects interacted in small groups and information was publicly known, subjects adopted cooperative strategies generally inconsistent with, but more profitable than, the non-cooperative reporting strategies analytical theory predicts. Prior work has shown that when sender-subjects signal their intentions, then user-subjects reciprocate (McCabe et al. (2003)). In much of these studies, the subjects are better off economically if they can cooperate. Is subjects observed behavior a result of heuristic behavior, or driven by social norms adopted in order to reap monetary benefits otherwise unavailable?

Experimental studies have shown when there is private information, or when coordination is difficult to implement, subjects tend to adopt non-cooperative or deceptive strategies (Hoffman et al. (1998)). Furthermore, several studies have shown that senders of information do adopt deceptive strategies when the interests of the receiver and sender diverge (Blume et al. (1998); Dickhaut et al. (1995); Cai and Wang (2006)). However, while subjects have shown a willingness to deceive others in strategic transmission games, a phenomenon of overcommunication has been documented. Senders send more information than predicted by analytical theory and receivers rely more upon the sent information than predicted (Cai and Wang (2006); Sanchez-Pages and Vorsatz (2006); Wang et al. (2006)).

The above work suggests the possibility of people who choose to be honest, and others who choose to trust. Theoretical work by Crawford (2003) examines the presence of habitually honest, deceptive, trusting, skeptic, and strategic players in a one-shot setting. Other experimental research uses a related framework of behavioral types to explain the departure of experimental results from analytical theory (Nagel (1995); Costa-Gomes et al. (2000); Cai and Wang (2006)). Other work also postulates that some subjects are non-strategic in that the subjects appear to make no attempt to use other players' incentives to predict their decisions. Using behavioral type analysis, these authors find better fit to subject behavior than game theoretic predictions for one-shot games (Costa-Gomes and Crawford (2004); Camerer et al. (2004); Stahl and Wilson (1994)).

In order to examine the overcommincation phenomenon, I design a computerized economic experiment with two treatments. In both treatments, the analyst receives a signal on the state of nature and releases a forecast to two shareholders who can then trade. If there is trade the analyst receives a commission. There are gains to trading available in half of the potential states. The grouping and game is repeated eight times. I examine whether behavioral type analysis, specifically the existence of non-strategic players, can explain behavior in a repeated game, or whether behavior is better explained by social norms such as trust and honesty.

In the first treatment the trading costs are set such that all forecasts could induce trade. This is theoretically possible as expected gains to trade are available even if analysts use deceptive strategies. Using the said example, imagine the analyst only releases a forecast of A or D, but shareholders interpret the forecast as the state is an element of A or B or C or D. If the costs of trading are small enough, then the shareholders are willing to trade, as there are expected gains.

In the second treatment the trading costs are set so the analyst cannot theoretically adopt a deceptive strategy that always induces trade. Returning to the example, if the analyst releases forecast of A, but this is understood by sequentially rational shareholders to mean the state is either A or B. If the net gains from trading in state A are overwhelmed by the loss of trading in state B, then this deceptive strategy cannot induce trade. So the game-theoretic prediction is that no forecast will induce trade–despite that half the time there are potential gains to trading.

In a repeated game, are subjects non-strategic or can unpredicted behavior be explained by social norms of truth telling and trust? If some subjects are non-strategic, then I posit subjects use non-adaptive heuristics; that is, they play the same strategy repeatedly even if it results in economic detriment. If subjects instead are motived by social norms of honesty and trust, I posit behavior will be contingent upon the other subjects' past behavior.

The experimental results support the notion underlying behavioral type analysis: people act differentially. A majority of subjects tend to use deceptive forecasting strategies when analyst, but a minority are generally honest. Despite the prevalence of deceptive forecasting strategies, most subjects tend to use trusting trading strategies, relying upon forecasts more than appropriate given the frequency of deceptive reporting. These trusting strategies are not revised when playing against deceptive analysts despite evidence the analyst is using a deceptive strategy. Taken together, these results suggest that normative social behavior is not driving the overcommunication phenomenon.

Organization of Thesis

The remainder of the thesis is organized as follows: In Chapter 2, I review and contrast related research to this thesis. In Chapter 3, I present a finite-state finitely-repeated model of information transmission and construct the fully-strategic equilibrium forecasting strategy. In Chapter 4, I describe the experimental design and procedures. In Chapter 5, I present a framework of bounded rationality, apply it to my experimental setting, and construct competing hypotheses of subjects' behavior. Last, in Chapter 6, I present experimental results and conclude.

Chapter 2

Related Research

My works builds upon theoretical and experimental work. The chapter is organized as follows: First, I review the predictions of theoretical models of information transmission and analyst reporting. Second, I introduce bounded rationality theoretical work that departs from the assumption that all agents are fully strategic. Third, I review the findings of experimental work that test the predictions of theoretical work with boundedly rational agents. Last, I review the findings of experimental work that test theories of information transmission.

2.1 Theoretical Models Using Sequentially-Rational Agents

My work is related to analytical models of strategic information transmission. The seminal theory is Crawford and Sobel's strategic information transmission (Crawford and Sobel (1982)). Related to my topic of analyst reporting, the prediction is that unless the analyst's and shareholders' interests are always aligned in all states of nature, there will be some deception by the analyst in her reporting. Unlike Crawford and Sobel, I assume that interests are aligned in disjoint regions: in states that are sufficiently lower or higher than the expected medium states.

My work differs from other theoretical models of analyst reporting including Chen et al. (2005), Trueman (1994), Ottaviani and Sorensen (2006), and Morgan and Stocken (2003); Guttman (2005). In the aforementioned papers, the analyst not only has private value relevant information, but some other characteristic that distinguishes her from other analysts. Unlike my work, in these models the precision or quality of the analysts information is not uniform across analysts, so the analyst's forecasting strategy also tries to signal this quality. In Chen, Francis, and Jiang, the analysts with greater predictive ability only forecast extreme signals and are silent when the signal indicates normal states. Similar to my work, the analyst cares about moving the market, that is, about changing shareholders' expectations of the firm's value. However, unlike my model, the analyst is allowed to pre-commit to a reporting strategy.

In Trueman's model, an analyst with lesser ability under-weighs her private information and thus tempers her forecast towards the ex-ante mean. This result is similar to work done by Ottaviani and Sorenson, where shareholders use the analyst's forecasts to access the analyst's predictive ability. Like Trueman, Ottaviani and Sorenson do not give the analyst any incentive to change shareholder's expectation of the underlying firm, only an incentive to increase shareholders' perception of the analysts predictive ability. Unlike my work, Ottaviani and Sorenson also predict when the analyst will issue a forecast given the presence of other analysts. The choice of when to forecast is also present in Guttman, where analysts with greater predictive ability choose to forecast earlier. Guttman introduces an exogenous desire to bias the forecast in the analyst's utility. Uncertainty of the form of the analysts utility is also present in Morgan and Stocken, where shareholders are uncertain about the analyst's incentives. In both papers, shareholders use the forecast to both infer the state and the bias.

2.2 Theoretical Models Using Boundedly Rational Agents

Other papers in the literature introduce bounded-ration players into the communication game. Crawford (2003) analyzes a binary single-period game where the two players have opposite interests. One player can costlessly signal her decision. Standard analysis predicts this communication is ignored. Crawford introduces behavioral types; agents whom always play the same strategy. There are senders who always tell the truth, senders who always lie, receivers who always believe the message is true, and receivers who always believe the message is false and invertible. By introducing behavioral types in addition to strategic types, Crawford finds that the strategic sender can successfully misrepresent her decision when the probability of behavioral types is high enough.

Ottaviani and Squintani (2006) introduce trusting receivers into a game of information transmission with partially aligned incentives. Trusting receivers always believe the sender's message is true. Ottaviani and Squintani find in equilibrium the senders message is inflated, the action taken by the receiver is biased, trusting receivers are deceived, and information transmitted in equilibrium is greater than predicted by a standard fully-strategic model. Unlike Crawford Crawford

(2003), there is only one behavioral type as all senders are strategic.

Chen (2006) also perturbs a model of information transmission by adding habitually trusting receivers and habitually honest senders in addition to strategic agents. This addition leads to a unique equilibrium where the strategic sender always distorts her messages and messages cluster around the top (or bottom) few messages. The strategic receiver's strategy is non-monotonic, and he may take more conservative action in response to extreme messages.

The above work assumes bounded rational types exist, but is agnostic as to why agents act so. Agents may have preferences for, or are adhering to, social norms and honesty and trust. Alternatively, some agents may lack the strategic sophistication, or desire, to figure out how to play the game optimally, and thus use a heuristic.

2.3 Tests of Bounded Rationality

Other papers in the literature have tested models of bounded rationality in one-shot games. Costa-Gomes et al. (2000) analyze subjects' decisions in a series of two-person normal form games. The payoff matrix is hidden to subjects, but is revealed if they click upon the hidden portion of the matrix with a computer mouse. Each game is a one-shot game where subjects are repaired after every round in the series. Their analysis suggests that strategic behavior can be understood by categorizing subjects' behavior into different types. Of the nine behavioral types, only two are strategic. The authors find boundedly rational strategic types, rather than strategic types, better describe the majority of subjects' behavior.

Stahl and Wilson (1994) characterize heterogeneous behavior over ten one-shot games with a hierarchical model of strategic thinking, where a level-0 type plays unpredictably, a level-1 type acts if others where level-0, a level-2 type acts as if other here all level-0 or level-1 types, and so on. The authors find three fourths of the players were level-1 or level-2, the remaining were Nash types. Camerer et al. (2004) perform a similar analysis of prior published results of coordination and market-entry games. They find, on average, a type of level-1.5 predicts subjects' behavior.

While the above work shows a framework of behavioral types fit the data better than the predictions of traditional analytical modeling, these tests of one-shot game behavior cannot distinguish whether agents are following social norms, acting heuristically, or adapting behavior to the environment at hand. In my work, subjects play a repeated game, giving me the opportunity to measure if subjects adapt, or simply use heuristics repeatedly.

2.4 Tests of Information Transmission

There are experimental studies that have corroborated the prediction that less information about the true state is transmitted as the preferences of the sender and receiver become less aligned. Both Dickhaut et al. (1995) and Blume et al. (1998) find less information is communicated as preferences diverge. However, both papers indicate an overcommunication tendency by subjects, where senders send messages finer than theory predicts, and also find receivers rely on messages more than predicted.

Later work duplicates these dual findings of deception and overcomminication. Cai and Wang (2006) use two measures of bounded rationality to understand the overcommunication finding, behavioral type analysis and quantal response equilibrium. Both measures have some success in interpreting their data where subjects played against each other only once. However, Cai and Wang cannot answer whether certain behavioral types, such as consistently honest senders and trusting receivers, are driven by social norms of trust and honesty, or are simply playing a repeated heuristic.

Sanchez-Pages and Vorsatz (2006) also randomly assigned subjects into pairs for a one-shot game. They add a punishment phase in one of two treatments where the receiver can elect to reduce both subject's payoffs to zero after observing the state and payoff. They find receivers who choose to costly punish are honest when senders. They argue this set of subjects are motivated by socially normative behavior and when this group is excluded the phenomena of overcommunication vanishes. While the authors provide an appealing explanation that trusting receivers punish based on normative behavior, other work shows that subjects simply enjoy punishing others (Fon and Parisi (2005); Fehr and Gaechter (2003); Hopfensitz and Reuben (2005); Fehr and Fischbacher (2004)). In some of the papers, subjects costly punish others even when the punished target's action did not affect the punishing subject's payoff.

Wang et al. (2006) hypothesize that determining how much to deceive another player is cognitively difficult, and measure subjects' pupil dilation, as pupils dilate under stress and cognitive difficulty. They find senders' pupils dilate when they send deceptive messages, and thus argue the data is consistent with the hypothesis that deception is cognitively difficult. While deceptive may be cognitively difficult, it remains an open question as to whether subjects avoid deception, trying to minimize cognitive difficulty, or choose honest behavior for other reasons.

Chapter 3

Model of Forecast and Trade

3.1 Fully-Strategic Analytical Model

The forecasting and trading game is similar, but not identical to, games of strategic information transmission. This difference is driven by a commission based on volume, as opposed to the bias included in quadratic loss function that generally characterize utility functions information transmission games. This gives rise to disjoint regions of the state space where the incentives of the analyst and shareholders are aligned. First I will examine the base game of the finitely repeated game, solving for the forecasting strategy of the analyst assuming all agents are sequentially rational. The base game is a one-shot game, but as I will test this experimentally, I need to consider what happens in the repeated stage game so that the theory matches the manner in which the experiments are ran.

Within the analysis of the base game, I model only the agents' payoffs without any representation of reputation. I show that in equilibrium all forecasts induce trade, or no forecast can induce trade. Thereafter I claim that a model with reputation represented in the analyst's utility function collapses into the base game equilibrium in all stages, and is thus the subgame perfect equilibrium.

3.1.1 The Base Game Equilibrium

Consider a single period game played by three risk-neutral agents, consisting of one analyst and two shareholders, indexed $i = 1,2$.

Shareholders start each stage with an endowment of one share which liquidates at the end of that stage. At the end of the stage, each share entitles its owner to a realization of θ. θ belongs to the finite set $\Theta = \{\theta_1, ..., \theta n\}$, is distributed symmetrically around mean μ, and had probability $g(\theta)$ which is strictly positive for all $\theta \in \Theta$.

At the beginning of the stage, the analyst receives a private signal ς on the liquidating value θ. ς belongs to the finite set $S = \{\varsigma_1, ..., \varsigma_m\}$ where $m \geq n$. The signal is generated through a joint probability $f(\varsigma, \theta)$ satisfying first order stochastic dominance in the sense that $\sum_{k=1}^{j} f(\varsigma_k, \theta) \geq \sum_{k=1}^{j} f(\varsigma_k, \hat{\theta}) \quad \forall j \in \{1, ..., m\}, \hat{\theta} > \theta$. Hence higher signal values ς are indicative of higher values of the liquidation value θ.

After viewing the signal ς, the analyst releases a public forecast of the liquidation value θ. After the analyst releases the forecast, but before the realization of θ is observed, a market opens and the two shareholders can trade.

Date 0	*Date* 1	*Date* 2
ς drawn. Analyst releases forecast.	Market opens for trade.	θ realized. All agents receive payoffs.

Figure 3.1: Base Game Timeline

In the above game, a market exists only if there are potential gains to trading between the owners. To capture gains to trade I assume each shareholder has different incremental tax rates. Specifically, I assume each shareholder incurs a tax expense (benefit) from the realization of θ if that realization is greater (less) than μ. I assume the tax rates are constant for all realizations of the liquidation value even if a single owner owns the entire firm.

Let $\beta_i^t \equiv 1 - r_i^t$, where r_i^t is the incremental tax rate for shareholder i at date t. Then the ex-post utility from the realization of θ is

$$v_i(\theta) = (\theta - \mu)\beta_i + \mu \tag{3.1}$$

Without loss of generality, I normalize $\beta_1^{t=1} = \beta_2^{t=1} = \beta_2^{t=2} = 1$. Let $\beta_1^{t=2} = \beta < 1$. Since $\beta_1^{t=2} \neq \beta_2^{t=2}$ there are potential gains to trade for small or large enough values of θ.

If there is trade, the costs of trading λ are bore by the owner of the firm, and thus, the shareholder owning the firm at the end of the period. I assume that $\lambda = 0$ if the shareholders elect not to trade.

When is θ small enough so that shareholder 1 is willing to buy shareholder 2's share and incur the costs of trading? When

$$v_1(\theta) - \lambda > v_2(\theta) \Rightarrow \beta(\theta - \mu) + \mu - \lambda > \theta \Rightarrow \theta < \mu - \frac{\lambda}{1-\beta} \tag{3.2}$$

Similarly, when θ is large enough so that shareholder 1 is willing to sell her share to shareholder 2? Shareholder 2 would be willing to buy the share and incur trading costs if only if

$$v_2(\theta) - \lambda > v_1(\theta) \Rightarrow \theta - \lambda > \beta(\theta - \mu) + \mu \Rightarrow \theta > \mu + \frac{\lambda}{1-\beta} \tag{3.3}$$

Defining the critical values of θ from equations (3.2) and (3.3) let $\theta_L \equiv \mu - \frac{\lambda}{1-\beta}$ and $\theta_H \equiv \mu + \frac{\lambda}{1-\beta}$. I make the joint assumption that $1 - \beta$ is large enough and λ is small enough such that $\theta_1 < \theta_L < \theta_H < \theta_n$.

If there is trade, the analyst is paid a constant κ, where κ is a component of the trading costs λ. Thus $\kappa \leq \lambda$. If there is no trade the analyst earns nothing. For $\theta > \theta_H$, both shareholders and the analyst benefit if shareholder 2 buys the share from shareholder 1. Likewise, for $\theta < \theta_L$ all agents benefit if shareholder 2 sells to shareholder 1.

3.1.2 Equilibrium Allocations

Before analyzing the analyst's reporting strategies and shareholders' trading behavior, I first determine obtainable allocations using a direct mechanism. After determining feasible allocations, I analyze reporting strategies that yield the same allocations. The agents participate in some sequential reporting-bargaining game, or mechanism, to determine first, what the analyst's report will reveal about the state of nature θ, second, whether there should be a redistribution of shares between the shareholders and, third, how much should the buyer pay to the seller. In order for the shareholders to be willing to participate, the appropriate individually rationality constraint is that the mechanism give non-negative expected gains from trade given the released analyst's forecast.

Invoking the revelation principle, I can without loss of generality restrict my analysis to a direct mechanism that directs trade between the shareholders given the analyst truthfully reports her private information. The incentive compatible direct mechanism yields the same allocations (share holdings, transfer payment, and reporting commissions) as any other mechanism or game (Myerson (1979, 1981)). This mechanism elicits the private information from the analyst and then implements the outcome, in this case redistribution of shares and transfer of money, as in the given game shown in Figure 3.1.

For ease of notation, I will design the mechanism so the analyst reports $\hat{\theta} \equiv E[\theta|\varsigma]$: her expectation of θ given her private signal ς. This is equivalent to asking the analyst to report her signal and then computing the expectation given the signal as the signal distribution is known to all agents. Since $f(\varsigma, \theta)$ satisfies first order stochastic dominance, higher reported $\hat{\theta}$ indicate higher values of ς were realized. The mechanism will consist of:

$T(\hat{\theta})$	The transfer price paid to selling shareholder by buying shareholder satisfying budget balancing
$\phi(\hat{\theta})$	The probability of trade
$\rho(\hat{\theta})$	The probability that shareholder 1 is the seller given trade

After the analyst reports $\hat{\theta}$ at date 0, the mechanism will announce a triplet $\{T(\hat{\theta}), \phi(\hat{\theta}), \rho(\hat{\theta})\}$ at date 1 to the shareholders. The shareholders do not see the analyst's private information, $\hat{\theta}$, but only the function values. Let I_m be the shareholders information at date 1; the realization of mechanism triplet. As in the given game (Figure 3.1), at date 1 the shareholders can commit to the mechanism, or choose to walk away from it.

The mechanism designer considers the shareholders' expected utilities given the analyst input of $\hat{\theta}$. As the utilities are linear in θ, the expected utilities are linear in $\hat{\theta}$

$$v_1(\hat{\theta}) + \phi(\hat{\theta}) \left[(1 - \rho(\hat{\theta}))(v_1(\hat{\theta}) - \lambda) - \rho(\hat{\theta})v_1(\hat{\theta})\right] + (2\rho(\hat{\theta}) - 1)T(\hat{\theta}) \tag{3.4}$$

$$v_2(\hat{\theta}) + \phi(\hat{\theta}) \left[\rho(\hat{\theta})(v_2(\hat{\theta}) - \lambda) - (1 - \rho(\hat{\theta}))v_2(\hat{\theta})\right] + (1 - 2\rho(\hat{\theta}))T(\hat{\theta}) \tag{3.5}$$

The transfer payment in equations (3.4) and (3.5) is not multiplied by the probability of trade. If the transfer price is to be zero in equilibrium, this characterization will emerge endogenously rather than by construction. When combining the above equations, the transfer price drops out entirely; due to balance budgeting it has no effect on total expected gains to trade. Thus, the mechanism designer's programming problem maximizes the combined expected gains to trading, net of transaction costs.

$$\begin{aligned} &\underset{T(\hat{\theta}),\phi(\hat{\theta}),\rho(\hat{\theta})}{Max} && (1-\phi(\hat{\theta}))\left[v_1(\hat{\theta})+v_2(\hat{\theta})\right]+ \\ & && \phi(\hat{\theta})\left[\rho(\hat{\theta})2v_2(\hat{\theta})+(1-\rho(\hat{\theta}))2v_1(\hat{\theta})-\lambda\right] \end{aligned} \tag{3.6}$$

subject to:

$$\text{IC} \quad \phi(\hat{\theta})\kappa \geq \phi(\tilde{\theta})\kappa \qquad \forall \hat{\theta},\tilde{\theta} \tag{3.7}$$

$$\begin{aligned} \text{IR 1} \quad & \phi(\hat{\theta})\left[(1-\rho(\hat{\theta}))(E[v_1(\theta)|I_m-\lambda)-\rho(\hat{\theta})E[v_1(\theta)|I_m]\right]+ \\ & (2\rho(\hat{\theta})-1)T(\hat{\theta}) \geq 0 \end{aligned} \tag{3.8}$$

$$\begin{aligned} \text{IR 2} \quad & \phi(\hat{\theta})\left[\rho(\hat{\theta})(E[v_2(\theta)|I_m]-\lambda)-(1-\rho(\hat{\theta}))E[v_2(\theta)|I_m]\right]+ \\ & (1-2\rho(\hat{\theta}))T(\hat{\theta}) \geq 0 \end{aligned} \tag{3.9}$$

The programming problem (3.6) is the designer's sum of the expected shareholders' utilities (3.4) and (3.5). The incentive compatible (IC) constraint (3.7) pertains to the analyst, who is paid only if there is trade. The individually rationality (IR) constraints (3.8) and (3.9) are the shareholders' expected net utilities if they choose to commit to the mechanism after learning the mechanism's triplet.

First, I solve a relaxed problem omitting the two IR constraints. After finding solutions for $\phi(\hat{\theta})$ and $\rho(\hat{\theta})$, I will use the IR constraints to solve for obtainable values of $T(\hat{\theta})$.

Incentive Compatibility For The Analyst

As the analyst is paid only if there is trade, her constraint is

$$\text{IC} \quad \phi(\hat{\theta})\kappa \geq \phi(\tilde{\theta})\kappa \quad \forall \hat{\theta},\tilde{\theta} \tag{3.10}$$

Proposition 1. *In equilibrium, there will always be trade, or there will never be trade.*

Proof. In order to satisfy IC $\phi(\hat{\theta})$ must be a constant $\forall \hat{\theta}$. Hereafter I will alter the notation to omit any arguments to ϕ. As per (3.6), the programming problem is linear in ϕ, so ϕ will take on a value of either 1 or 0. □

If the designer had knowledge of ς, then $\forall E[\theta|\varsigma] \leq \theta_L$, shareholder 2 would sell to shareholder 1. Similarly, $\forall E[\theta|\varsigma] \geq \theta_H$, shareholder 1 would sell to shareholder 2. In both these cases, the expected difference between the shareholders' utilities is at least λ. For interior values of θ such that $\theta_L < E[\theta|\varsigma] < \theta_H$, there would be no trade.

Corollary 1. *The mechanism is ex-ante inefficient.*

While the programming problem (3.6) is linear in ϕ, it also depends upon $\rho(\hat{\theta})$. Before characterizing the parameters for which there will always be trade or there will never be trade, I need the optimal values for $\rho(\hat{\theta})$.

Proposition 2. *The sum of the expected shareholders' utilities (3.6) is maximized by*

$$\rho(\hat{\theta}) = \begin{cases} 1 & \text{if } \hat{\theta} > \mu \\ \varrho & \text{if } \hat{\theta} = \mu \\ 0 & \text{otherwise} \end{cases} \quad \text{where } 0 \leq \varrho \leq 1 \tag{3.11}$$

Proof. From equation (3.1) the shareholders value realizations of θ differently at all values other then μ. Given trade, the programming problem (3.6) is maximized when $\hat{\theta} < \mu$ if shareholder 1 owns both shares. For values of $\hat{\theta} > \mu$, shareholder 2's expected value of owning both shares exceeds shareholder 1's expected value. □

For all values of θ other than μ, one shareholder values the stock more than the other. If there is going to be trade, the programming problem is maximized when allocating the share to the shareholder who values it more [1]. If $\rho(\hat{\theta}) \notin \{0, 1\}$, the shareholders will correctly infer $\hat{\theta} = \mu$, and will not trade despite the value of ϕ. So, ϱ must be either 0 or 1. So let $\varrho = 1$ with probability of $\frac{1}{2}$ and 0 otherwise.

Given (3.11), ϕ can be determined ex-ante. The mechanism designer determines if gains to unconditional trading exceed the costs of trading. ϕ maximizes (3.6) when

$$\phi = \begin{cases} 1 & \text{if } E\left[v_1(\theta) + v_2(\theta)\right] < \\ & E\left[2v_1(\theta)|\hat{\theta} < \mu\right](G(\mu) - \frac{g(\mu)}{2}) + E\left[2v_2(\theta)|\hat{\theta} \geq \mu\right](1 - G(\mu) + \frac{g(\mu)}{2}) - \lambda \\ 0 & \text{otherwise} \end{cases} \tag{3.12}$$

[1] Consider the case where there is one divisible share in the firm, each shareholder starts the game with an equal proportion of the share, and terminal pay-out is 2*θ. In this divisible share scenario, $\rho(\hat{\theta})$ is interpreted as the proportion of shares that shareholder 1 sells to shareholder 2. Then per Proposition 2 the mechanism would specify that all shares be transferred if there is to be trade. All other results of the model hold in this divisible case of divisible shares.

Let $G(\mu)$ be the probability that analyst will receive a signal ς that induces him to believe θ is μ or less. As per the law of iterated expectations, $G(\mu) = \sum_1^z g(\theta_j)$ where z is the index of the $\theta_z = \sup\{\theta \in \Theta : \theta \leq \mu\}$. Equation (3.12) shows if the gains to unconditional trading exceed the costs of trading, given the mechanism can dictate the direction of trade, then trade should always occur. Otherwise, trade will never occur.

Individual Rationality For The Shareholders

Both shareholders view the mechanism output I_m. Later I address the case where there is no trade, but first assume that $\phi = 1$. As the value of $\rho(\hat{\theta})$ is either 0 or 1, each shareholder is faced with the choice of accepting his role as buyer or seller for the announced transfer price $T(\hat{\theta})$, or keeping his endowed share and not trading. While the shareholders do not view the analyst's private information, they can condition their expectation of θ upon $\rho(\hat{\theta}) \in I_m$. Given Proposition (3.11), I can rewrite equations (3.8) and (3.9) conditional upon the value of $\rho(\hat{\theta})$

(IR) case: $\rho(\hat{\theta}) = 0$

$$T(\hat{\theta}) \geq E\left[v_2(\theta)|\hat{\theta}\rho(\hat{\theta}) = 0\right] \tag{3.13}$$

$$E\left[v_1(\theta)|\rho(\hat{\theta}) = 0\right] - \lambda \geq T(\hat{\theta}) \tag{3.14}$$

(IR) case: $\rho(\hat{\theta}) = 1$

$$T(\hat{\theta}) < E\left[v_2(\theta)|\rho(\hat{\theta}) = 1\right] - \lambda \tag{3.15}$$

$$E\left[v_1(\theta)|\rho(\hat{\theta}) = 1\right] < T(\hat{\theta}) \tag{3.16}$$

Combining (3.13) and (3.14) yields

$$E\left[v_1(\theta)|\rho(\hat{\theta}) = 0\right] - \lambda \geq T(\hat{\theta}) \geq E\left[v_2(\theta)|\rho(\hat{\theta}) = 0\right] \tag{3.17}$$

Analogously, combining (3.15) and (3.16) yields

$$E\left[v_1(\theta)|\rho(\hat{\theta}) = 1\right] \leq T(\hat{\theta}) \leq E\left[v_2(\theta)|\rho(\hat{\theta}) = 1\right] - \lambda \tag{3.18}$$

Proposition 3. *If $E\left[\theta|\hat{\theta} > \mu\right] > \theta_H$ and $E\left[\theta|\hat{\theta} \leq \mu\right] < \theta_L$, IR will be met and the shareholders will participate when directed to trade.*

Proof. Equation (3.17) is true when $E\left[\theta|\rho(\hat{\theta}) = 0\right] < \theta_L$. If so, there exists a transfer price agreeable to both shareholders. If $E\left[\theta|\rho(\hat{\theta})=1\right] > \theta_H$, then (3.18) is true and there again exists a transfer price agreeable to both shareholders. Together (3.17) and (3.18) are sufficient conditions for ϕ to be 1 as per equation (3.12). As ϕ is constant, the analyst's IC requirement is satisfied. □

Since θ is distributed symmetrically around μ, either (3.17) and (3.18) are true, or neither (3.17) or (3.18) is true [2]. Notice that if neither (3.17) or (3.18) are true, then ϕ is 0, which satisfies IC. The mechanism again satisfies IR as there is no transfer price agreeable to both shareholders, so they are unwilling to trade.

Corollary 2. *If neither (3.17) or (3.18) are true, then there will be no trade.*

Characterization Of The Transfer Price

If neither (3.17) nor (3.18) is true, then the mechanism dictates that the shareholders should not trade (ϕ is 0). The IR constraints (3.8) and (3.8) are satisfied only with a transfer price of 0. Trivially, the only method to ensure the shareholders participate in the mechanism when directed not to trade is to set the transfer price be zero. All subsequent discussion in this section assumes that (3.17) and (3.18) are true.

Given IR is satisfied, the shareholder's expected utilities differ by at least λ. The transfer price $T(\hat{\theta})$ must at minimum take on two values, conditional upon $\rho(\hat{\theta})$ so that (3.17) and (3.18) are true. The transfer price in this case can be characterized by

$$\alpha\{E\left[v_1(\theta)|\rho(\hat{\theta})\right] - (1-\rho(\hat{\theta}))\lambda\} + (1-\alpha)\{E\left[v_2(\theta)|\rho(\hat{\theta})\right] - \rho(\hat{\theta})\lambda\} \tag{3.19}$$

where $0 < \alpha < 1$.

In order to satisfy IC and IR, any mechanism that induces trade does so by coarsening the analyst's private information such that values of $\hat{\theta}$ near μ are pooled with lower or higher values so that the shareholders of expectation of θ is less the θ_L or greater than θ_H. The transfer price $T(\hat{\theta})$ characterized by equation (3.19) is the coarsest partition.

[2] If θ is not symmetric around μ, it may be the case the only (3.17) or (3.18) is true and the other false. If both are true, or both are false, then Proposition 3 and the accompanying corollary hold. In this case, it is still possible to design a mechanism that induces trade by selecting a cut-off for $\tilde{c} \neq \mu$ for $\rho(\hat{\theta})$ such that $E\left[2v_1(\theta)|\hat{\theta} \leq \tilde{c}\right] G(\tilde{c}) + E\left[2v_2(\theta)|\hat{\theta} \geq \tilde{c}\right](1-G(\tilde{c})) - \lambda > E\left[v_1(\theta)+v_2(\theta)\right]$, and the rationality constraints, $E\left[v_1(\theta)|\hat{\theta} \leq \tilde{c}\right] - \lambda \geq T(\hat{\theta}) \geq E\left[v_2(\theta)|\hat{\theta} \leq \tilde{c}\right]$, and $E\left[v_1(\theta)|\hat{\theta} \geq \tilde{c}\right] \leq T(\hat{\theta}) \leq E\left[v_2(\theta)|\hat{\theta} \geq \tilde{c}\right] - \lambda$ are true. Even in the case when θ is symmetric around μ, a different cut-off can be supported as an equilibrium.

Alternatively, a finer partition might be possible given equations (3.17) and (3.18) are true. Let $\underline{\theta} = \sup\{\theta \in \Theta : E\left[\theta|\underline{\theta} \leq \hat{\theta} \leq \mu\right] \leq \theta_L\}$. Note that by equation (3.17) $\underline{\theta}$ exists, although it may be that $\underline{\theta} = \theta_1$. Denote the shareholder expectation in this region as $v_* \equiv E\left[\theta|\underline{\theta} \leq \hat{\theta} \leq \mu\right]$. Symmetrically let $\overline{\theta} = \inf\{\theta \in \Theta : E\left[\theta|\mu \leq \hat{\theta} \leq \overline{\theta}\right] \geq \theta_H\}$, and if equation (3.18) is true, $\exists\overline{\theta} \leq \theta_n$. Let v^* denote the shareholder expectation of θ in this region: $v^* \equiv E\left[\theta|\mu < \hat{\theta} \leq \overline{\theta}\right]$. Consider the transfer price

$$T(\hat{\theta}) = \begin{cases} \alpha\{v_1(\hat{\theta}) - \lambda\} + (1-\alpha)\{v_2(\hat{\theta})\} & \text{if } \hat{\theta} < \underline{\theta} \\ \alpha\{v_1(v_*) - \lambda\} + (1-\alpha)\{v_2(v_*)\} & \text{if } \underline{\theta} \leq \hat{\theta} \leq \mu \text{ and } \rho(\hat{\theta}) = 0 \\ \alpha\{v_1(v^*)\} + (1-\alpha)\{v_2(v_*) - \lambda\} & \text{if } \mu \leq \hat{\theta} \leq \overline{\theta} \text{ and } \rho(\hat{\theta}) = 1 \\ \alpha\{v_1(\hat{\theta})\} + (1-\alpha)\{v_2(\hat{\theta}) - \lambda\} & \text{if } \hat{\theta} > \overline{\theta} \end{cases} \tag{3.20}$$

For values of $\hat{\theta}$ in a neighborhood of μ, specifically $\underline{\theta} \leq \hat{\theta} \leq \overline{\theta}$, the transfer price given by (3.20) is equivalent to the price given by (3.19). For values of $\hat{\theta}$ outside this neighborhood, the transfer price is a weighted average of the two shareholders' utilities evaluated at the expected value of θ. Note this transfer price satisfies interim IR and there is always trade between the two shareholders.

3.1.3 Allocations In the Original Game

Given the equilibrium allocations found in the last section, I examine the game shown in Figure 3.1 without a direct-revelation mechanism. The analyst will form a reporting strategy that yield the same allocations as the mechanism. Instead of being directed to trade, the shareholders will jointly decide if there is trade. If there is trade, the shareholders will also determine whom will be the seller and negotiate a transfer price.

Forecast Strategy

Formally, the analyst's reporting strategy, R, maps private signals into a forecast: $R : \varsigma \mapsto f$. Without loss of generality, I assume the forecast space is restricted to the signal space. When the analyst reported to a mechanism, Proposition 1 dictated that there was always trade. This leads to the first result.

Proposition 4. *In equilibrium, all forecasts released must induce trade, or no forecast induces trade.*

Proof. The analysts earn κ if the shareholders trade, else nothing. Posit an equilibrium where there are some forecasts that induce trade, while others messages do not. In period τ, the analyst will never choose to release a forecast that does not not induce trade, as she will always be better off conveying a forecast that does induce trade. Likewise, in period $\tau - 1$, the analyst will again never choose to release a forecast that does not induce trade. The same holds for all earlier periods. □

Akin to the mechanism, let $\hat{\theta} \equiv E[\theta|\varsigma]$, the analyst's expectation of θ given his private information. The analyst's reporting strategy induces trade by pooling $\hat{\theta}$ near μ with lower or higher values so that the shareholders' expectation of θ is less the θ_L or greater than θ_H.

Proposition 5. *In order to induce trade, every forecast f must satisfy*

$$E\left[v_1(\theta)|f\right] - \lambda \geq E\left[v_2(\theta)|f\right], \quad \text{or} \tag{3.21}$$

$$E\left[v_1(\theta)|f\right] \leq E\left[v_2(\theta)|f\right] - \lambda \tag{3.22}$$

Proof. If the forecast f neither satisfies equations (3.21) or (3.22), then the shareholders' expectation of θ is such that $\theta_L < E[\theta|f] < \theta_H$. In this region the difference in the shareholders' expected utilities is less than the costs of trading λ. Thus, the shareholders will not willingly trade. □

As in the mechanism, if equations (3.17) and (3.18) are true, then a reporting strategy that always induces trade is feasible. Consider a simple trigger strategy where $\dot{R} : \varsigma \mapsto \{\dot{f}, \ddot{f}\}$. Let $\dot{\varsigma} = \sup\{\varsigma : \hat{\theta} \leq \mu\}$ be the cutoff signal such that $\forall \varsigma \leq \dot{\varsigma}$, the analyst releases the forecast $\dot{f}$, else $\ddot{f}$ is released. Given this reporting strategy, both shareholders would be willing to trade. When the forecast $\dot{f}$ is released, shareholder 2 would sell, and when $\ddot{f}$ is released shareholder 1 would sell. The transfer price is characterized by equation (3.19).

A reporting strategy need not be as coarse as the aforementioned simple trigger strategy to induce trade. Again, assume equations (3.17) and (3.18) are true, then the following reporting strategy will always induce trade.

$$R^*(\varsigma) = \begin{cases} v_* & \text{if } \underline{\theta} \leq \hat{\theta} \leq \mu \\ v^* & \text{if } \mu < \hat{\theta} \leq \overline{\theta} \\ \hat{\theta} & \text{otherwise} \end{cases} \tag{3.23}$$

By construction, the reporting strategy R^* is the finest partition feasible if the analyst wants shareholder 1 to sell to shareholder 2 when $\hat{\theta} > \mu$, while $\dot{R}$ is the coarsest. In designing the mechanism, the direction of trade was dictated by (3.11). Clearly this was preferred by all agents given equations (3.17) and (3.18) are true. However, if only one of these equations are true, which is only possible when θ is not symmetrically distributed about μ, then trade may still be possible.

Definition 1. *Let R' be a simple trigger strategy, $R' : \varsigma \mapsto \{f,' f''\}$ where there is a cutoff signal, ς', such that $\forall \varsigma \leq \varsigma'$, the analyst releases forecast f', else the analyst releases f''. R' is feasible if it always induces trade. R' induces trade if f' satisfies equation (3.21) and f'' satisfies equation (3.22).*

The reporting strategy $\dot{R}$ is by definition a special case of R'. However, the analyst's expectation of θ, need not be in a neighborhood of μ at the cutoff. The cutoff for R' is not necessarily unique, so $\varsigma' \neq \dot{\varsigma}$ for all feasible R'.

Proposition 6. *If $\dot{R}$ is feasible, then it is preferred by both shareholders to any other feasible reporting strategy R' where $\varsigma' \neq \dot{\varsigma}$.*

Proof. By construction, for values of $\hat{\theta} > \mu$, the analyst releases $\dot{f}$, else $\ddot{f}$ is released. A forecast of $\dot{f}$ prompts shareholder 1 sell to shareholder 2, and $\ddot{f}$ prompts shareholder 2 to sell to shareholder 1. This is equivalent to the mechanism trade parameter $\rho(\hat{\theta})$ in equation (3.11) which maximized the two shareholders' gains to trade. □

Proposition 7. *If R' is not feasible, then no other reporting strategy can induce trade.*

Proof. If R' is not feasible, then the analyst must deviate from a single cutoff strategy. She must use a higher cutoff for f'' so that equation (3.22) evaluated at f'' is true and/or a lower cutoff for f' so that equation (3.21) evaluated at f' is true. This leaves at a subset of ς interior that will map to at least one forecast that conveys $\theta_L < \hat{\theta} < \theta_H$, thus shareholders will not trade. By Proposition 4 this is not a feasible reporting strategy. □

Trade and Transfer Prices

As long as the forecast, given the forecast strategy, satisfies equation (3.21), shareholder 2 will sell to shareholder 1. If the forecast induces an expectation of θ such that equation (3.22) is true, then shareholder 1 will sell to shareholder 2.

Proposition 8. *For any reporting strategy that induces trade, the transfer price $T(f)$ can be characterized by*

$$\alpha\{E\left[v_1(\theta)|f\right] - (1 - I_f)\lambda\} + (1 - \alpha)\{E\left[v_2(\theta)|f\right] - I_r\lambda\} \tag{3.24}$$

where I_f is an indicator variable that is 1 if the forecast conveys $E[\theta] \geq \mu$ and 0 otherwise.

Proof. As both shareholders view the same forecast, their expectation of θ will be homogeneous. Any price negotiated must be interior to the shareholders' expected utilities, which are common knowledge. Since equation (3.24) is simply a weighted average of the expected utilities, it characterizes any negotiated transfer price $T(f)$ that satisfies

$$E[v_1(\theta)|f] - \lambda \geq T(f) \geq E[v_2(\theta)|f], \quad \text{or}$$
$$E[v_1(\theta)|f] \leq T(f) \leq E[v_2(\theta)|f] - \lambda$$

□

Corollary 3. *Any negotiation of the transfer price between the shareholders can be reduced to a negotiation upon α.*

3.1.4 Representing Reputation in the Repeated Game

To recap, I model three agents: an analyst with private signal on the state of nature, and two shareholders play a repeated finite game. For sufficiently low states of nature, one shareholder would be better off selling to the other. For sufficiently high states the opposite is true; one shareholder would be better off buying from the other. For intermediate states there are no gains to trades. If there is trade, the analyst earns a commission.

Within the analyst's utility function, I characterize reputation as an expectation of future periods' profits based on current actions. So, in any stage t other than the terminal stage T, the analyst's utility would include a short-term and long-term component. Let the analyst's utility at stage t be a linear function of short and long-term consequences of her forecast f:

$$U_t(f) + \omega_t V_t(f)$$

Where $U_t(f)$ is the expected commission based on the forecast f, $V_t(f)$ represents the future commissions based on the forecast f in stage t, and ω is the weight placed on future as of stage t.

I argue ω should be a decreasing function of the stage. Furthermore, I argue that $\omega_{t=T}$ is zero. As such, the game in the last stage is identical to the game modeled in Section 3.1.1. Sequentially rational shareholders would anticipate this in the last period, and trade accordingly. Anticipating this, the analyst now plays stages $t \in \{1, 2, ..., T-2\}$ with short-term and long-term components in her utility function, and simply chooses her forecast f to maximize commissions in the last two stages. However, shareholders anticipate this, and soon the game collapses into T identical subgames where the analyst maximizes commissions and ignores reputation in all stages.

Chapter 4

The Experiment

4.1 Experimental Design

The experimental design is based upon information transmission games with multiple receivers. One player with private information sends a message to two other players, who jointly take an action that affect both their own and the message sender's payoff. This base game is repeated eight times.

4.1.1 Base Game

This is a three-person sequential move game that consists of three stages. At the beginning of the game an analyst is endowed with a private signal and two shareholders are endowed with one share that liquidates in the last stage. This liquidation value is a function of nature, which is an element of the set $\{A, B, C, D\}$. The private signal, $s \in \{A, B, C, D\}$, is uniformly distributed. The signal is informative of the state of nature as per Table 4.1. The signal structure and all distributions are public knowledge, while the realization of the analyst's signal is private.

Table 4.1: Chances of State Given Private Signal

Signal Realized	A	B	C	D
Probability State is A	.9	.1		
Probability State is B	.1	.9		
Probability State is C			.9	.1
Probability State is D			.1	.9

Stage 0	*Stage* 1	*Stage* 2
Analyst releases forecast.	Market opens for trade.	State realized. All subjects receive payoffs.

Figure 4.1: Base Game Stages

In the first stage, the analyst releases a forecast, $f \in \{A, B, C, D\}$. If the realization of the private signal s $\in \{A, B\}$ the analyst can release of forecast of either A or B. Likewise if $s \in \{C, D\}$, the analyst can release a forecast of either C or D. This forecasting technology is common knowledge, so after the forecast is released shareholders know the state is either in $\{A, B\}$ or in $\{C, D\}$.

In the second stage, a market opens and the shareholders can trade. If there is trade, a shareholder buys the other's share at a transfer price P and pays trading costs λ including a \$13 commission paid the to analyst. If there is no trade, each shareholder retains his endowment and the analyst earns nothing.

In the last stage, the state of nature is drawn, conditional upon the private signal, and shares are liquidated. The shareholders have different payoffs from the liquidated shares. For each share held by the shareholder, his payoff is shown in Table 4.2.

Table 4.2: Payoffs Given State

State	A	B	C	D
Payoff to Shareholder 1	30	50	70	90
Payoff to Shareholder 2	0	40	80	120
Difference	30	10	10	30

Trading costs vary across treatments. In the low-cost treatment trading costs, λ, consist only of the \$13 commission. In the high cost treatment trading costs are \$21 and include the \$13 commission. Imagine the analyst honestly revealed her private signal via the forecast. Then risk-neutral shareholders would have expectation of payoffs shown in Table 4.3 after the release of the honest analyst's forecast. Note the trading costs are high enough such that if the shareholders believe the state were likely B or C, they would not rationally trade. This is illustrated in Figure A. Thus the analyst has incentive to convince the shareholders that the state is likely A or D.

Table 4.3: Honest Forecasting Strategy

Forecast f	Shareholder 1's Expected Payoff	Shareholder 1's Expected Payoff	Absolute Difference
A	32	4	28
B	48	36	12
C	72	84	12
D	88	116	28

Rational shareholders would anticipate this deception, viewing each forecast as an indication that the state is either $\in \{A, B\}$ or $\in \{C, D\}$. The shareholder would then have the expectation of payoffs shown in Table 4.4 after the release of the forecast. In the low-cost treatment shareholders could trade, as the difference in expected payoffs exceed the trading costs. This relationship is illustrated in Figure A for the low-cost treatment. However, in the high-cost treatment, the costs of trading exceed the difference in expected payoffs, and the shareholders would not trade.

Table 4.4: Deceptive Forecasting Strategy

Forecast f	Shareholder 1's Expected Payoff	Shareholder 1's Expected Payoff	Absolute Difference
$\dot{f}$	40	20	20
$\ddot{f}$	80	100	20

This suggests that if agents are sequentially rational, every forecast would induce trade in low-cost treatment and no forecast would induce trade in the high-cost treatment. These arguments are presented formally in Section 3.1.

4.1.2 Subjects' Decisions

Within the experiment, an economy is defined as a grouping of an analyst and two shareholders. Each round the economy plays the base game. The economy remains intact for a set of eight rounds. Each economy was independent from other in the laboratory. Each participant in the economy only witnessed results for their own economy.

Subjects playing the role of analysts entered their forecasting strategy before seeing a realization of the private signal. That is, for all possible values of the private signal, the analyst made a binary choice: to forecast honestly or attempt to deceive the shareholders. See Figure B.

Meanwhile subjects playing the role of shareholders entered their trading decision. Four trading prices were presented as per Table 4.5. If the shareholders believe the analyst is honest, then P1 and P4 are the feasible prices when the forecast is A and D, respectively. If the shareholders believe the analyst is deceptive, then P2 and P3 are the feasible prices when the forecast is $\{A, B\}$ or $\{C, D\}$, respectively, in the low-cost treatment. In the high-cost treatment, no price is feasible when the analyst is deceptive.

Table 4.5: Prices for Both Treatments

$P1$	7
$P2$	23
$P3$	83
$P4$	91

For each possible forecast value two trading prices where made available to shareholders. For forecast values of A and B, prices P1 and P2 where presented, and for forecast values of C and D prices P3 and P4 where presented. In addition, shareholder 1 was restricted to buying or not trading if the possible state was A or B, and selling or not buying if the possible state was C or D. Shareholder 2 was symmetrically restricted.

Like the analysts, a shareholder entered his entire trading strategy before seeing the analyst's forecast. That is, for all possible values the analyst's forecast might take, the shareholder elected to trade at the high price, the low price, or elected not to trade. See Figures B and B.

4.1.3 Trading and Reporting History

The computer drew a realization of the state and private signal using the aforementioned joint distributions. Using the forecast strategy input by the analyst, the computer determined the analyst's forecast for the given realization of the private signal. Using the shareholders' input trading strategies, the computer determined if there was trade for the released forecast. If the ask exceeds the bid, or either shareholder elected not to trade, the shareholders kept their endowed share and the analyst earned nothing. If the shareholders agreed to trade at both the high and low price (high bid and low ask), then the lower price was use. Using the state and agreed upon price, the computer determined the shareholders' and analyst's payoffs. The computer drew ten such realizations and reported the results in a summary table (see Figure B). This summary table showed the realizations for all rounds played within the economy. Subjects were able to scroll down and see realizations in earlier rounds.

Using the summary report for the ten realizations, shareholders could partially infer the forecasting strategy of the analyst in the prior round. The shareholders' trading strategies were explicitly displayed for all forecasts released, but the analyst's strategy could be inferred from the numbers of mismatches between the forecast and the realized state. Knowing the analyst's signal structure (shown in Table 4.1), shareholders can, in theory, calculate the likelihood of the analyst's forecast strategy given the probabilities shown in Table 4.6.

Table 4.6: Probability of Forecast and State Mismatch Given Forecast Strategy

Number of Forecasts	Probability Given	
not Equal to State	Honest Strategy	Deceptive Strategy
0	89%	1%
1	10%	2%
2	1%	4%
3	<1%	9%
4	< .1%	17%
5	< .1%	34%
6	< .1%	17%
7	< .1%	9%
8	< .1%	4%
9	< .1%	2%
10	< .1%	1%

4.2 Experimental Procedures

The experiments were conducted in Montreal, Canada by CIRANO over two sessions during April and May 2007. Participants were recruited by CIRANO from a standard subject pool and remain anonymous to the author. Participants interacted with each other anonymously over a local computer network. The experiment was programmed and conducted using z-Tree software that was specifically designed for economic experiments (Fischbacher (2007)). The computers were placed in such a way that all participants could only view their own computer screen.

The treatments lasted approximately two hours, and were sequenced as follows.

1. An experimenter read the instructions aloud while each participant followed along with their own copy of the instructions (see Appendix D). The instructions explain the experimental procedures and the information structures used in the experiment. While going over the instructions, participants were asked to write down their answers to several questions to ensure that they understand the instructions. Participants' answers remained confidential. The experimenter reviewed the correct answers. During and after the instructions were read, participants were prompted to ask the experimenter any questions regarding the experiment procedures.

2. Each participant was randomly and anonymously grouped with two other participants. The identities of group members were not revealed to any participant. An analyst was randomly selected within each group and others were assigned the role of shareholders, each owning a stock that paid a dividend.

3. The analyst input her forecasting strategy for each possible value of the private signal of the stock dividend. Meanwhile, the shareholders entered their trading strategies for each possible value of the analyst's forecast. Each participant had one-minute to enter his or her decision. If the analyst failed to enter a decision, the computer used her last input decision, or, if in the first round of set, randomly determined the forecasting strategy. If the analyst failed to make a decision no commission was paid regardless of trade. If the shareholder failed to enter a decision, the computer assigned a strategy of no trade and the shareholder kept his endowment of stock.

4. Using the input strategies, the computer drew ten realizations of the analysts private signal and nature. A summary screen showed each participant the results, including their own payoff, of these ten realizations.

5. The preceding two steps constitute a round. Each grouping of particpants played eight rounds.

6. Participants will be regrouped, as described above, and played another. Another set of eight rounds. The participants played eight sets in total.

7. Each participant was paid a $10 participation fee and the payoffs of ten randomly drawn realizations over the sixty-four rounds.

8. Each participant signed and dated a payment receipt form and received payment.

Chapter 5

Hypotheses Development

5.1 Alternative Theories of Behavior

Prior research clearly shows that senders will attempt to deceive receivers in the standard information transmission game (Dickhaut et al. (1995); Blume et al. (1998); Cai and Wang (2006); Sanchez-Pages and Vorsatz (2006); Wang et al. (2006)). At the same time, these papers clearly show that there is over-communication compared to the fully strategic analysis. In this setting, this would dictate that some analysts forecast honestly, fully revealing their private signal, and some shareholders trust the analysts' reports. I look to two potentially competing notions of behavior that might drive off-equilibrium behavior: boundedly rational behavior and normative social behavior.

5.1.1 Boundedly Rational Behavior

Behavioral type analysis is an approach used by several authors (Costa-Gomes et al. (2000); Stahl and Wilson (1994); Camerer et al. (2004)) to predict how subjects will behave. This approach posits that subjects may differ in the extent they analyze the game and other subjects' incentives, leading to heterogeneous levels of sophistication. These and other papers claim that subjects in experiments behave in a specific boundedly rational manner, using the same strategy, or heuristic, again and again. Furthermore, the analysis assumes that subjects of differing levels of sophistication have non-equilibrium beliefs of other subjects' level of sophistication.

It is not straightforward to apply the cognitive hierarchy of Camerer et al. (2004) to communication games. To apply the behavioral type approach, Crawford (2003) cites early experiment evidence (Blume et al. (1998)) and argues the system of types should anchored on the honest type of sender (analyst) and the trusting type of receiver (shareholder). In this work, I also anchor my schema upon norms of honesty and trust. Agents of certain level of sophistication adapt behavior assuming all other players have lesser levels of sophistication.

Applying this schema to agents playing the base game in section 4.1 yields the following strategies and beliefs:

Level-0 An agent of this level of sophistication forecasts honestly and trusts that others forecast honestly when trading.

Level-1 An agent of this level of sophistication believes that all shareholders trust the forecast released, and thus releases a deceptive forecast. However, this agent also believes that all analysts are forecasting honestly, so trusts the forecasts when trading.

Level-2 An agent of this level of sophistication believes that all shareholders trust the forecast released, so then releases a deceptive forecast. This agent believes that all analysts are deceptive, become is skeptical of forecasts when trading.

Level-3 The level of sophistication is identical to the sequentially rational beliefs and actions as per Section 3.1. Note the subject's observed actions at this level are identical to actions of the prior level, but the agentst's beliefs differ.

Entertaining that agents may be boundedly rational introduces a richer set of predictions than the sequentially rational model. Table 5.1 reports the predicted results if overall subjects' behavior can be explained by one of the above levels of strategic sophistication. While one combination of types, the honest analyst and strategic (skeptical) shareholders, does not appear in the schema above, it is included to provide a complete set of combinations. The honest analyst and trusting shareholder, play a heuristic repeatedly, while a strategic player correctly infers the others' types. For instance, a strategic shareholder determines whether the analyst is honest or not, and trades accordingly.

	Likelihood of Trade		Likelihood of Capturing Gain		Consistency in Forecasting	Allocation of Gains to Trade
	Low-cost Treatment	High-cost Treatment	Low-cost Treatment	High-cost Treatment		
Honest Analyst and Trusting Shareholders (Level-0)	50%	50%	90%	90%	Maximum	Equal Between Shareholders
Honest Analyst and Strategic Shareholders	50%	50%	90%	90%	Maximum	Equal Between Shareholders
Strategic Analyst and Trusting Shareholders (Level-1)	Dependent Upon Forecasting Consistency				Varies	Greater for Shareholder 1
Strategic Analyst and Strategic Shareholders (Level-2, Level-3)	100%	0	100%	0	Varies	Equal Between Shareholders

Table 5.1: Predictions for Levels of Strategic Sophistication

The predictions are based on the experimental parameters discussed in Section 4.1. The calculations for the predictions in Table 5.1 are discussed below.

Honest Analyst and Trusting Shareholders

Within the cognitive hierarchy framework, the combination of honest forecasting and trusting shareholders is the base level of sophistication. When the analyst honestly reveals her signal, the difference in shareholders' payoffs is greater than trading costs only when the signal is A or D (see Table 4.3). This is true in both the high-cost and low-cost treatments. When the analyst is honest and shareholders are trusting, there is trade at prices P1 and P4 when the forecast is A or D, respectively. Given the analyst's signal structure, there is trade half the time.

Gains to trade are captured if the shareholders trade in states A and D. The likelihood that gains to trade are captured is less than unity due to the structure in the analyst's signal shown in Table 4.1. For both signals A and D, there is a 10% chance that the state is B and C, respectively.

Within the framework of this analysis, rooted in differing levels of sophistication, I argue that an honest analyst uses a reporting strategy that simply reveals her signal, as opposed to a reporting strategy that is invertible and thus fully-revealing. Given this argument, the forecast mapping is consistent from round to round.

Each shareholder captures a portion to the gains of trade when trading. The combined gains from trading when the signals are A and D are split equally and shown in Table 5.2.

Table 5.2: Expected Gains to Trading Given Honest Reporting Strategy

	Forecast A	Forecast D	Average
Low-cost Treatment			
Shareholder 1	12	3	7.5
Shareholder 2	3	12	7.5
High-cost Treatment			
Shareholder 1	4	3	3.5
Shareholder 2	3	4	3.5

Honest Analyst and Strategic Shareholders

The aforementioned allocations are also obtained when the analyst is honest and strategic shareholders correctly determine the analyst's behavior. The combination is not included in the cognitive hierarchy above, but is included for comparison and is later relevant when normative behavior is considered. Note that if the analyst deviated, for instance in the last round, to a deceptive strategy, then the consistency would decrease, but the likelihood of trade would increase.

Strategic Analysts and Trusting Shareholders

A combination of a strategic analyst and trusting shareholders is equivalent to level-1 of the cognitive hierarchy. When the analyst is strategic, she uses one of the forecast strategies described in section 3.1.3. The analyst might always release a forecast of A or D, but may also elect to use a random reporting strategy, randomly forecasting either A or B when receiving both signal A and B, and randomly forecasting either C or D when receiving both signal C and D. However, since shareholders are trusting–taking the forecast at face value, they only trade when the forecast is A or D.

This combination is uniquely different from all other combinations. First, the shareholders do not equally split gains to trade. If the analyst forecasts A when her signal is B, then shareholder 2 sells for too little benefiting shareholder 1. If the analyst forecasts D when her signal is C, then shareholder 2 buys for too much, again benefiting shareholder 1. Second, if the analyst forecasts B or C, then there is no trade. Thus, employing a mixed reporting strategy is detrimental to the analyst when facing naïve shareholders. Therefore, consistency in the mapping of signals to forecasts, as well as the forecasting strategy, dictates the likelihood of trade and the likelihood of capturing gains to trade.

Strategic Analysts and Strategic Shareholders

The last combination is equivalent to the predictions of the fully-ration model in Section 3.1. The allocations are equivalent to both level-2 and level-3 of the hierarchy described above. There is always trade, or there is never trade, dependent upon trading costs. Since the shareholders are skeptical of the forecast, the consistency in the forecast mapping is irrelevant.

For a forecast of A or B, the shareholders will trade at P2 in the low-cost treatment, and not trade at all in the high-cost treatment. For a forecast of C or D, the shareholders will trade at

P3 in the low-cost treatment, and not trade at all in the high-cost treatment. The prices are such that the overall expected gains are split equally and shown in Table 5.1.1

Table 5.3: Expected Gains to Trading Given Deceptive Reporting Strategy

	Forecast A or B	Forecast C or D	Average
Low-cost Treatment			
Shareholder 1	4	3	3.5
Shareholder 2	3	4	3.5
High-cost Treatment			
Shareholder 1	0	0	0
Shareholder 2	0	0	0

5.1.2 Socially Normative Behavior

Evolutionary biologists and psychologists have devised explanations for the levels of altruism and reciprocity found among some creatures. For example, natural selection can select for genes that encourage relatives to help one another. If I sacrifice a bit of food to relatives or defend them from attack, I am helping some of my genes survive. These biologists and psychologists argue social norms of trust and honesty, paired with the ability to detect deception, have perpetuated as these norms have enabled coordination and economic prosperity (Dawkins (2006); Cosmides and Tooby (2005); Axelrod (1981)).

These norms suggest that some shareholders might elect to trust the analyst initially, and, assuming the analyst is indeed perceived as honest, hold this belief absent evidence to the contrary. However, given evidence that suggests the analyst is deceptive, the shareholders would revise their trading strategies. If social norms of honesty and trust drive off equilibrium behavior, then the allocations would be equivalent to the second and last combinations of types shown in Table 5.1. Either the analyst is honest and strategic shareholders sensibly trust her forecasts, or the analyst is deceptive and shareholders are skeptical of her forecasts. Unlike the boundedly rational framework, there is no prediction of shareholders naïvely trusting in a repeated game. Shareholders would detect deceptive behavior and alter their trading behavior.

During the experiment a rich history set is provided to subjects so that a strategic subject might be able to reasonably infer others' strategies. By providing history in a repeated game, I will be able to assess whether subjects are indeed boundedly rational. In particular, do trusting types not alter their beliefs in light of conflicting information, or do subjects adapt to other players' strategies?

Instead of acting on social norms of trust and honesty, do subjects lack the strategic sophistication to figure out how to play the game optimally, and thus use the same trading heuristic repeatedly? If so, I should see shareholders using the same strategies independent of the analyst's strategy. Prior work has shown that subjects tend to use simple heuristic behavior as the game becomes more complex (Costa-Gomes et al. (2000)).

5.2 Hypotheses

Three competing hypotheses are derived based on the overall predicted strategies, likelihood of trade, and the allocation of gains to trading presented in Table 5.1.

If most subjects are honest and trusting, whether driven by social norms or lacking in sophistication, then the allocations will be as shown in the first (or second) row of Table 5.1. The analysts will reveal their private information and shareholders will trade when the extreme states are likely.

Hypothesis 1. *If, overall, subjects form strategies based on honesty and trust, then the likelihood of trade is equal over the two treatments, and the likelihood of capturing gains is high and equal over treatments. Analysts' forecasts will be consistent, and the benefit of trading will be equal and positive for both shareholders.*

If most subjects are boundedly rational, to a level less than required to reach the sequentially-rational equilibrium, but greater than the base level, then allocations will be similar to those shown in the third row of Table 5.1.

Hypothesis 2. *If overall, subjects are boundedly rational, then most subjects will chose deceptive forecasting strategies and trusting trading strategies. The likelihood of trade and capturing gains will be correlated with the consistency of forecasting. Shareholder 1 will benefit from trading at shareholder 2's expense.*

This case is incompatible with behavior driven by social norms of honesty and trust, as shareholders do not detect deception in the analyst's forecasts.

Last, if most subjects are fully rational, then allocations will be similar to the last row of Table 5.1. Overall shareholders will anticipate deceptive analyst forecasts and the likelihood of trade will depend upon the cost of trading.

Hypothesis 3. *If overall, subjects are fully rational, then the likelihood of trade and the likelihood of capturing gains to trade will be greater in the low-cost treatment than in the high-cost treatment. The benefit of trading will be equal and positive for both shareholders in the low-cost treatment, and non-negative for shareholders in the high-cost treatment.*

Chapter 6

Findings and Conclusion

6.1 Behavioral Analysis and Descriptive Statistics

6.1.1 Description of the Data Sets

Two data sets are analyzed in this section. Both sets come from laboratory experiments conducted in Montreal, Canada by the Center for Interuniversity Research and Analysis of Organizations (CIRANO). The first set is from the low-cost treatment conducted April 7, 2007, and the second from the high-cost treatment conducted May 2, 2007. Experiment parameters are described in Section 4.1.

Twenty-four subjects where randomly assigned to the two treatments from a subject pool consisting primarily of university students. 15 of the 24 subjects were female. The subjects' median age was 25 years, the youngest 20, and eldest 34.

Each treatment was conducted with 12 subjects and total of 8 sets. Each set consisted of 8 rounds. At the start of each set, subjects were randomly grouped into economies consisting of an analyst and two shareholders. Within the economy subjects were randomly assigned roles of analyst or shareholder. As such, there were 32 independent economies within each treatment. Since roles were randomly assigned, not every subject played each role an equal number of times. However, every subject played the role of analyst at least once and the role of shareholder at least twice.

Table 6.1: Subject Demographics

Session	Gender	Birth Date	Domain of study
Apr 7 1 p.m.	male	1984	Anthropology
Apr 7 1 p.m.	female	1979	Other
Apr 7 1 p.m.	female	1985	Biology
Apr 7 1 p.m.	male	1980	Cinema
Apr 7 1 p.m.	female	1984	Accounting
Apr 7 1 p.m.	female	1983	Law
Apr 7 1 p.m.	female	1979	Cultural Studies
Apr 7 1 p.m.	female	1979	Linguistic
Apr 7 1 p.m.	female	1986	Music
Apr 7 1 p.m.	male	1987	Music
Apr 7 1 p.m.	female	Not Reported	Psychology
Apr 7 1 p.m.	female	1981	Health
May 2 2 p.m.	male	1984	Business Administration
May 2 2 p.m.	female	1985	Biology
May 2 2 p.m.	male	1982	Engineering
May 2 2 p.m.	male	1984	Engineering
May 2 2 p.m.	female	1986	Management
May 2 2 p.m.	male	1980	Computer Science
May 2 2 p.m.	female	1983	Computer Science
May 2 2 p.m.	male	1973	Marketing
May 2 2 p.m.	female	1982	Math and Statistics
May 2 2 p.m.	female	1987	(Unemployed)
May 2 2 p.m.	male	1978	Economics
May 2 2 p.m.	female	1983	Political Science

6.1.2 Strategic Behavior Classification

Subjects were classified into one of three types based upon their input strategy for every round. Based upon the eight round-based classifications, subjects were classified into three behavioral types for the every set. The set-types are based on the cognitive hierarchy described in Section 5.1.

Analyst Round Behaviors

Honest The subject revealed her private signal in the forecast, forecasting A when the signal was A, forecasting B when the signal was B, and so forth.

Deceptive The subject chooses one forecast value for both signals A and B, and another forecast value for both C and D.

Partially Deceptive The subject revealed her private signal when the signal was A or B (C or D), but chooses one forecast value for both signals C and D (A or$B$).

Shareholder Round Behaviors

Trusting The subject had different trading decisions (bids or asks) for forecasts of A versus B and/or different trading decisions for forecasts of C versus D. [1]

Skeptic The subject had one trading decision for both forecasts of A and B, and another trading decision for forecasts of C and D.

Non-trading The subject elected not to trade for every forecast value. This was also default action if the shareholder failed to input a decision in the allotted time.

The frequencies of the subjects' behavior in each round are shown in Tables 6.2 and 6.3. In both treatments analysts tend to use deceptive or partially deceptive forecasting behavior. In both treatments, shareholders tend to use trusting behaviors.

Table 6.2: Frequency of Round Behaviors: Low-Cost Treatment

Analyst	Honest	Partially Deceptive	Deceptive
	113	39	104
Shareholder 1	Trusting	Non-cooperative	Skeptic
	128	47	81
Shareholder 2	Trusting	Non-cooperative	Skeptic
	176	26	54

Table 6.3: Frequency of Round Behaviors: High-Cost Treatment

Analyst	Honest	Partially Deceptive	Deceptive
	88	27	141
Shareholder 1	Trusting	Non-trade	Skeptic
	139	20	97
Shareholder 2	Trusting	Non-trade	Skeptic
	161	28	67

[1]This classification does not require trading decisions monotonic in the forecast. Only one observation of the 604 classified as trusting had a higher price for B than A and a higher price for C than D. Less than 10% had a higher price for B than A or a higher price for C than D. Accordingly, this coding seems to capture subjects whom believed the forecast, rather than traded erratically.

Analyst Set Behaviors

Given the eight round strategies, a subject's forecasting set behavior is classified using two different methods First, the set behavior is determined using a simple majority replicating the methods of Cai and Wang (2006); Costa-Gomes et al. (2000). Borrowing from this work, the majority is defined as 5 out of 8 rounds. Second, set behavior is classified using a deceptive score. The score is assigned in each of the eight rounds as per Table 6.4. The average of the eight rounds' scores is used to determine the set behavior. Figures C.1 and C.2 show the distribution of the average deceptive scores.

Table 6.4: Round Deceptive Score

Round Strategy	Honest	Partially-Deceptive	Deceptive
Deceptive Score	0	0.5	1

Following are the analyst behavior types based on the two classification methods. The portion in brackets describes the numeric cutoffs for the average deceptive score.

Honest The subject used an honest round strategy a majority of the time [Deceptive score less than or equal to 0.25].

Deceptive The subject used a deceptive and/or partially deceptive round strategies a majority of the time, or partially deceptive and honest strategies an equal number of times [Deceptive score greater then 0.5].

Mixed The subject used honest and deceptive strategies an equal number of times. In this case, the forecast may or may not have any information content [Deceptive score greater than 0.25 and less than or equal to 0.5]

The analyst set behavior frequencies for each subject are shown in Tables 6.5 and 6.6 using both classification methods. As roles of analyst and shareholders where randomly assigned, not all subjects played the role of analyst an equal number of times. A majority of subjects' set behavior (70%) was classified as deceptive or mixed. Most subjects had behavior that was classified into one type a majority of the time (shown in bold in Tables 6.5 and 6.6).

Both classification methods yielded 5 subjects who where honest a majority of the sets [2]. Only 1 subject used an honest forecasting strategy in every round of every set.

[2]Changing the majority criteria from to 5 of 8 rounds to 6 of 8 rounds yielded only 4 subjects who where classified as honest.

	Low-Cost Treatment				High-Cost Treatment		
Subject	**Honest**	**Deceptive**	**Mixed**	**Subject**	**Honest**	**Deceptive**	**Mixed**
1		**2**	1	13	1	1	
2	1	1		14		**3**	
3		**2**		15		1	1
4		**3**		16		**3**	
5		**1**		17	**2**		
6		**1**		18		**3**	
7	1	1		19		**2**	
8		**2**		20		4	
9	**4**	1		21	**1**		
10	1	1		22		**4**	
11	1		**2**	23	**3**		1
12	**4**	2		24			**2**
Totals	12	17	3	Totals	7	21	4
Frequency	37.5%	53.1%	9.4%	Frequency	21.9%	65.6%	12.5%

Table 6.5: Subject Forecasting Behavior in Every Set (Using Majority Round Behavior as Criteria)
Prevalent strategy shown in bold

	Low-Cost Treatment				High-Cost Treatment		
Subject	**Honest**	**Deceptive**	**Mixed**	**Subject**	**Honest**	**Deceptive**	**Mixed**
1		**2**	1	13	1	1	
2	1	1		14		**3**	
3		**2**		15	1	1	
4		**3**		16		**3**	
5		**1**		17	**2**		
6		**1**		18		**3**	
7	1	1		19		**2**	
8		**2**		20		**4**	
9	**5**			21	**1**		
10	1	1		22		**4**	
11	1		**2**	23	**4**		
12	**4**		2	24			**2**
Totals	13	14	5	Totals	9	21	2
Frequency	40.6%	43.8%	15.6%	Frequency	28.1%	65.6%	6.3%

Table 6.6: Subject Forecasting Behavior in Every Set (Using Average Deceptive Score as Criteria)
Prevalent strategy shown in bold

Shareholder Set Behaviors

Similar to the classification of forecasting behavior, each subject's trading set behavior was classified using two different methods. First, classification was based on a simply majority of behavior based on 5 out of 8 rounds. If there was no majority, the mode round behavior was used as the basis of classification. Second, set behavior is classified using a trust score. The score is assigned in each of the eight rounds as per Table 6.7. The average of the eight rounds' scores is used to determine the set behavior. Figures C.3 and C.4 show the distribution of the average deceptive scores.

Table 6.7: Round Trust Score

Round Strategy	Skeptical	Non-trading	Trusting
Trust Score	0	0.5	1

Following are the set behaviors based on the two classification methods. The portion in brackets describes the numeric cutoffs for the average trust score. As non-trading behavior maybe a function of skepticism, or maybe a signal to the other shareholder to alter his bid, the result of procedural time-out, all numeric cutoffs where formed interpreting non-trading as neutral in respect to trust.

Trusting The subject used a trusting round strategy a majority of the time [Trust score greater than or equal to 0.675].

Skeptic The subject used a skeptic and non-trading round strategy a majority of the time [Trust score less than or equal to 0.375].

Mixed The subject used trusting and skeptic strategies an equal number of times [Trust score greater than 0.375 and less than 0.675].

The shareholder set behavior frequencies for each subject are shown in Tables 6.8 and 6.9. These frequencies include both the behavior when the subject played shareholder 1 and shareholder 2, thus each table has two observations for each economy. A majority of the subjects' set behavior (69%) was classified as trusting. Most subjects had trading behavior that was classified into one type a majority of the time (shown in bold in Tables 6.8 and 6.9). Both classifications approaches yielded 16 or more subjects who where trusting in a majority of the sets [3]. Only two subjects used

[3] Changing the criteria to require a majority, rather than using mode in absence of majority, yielded 15 subjects classified as trusting

trusting strategies in all rounds of all sets.

	Low-Cost Treatment				High-Cost Treatment		
Subject	**Trusting**	**Skeptic**	**Mixed**	**Subject**	**Trusting**	**Skeptic**	**Mixed**
1	**3**	2		13	**4**	2	
2	3	3		14	**5**		
3	**5**		1	15	**5**		1
4	**4**	1		16	**3**	2	
5	**7**			17	**6**		
6	**5**	2		18	1	2	2
7	**6**			19	**6**		
8		**6**		20	**3**		1
9	**3**			21	3	3	1
10	**6**			22	2	2	
11	2	**3**		23	**4**		
12	**2**			24		**5**	1
Totals	46	17	1	Totals	42	16	6
Frequency	71.9%	26.6%	1.6%	Frequency	65.6%	25.0%	9.4%

Table 6.8: Subject Trading Behavior in Every Set (Using Majority Round Behavior as Criteria)
Prevalent strategy shown in bold

	Low-Cost Treatment				High-Cost Treatment		
Subject	**Trusting**	**Skeptic**	**Mixed**	**Subject**	**Trusting**	**Skeptic**	**Mixed**
1	**3**	1	1	13	4	2	
2	3	3		14	**5**		
3	**5**		1	15	**5**		1
4	**5**			16	**3**	2	
5	**7**			17	**6**		
6	**5**	2		18		1	**4**
7	**6**			19	**6**		
8		**5**	1	20	2		2
9	**3**			21	2	**3**	2
10	**6**			22	2	2	
11	1	3	1	23	4		
12	**2**			24		**5**	1
Totals	46	14	4	Totals	39	15	10
Frequency	71.9%	21.9%	6.3%	Frequency	60.9%	23.4%	15.6%

Table 6.9: Subject Trading Behavior in Every Set (Using Average Trust Score as Criteria)
Prevalent strategy shown in bold

6.1.3 Subject Overall Behavior

Table 6.10 reports each subject's set behaviors were classified using the majority of round behavior described in Section 6.1.2 and Section 6.1.2. For every row within the table, each subject's behavior in sets 1 through 8 is classified as either honest, deceptive or mixed when playing the role of analyst, and trusting, skeptical, or mixed when playing the role of shareholder. For example, subject #1 played the role of shareholder 1 in sets 1 and 2, submitting bids consistent with skeptical behavior. In sets 3 and 4 subject #1 played the role of analyst and used a combination of honest and deceptive forecasting strategies, but in set 4 used primary deceptive forecasting strategies. Thereafter, subject #1 played the role of shareholder 2 in sets 5 and 6, and submitted bids consistent with trusting behavior. In set 7, subject #1 again used primary deceptive forecasting strategies in the role of analyst. In the final set, subject #1 again submitted bids consistent with trusting behavior in the role of shareholder 1.

	Reporter			Shareholder 1			Shareholder 2		
Subject	**Honest**	**Deceptive**	**Mixed**	**Trusting**	**Skeptic**	**Mixed**	**Trusting**	**Skeptic**	**Mixed**
1		4,7	3	8	1,2		5,6		
2	3	6			4,5,7		1,2,8		
3		6,8		5		7	1,2,3,4		
4		1,2,5		3,4,6			7	8	
5		6		1,4,5,7,8			2,3		
6		4		1,2,3,5	6,8		7		
7	5	1		2,3			4,6,7,8		
8		1,3			2,7			4,5,6,8	
9	2,3,7,8	1		4			5,6		
10	8	7		3,6			1,2,4,5		
11	2		4,5	8	1,6		7	3	
12	5,6,7,8	2,4					1,3		
13	3	6			4		1,2,5,8	7	
14		1,2,5		6			3,4,7,8		
15		6	5	1,2,4,8			7		3
16		3,5,8		4	1		6,7	2	
17	6,7			1,3			2,4,5,8		
18		1,3,4			5	7	2	8	6
19		4,8		1,2,3,5,7			6		
20		1,6,7,8		2			3,5		4
21	2			5,6	8	7	1	3,4	
22		2,4,7,8			3,6		1,5		
23	2,3,7		1	4,5,8			6		
24			4,5		2,3,6,7,8				1

Table 6.10: Subject's Behavior in all Rounds of Both Treatments
The table reports each subjects behavior in each of the eight sets. Subjects #1 through #12 played in the low-cost treatment, and subjects #13 through #24 played in the high-cost treatment. Each row lists behavior in sets 1 through 8.

Using the levels of sophistication framework laid out in Section 5.1.1, each of the 24 subjects' behaviors was classified into one of three levels when possible. Two additional levels are reported for those subject's actions that straddle two adjacent levels. Looking at behavior in all sets, 4 subjects used primarily honest forecasting strategies and trusting trading strategies (equivalent to level-0 degree of sophistication). 3 subjects used honest and non-honest forecasting strategies an equal number of times, and used primary trusting trading strategies (between level-0 and level-1 degrees of sophistication). 10 subjects used primary non-honest forecasting strategies and trusting trading strategies (equivalent to level-1 degree of sophistication). 2 subjects used primarily non-honest reporting strategies and skeptical and trusting trading strategies an equal number of times (between level-1 and level-2 degrees of sophistication). 6 subjects used primarily non-honest reporting strategies and skeptical trading strategies (equivalent to level-2 degree of sophistication). 23 of 24 subjects' behavior were classifiable into levels of sophistication framework; the remaining subject was honest as a reporter and skeptical as a shareholder. These results are reported in the first column of Table 6.11.

To gain insight as to whether subjects' behavior changed over time, each subject's behavior was alternatively classified using only select sets within the experiment. In the second column of Table 6.11, subjects' behavior is reported using only the last six of the eight sets. When dropping the first two sets, 23 of 24 subjects' behaviors were classifiable into the aforementioned framework. In the last column of 6.11, subjects' behavior is classified based upon overall analyst behavior and shareholder behavior is classified only after playing the role of analyst. That is, only the sets subsequent to playing the role of analyst were used to classify shareholder behavior. 23 of 24 subjects' behavior were classifiable into levels of sophistication framework using this modification. Using this later classification, 3 subject's shareholder behavior changed from being classified as trusting to mixed/skeptical. Two of these subjects used trusting strategies and mixed/skeptical strategies an equal amount of times. The other had only one observation of shareholder behavior using a mixed strategy.

To determine whether playing the role of analyst altered strategies, each subjects' trading behavior was classified into two categories: trading behavior before playing the role of analyst, and trading behavior after. The portion using skeptical strategies was calculated across all subjects and is reported in Table 6.12. Trading behavior subsequent to playing the role of analyst was further categorized conditional upon the forecasting strategy, coding forecasting behavior as honest, else deceptive. If the subject altered forecasting strategy over the 8 sets, the most recent forecasting

	Using All Sets	Using Sets 3 through 8	Using Post-Analyst Sets
Honest Analyst Trusting Shareholder (level-0)	4	3	4
Honest/Deceptive Analyst Trusting Shareholder	3	1	1
Deceptive Analyst Trusting Shareholder (level-1)	10	12	9
Deceptive Analyst Trusting/Skeptic Shareholder	2	1	3
Deceptive Analyst Skeptic Shareholder (level-2)	6	6	6

Table 6.11: Observed Levels of Sophistication
Each subject's observed set behavior was used to classify that subject's level of strategic sophistication. In the first column all sets are used. In the second column only the last six of eights sets are used. In the last column, shareholder behavior was classified only using sets after the subject played the role of analyst. Number of subjects fitting each classification are reported. 23 of 24 subjects fit into the listed categories.

Table 6.12: Proportion of Sets Where Subjects Used a Skeptical Trading Strategy

Treatment	Before Playing Role of Analyst		After Playing Role of Analyst		After Using Deceptive Strategy		After Using Honest Strategy	
Combined	9/46	19.57%	38/82	46.34%	28/65	43.08%	10/17	58.82%
Low-cost	3/25	12.00%	17/39	43.59%	12/30	40.00%	5/9	55.56%
High-cost	6/21	28.57%	21/43	48.84%	16/35	45.71%	5/8	62.50%

strategy was used for this secondary classification.

After playing the role of analysts, subjects appear to trade more skeptically when playing the role of shareholders. However, the differences within each treatment are statistically insignificant using the Chi-square test for equality of distributions (low-cost treatment: $X^2 = 2.96$, DF = 1, $p \leq 0.085348$; high-cost treatment: $X^2 = 0.58$, DF = 1, $p \leq 0.44631$) and only weakly significant using both treatments ($X^2 = 3.75$, DF = 1, $p \leq 0.052808$). Interestingly, subjects are more skeptical after using honest forecasting strategies versus deceptive forecasting strategies, but again, this difference is statistically insignificant.

6.1.4 Metrics for Trade and Forecasting Consistency

Metrics for trade are constructed rather than relying upon the random realizations generated during the experiment, as described in Section 4.1.3. These metrics are free of any stochastic noise inherent in the realizations. Parameters for the states, signals, and probabilities thereof are described in Section 4.1.

Likelihood of Trade

This metric is defined for a round in a single economy. This captures the probability the shareholders will trade given the input strategies of the subjects.

$$\sum_{States} \sum_{Signals} Pr(state) Pr(signal|state) I_f(f)$$

$I_f(f)$ is an indicator function, defined in every round for the three subjects within a single economy. The function is 1 if there is trade for the released forecast f and 0 otherwise. There is trade if the input bid for the released forecast is greater than, or equal to, the input asking price for the released forecast. The forecast released is the result of the analyst's input reporting strategy for the given signal.

Likelihood of Capturing Gains to Trade

This metric is also defined in every round in a single economy, but captures the probability the shareholders will trade given the state is favorable. While the prior metric was summed over all states, this metric is summed over states A and D.

$$\sum_{\{A,D\}} \frac{Pr(state) \sum_{Signals} Pr(signal|state) I_f(f)}{\sum_{\{A,D\}} Pr(state)}$$

This metric is agnostic as to the distribution of gains of trade, and is akin to a social welfare metric.

Benefit of Trading

This metric is defined for a round in a single economy, but unlike the measure above, captures the expected benefit to an individual shareholder. To avoid notational clutter, the metric is defined separately for both shareholders.

Shareholder 1

$$\sum_{States} \sum_{Signals} Pr(signal|state) I_f(f) I_{state} \{T(f) - v_1(state) + \frac{1 - I_{state}}{2} \lambda\}$$

Shareholder 2

$$\sum_{States} \sum_{Signals} Pr(signal|state) I_f(f) I_{state} \{v_2(state) - T(f) - \frac{1 - I_{state}}{2} \lambda\}$$

I_{state} is an indicator function capturing the direction of trade. The indicator is -1 if the state is A or B, and 1 otherwise. $T(f)$ is the transfer price–the minimum of the input bid and input ask for the released forecast. Both $v_1(.)$ and $v_2(.)$ are the payoffs to shareholder 1 and 2, respectively, given the state. The payoffs, and the costs of trading, λ, are treatment parameters described in Section 4.1.

Metrics of Forecasting Consistency

Two consistency metrics are constructed, each based on the analyst's forecasting strategy, using the specific mapping from the analyst's private signals to forecasts.

The first metric is constructed as follows. In the first round of a set, the metric is zero. In subsequent rounds, if the analyst's mapping is identical to mapping used in the last round, the metric is increased by one for the current round. If the mapping is different, the metric is reset to zero. For example, an analyst who used the same mapping in every round of a set would have scores of $\{0, 1, 2, 3, 4, 5, 6, 7\}$ in rounds one through eight respectively. An analyst who choose to forecast A and D always in first four rounds, but B and C always in the last four rounds would have scores of $\{0, 1, 2, 3, 0, 1, 2, 3\}$ in rounds one through eight, respectively.

The sum of these scores is used to arrive at a consistency metric for the set. For the first example above, the set metric would be 28 while the second would be 12. The metric captures how consistently the analyst chooses to forecast, not necessarily how consistently the analyst's set decision was classified. Note that in both examples above the analysts were deceptive in every round, but differ in how the deceptive strategy was executed.

The second metric is simply the longest string of consistent round-to-round mappings used in a set. In the examples above, the longest mapping strings are 7 and 3. This is equivalent to the maximum of the round scores.

Mean and Inter-quartile Ranges of Metrics

Statistics for constructed metrics are shown in Tables 6.13 and 6.14 for the low-cost and high-cost treatments respectively. Metrics are constructed and measured over an economy of three subjects, yielding 32 observations for each treatment. Each economy is independent of all others. Discussion of these metric values is included in Section 6.2.

	Mean	Minimum	1st Quartile	Median	3rd Quartile	Maximum
Longest Forecast Mapping String in Set		1.0	2.0	3.0	4.0	7.0
Consistency Metric For Set		0.0	1.5	3.5	9.0	28.0
Average Likelihood of Trade	25.9%	0.0%	14.1%	25.0%	40.6%	50.0%
Average Likelihood Gain Captured	29.1%	0.0%	19.1%	26.3%	43.8%	61.3%
Average Benefit of Trading for Shareholder 1	2.741	-0.469	1.109	2.094	3.875	8.938
Average Benefit of Trading for Shareholder 2	-0.673	-5.438	-1.594	-0.359	0.719	3.250

Table 6.13: Metric Statistics: Low-Cost Treatment

	Mean	Minimum	1st Quartile	Median	3rd Quartile	Maximum
Longest Forecast Mapping String in Set		0.0	1.5	3.5	7.0	7.0
Consistency Metric For Set		0.0	1.5	5.5	24.5	28.0
Average Likelihood of Trade	31.5%	3.1%	18.8%	28.1%	42.2%	75.0%
Average Likelihood Gain Captured	34.8%	0.6%	18.8%	32.8%	44.1%	86.3%
Average Benefit of Trading for Shareholder 1	1.548	-1.906	0.125	1.125	2.641	7.313
Average Benefit of Trading for Shareholder 2	-1.600	-7.719	-2.797	-1.506	0.078	3.688

Table 6.14: Metric Statistics: High-Cost Treatment

6.2 Experimental Results

In this Section, the behavior of subject, metrics and allocations listed in Table 5.1 are analyzed. After reviewing the results, the surviving hypothesis is selected from the competing hypotheses listed in Section 5.2.

Result 1. *Analysts tend to avoid honest forecasting strategies.*

Subjects adopted deceptive or mixed forecasting set strategies 70.3% of the time. 15 subjects tended to adopt deceptive and/or mixed forecasting strategies and 5 subjects tended to use honest forecast strategies. Using a Chi-square test for known distributions, I reject the null hypothesis that half the analysts use primarily trusting strategies ($X^2 = 7.04$, DF = 1, $p \leq 0.0079709$). This result in consistent with prior work that document subjects' use of deceptive messages (Dickhaut et al. (1995); Blume et al. (1998); Cai and Wang (2006); Sanchez-Pages and Vorsatz (2006)).

Result 2. *Shareholders tend to use trusting trading strategies.*

Analysts adopted deceptive or mixed forecasting set strategies 70.3% of the time. 15 subjects tended to adopt deceptive and/or mixed forecasting strategies and 5 subjects tended to use honest forecast strategies. These results are consistent with prior work that document subjects' use of deceptive messages (Dickhaut et al. (1995); Blume et al. (1998); Cai and Wang (2006); Sanchez-Pages and Vorsatz (2006)).

Clearly subjects did not exhibit homogeneous forecasting behavior. Using either a Chi-square test for known distributions, or a test of binomial proportions, I reject the null hypothesis that all subjects used primarily deceptive strategies and the null hypothesis that all subjects used honest strategies.

Despite heterogeneous behavior, can any level of strategic sophistication, as described in Section 5.1.1 explain overall observed behavior better than the other levels? Given that not all analysts were honest and not all analysts were deceptive, is it the case that the proportions are non-equal? Using a Chi-square test for known distributions, I reject the null hypothesis that half the analysts use primarily trusting strategies ($X^2 = 7.04$, DF = 1, $p \leq 0.0079709$).

As with analyst behavior, subjects clearly do not exhibit homogeneous trading behavior. Any null hypothesis of homogeneous trading behavior is openly rejected. However, the aim of this work is to determine whether one level of strategic sophistication can explain overall observed behavior.

Thus, I test whether the proportion of subjects using trusting strategies is greater than those using skeptical and/or non-trading strategies. Using a Chi-square test for known distributions, I reject the null that half the shareholders used skeptical strategies ($X^2 = 5.04$, DF $= 1$, $p \leq 0.0.024768$).

Result 3. *The likelihood of trade and the likelihood of capturing gains to trade were not equal across treatments.*

The mean likelihood of trade was 25.9% and 31.5% in the low-cost and high-cost treatments, respectively. The mean likelihood of capturing gains to trade, trading when the state is A or B, was 29.1% and 34.8% in the low-cost and high-cost treatment respectively.

If overall subject behavior was best captured by a level-0 degree of sophistication, (honest and trusting), the likelihood of trade and the likelihood of capturing gains to trade would be equal across treatments. I reject the null hypotheses that the medians of these measures are equal over the two treatments using a Wilcoxon Two Sample Test (likelihood of trade: $W = 61962.5$, $p \leq 0.02703$; likelihood of capturing gains: $W = 62010.5$, $p \leq 0.02908$).

If overall subject behavior was captured by a level-2 or higher level of sophistication (as per Section 5.1.1), then the likelihood of traded and likelihood of capturing gains would be higher in the low-cost treatment and lower in the high-cost treatment. As this relationship is reversed in the experiments, with the likelihoods being greater in the high-cost treatment than in the high-cost treatment, the prior Wilcoxon Two Sample Test is sufficient to reject the null hypothesis that the likelihoods are greater in the low-cost treatment. Within the next result I examine why the likelihood of trade is greater in the high-cost treatment than in the low-cost treatment.

Result 4. *The likelihood of trade is correlated with analyst forecast consistency.*

Table 6.15 reports the rank correlation coefficient of the two constructed consistency metrics against measures of trade. The p-values reported are for the null of no monotonic relation between the variables. While the coefficient is positive and significant over both treatments, it is larger and more significant when isolating deceptive analysts from honest, despite reducing the number of observations. As discussed in Section 5.1, consistency would be positively related to the likelihood of trade when a deceptive analyst is paired with trusting shareholders, but consistency would be negatively related to the likelihood of trade when a honest analyst is paired with trusting shareholders. The rank correlation is positive, albeit statistically insignificant, when isolating honest analysts from deceptive.

This relationship sheds light on why there is more trade, on average, in the high-cost treatment than the low-cost treatment. In the high-cost treatment, the consistency metric was greater than or equal to the metric in the low-cost treatment at each inter-quartile value. The differences in the consistency between treatments were statistically significant (Wilcoxon Signed-Ranks Test, W+ = 192.50, W- = 17.50, N = 20, $p \leq 0.0003948$).

Table 6.15: Correlation Between Consistency and Trade

	Likelihood of Trade	Likelihood of Capturing Gain
For both treatments (64 observations)		
Set Consistency Metric	0.3368 $p \leq 0.007505$	0.346 $p \leq 0.006022$
Longest Mapping String in Set	0.306 $p \leq 0.01515$	0.2705 $p \leq 0.03177$
For Deceptive Analysts in both treatments (47 observations)		
Set Consistency Metric	0.5572 $p \leq 0.000026$	0.5249 $p \leq 0.00015$
Longest Mapping String in Set	0.4543 $p \leq 0.001342$	0.3939 $p \leq 0.006233$
For Deceptive Analysts in low-cost treatment (22 observations)		
Set Consistency Metric	0.5802 $p \leq 0.0046$	0.434 $p \leq 0.04397$
Longest Mapping String in Set	0.4873 $p \leq 0.021234$	0.4072 $p \leq 0.060437$
For Deceptive Analysts in high-cost treatment (25 observations)		
Set Consistency Metric	0.5572 $p \leq 0.003793$	0.532 $p \leq 0.006242$
Longest Mapping String in Set	0.4428 $p \leq 0.026555$	0.3902 $p \leq 0.054078$

Rank correlation coefficient reported.
P-values are for null of no monotonic relationship.

Result 5. *Gains from trading were greater for shareholder 1 than for shareholder 2.*

The benefit of trade is greater for both shareholders in the low-treatment compared to high-cost treatment in all quartiles, due to the lower trading costs. The average benefit of trading for shareholder 1 was 2.741 and 1.548 in the low-cost and high-cost treatments, respectively. The average benefit of trading for shareholder 2 was -0.673 and -1.600 in the low-cost and high-cost treatments, respectively. I reject that null hypotheses that difference between shareholders gains in each round has a median value of zero using the Wilcoxon Matched-Pairs Signed-Ranks Test (low-cost treatment: W+ = 7992.50, W- = 1460.50, N = 137, $p \leq 2.298e^{-12}$; high-cost treatment: W+ = 9462, W- = 2166, N = 152, $p \leq 1.965e^{-11}$).

Overall trade benefited shareholder 1 more so than shareholder 2. This resulted from trading at prices P1 and P4 when the analyst is using a deceptive forecast strategy, so shareholder 2 sells for too little and buys for too much. The rank correlation coefficient between the shareholders' benefits of trading in each round is negative and significant in both treatments (low-cost treatment: R = -0.4234, N = 256, $p \leq 1.376e^{-11}$; high-cost treatment: R = -0.7803, N = 256, $p \leq 1.27e^{-35}$).

Result 6. *Shareholders who elected to initially trust analysts did not revise their strategies when facing deceptive analysts.*

15 of the 24 subjects adopted a trusting trading strategy in the first round in a majority of sets. For this subset of subjects, conditionally upon trusting in the first round, I coded the subjects' subsequent set-strategy using the trusting score shown in Table 6.7. If a subject had an average trusting score less than or equal to .5 for the subsequent rounds, the subjects was coded as continuing to use a trusting strategy. Otherwise, the subject was coded as using a revised strategy. The results of this coding, sorted by the analyst's strategy within the economy, are shown in Table 6.16. The sorting of analysts separates deceptive analysts who adopted honest and partially deceptive strategies an equal number of times within a set. These analysts are labeled as mimics in this analysis.

Using the Friedman Test, I fail to reject the null hypothesis that shareholders revise similarly against dissimilar forecasting behaviors (W = 0.11, Q = 0.33 $\sim X^2$, DF = 1, $p \leq 0.56566$). The result holds when grouping mimic analysts with deceptive rather than with honest, or omitting the observations with mimic analysts altogether. The result also holds when classifying subsequent trading behavior using the mode strategies across the set rather than the trust score.

Table 6.16: Subsequent Strategies of Trusting Shareholders

	Revised Strategy	Remained Trusting
Honest Analyst	4	13
Mimic Analyst	5	13
Deceptive (not Mimic) or Mixed Analyst	3	20

Hypotheses Selection

Taken collectively, the results support Hypothesis 2 from the competing hypotheses. Subjects' overall behavior is best described as boundedly rational, as most subject chose deceptive forecasting strategies, but in turn, subjects chose trusting trading strategies. This behavior is consistent with some level of, albeit a limited level, of strategic sophistication. The likelihood of trade and capturing gains was significantly correlated with forecasting consistency when the analyst adopted a deceptive forecasting strategy. Correlation the likelihood of trade with forecasting consistency is only coherent with a limited level of strategic sophistication. Since shareholders did not revise their trading trusting strategies when facing deceptive forecasting strategies, shareholder 1 benefited from trade at the expense of shareholder 2. Inequity of gains to trade is consistent with a limited level of strategic sophistication.

6.3 Conclusion

Departing from the homogenous reporting strategies predicted in sender-receiver games, Crawford (2003) shows the existence of honest sender types can cause a rational individual to mimic boundedly rational behavior. This and subsequent research raises the possibility of overall honest and trusting behavior in spite of incentives to deceive. This possibility presumes the ability of receivers to determine when the sender was deceptive. This same trait of cheat-detection is necessary for the preservation of socially normative behavior when selfish behavior would otherwise destroy social gains created by off-equilibrium behavior (Cosmides and Tooby (2005)).

To this end I study the behavior of a group of subjects in a moderately complex game of information transmission. In keeping with prior work, I find analysts send more information than predicted by models assuming sequentially-rational agents. Likewise, I find shareholders rely on information sent more than predicted. Unlike Sanchez-Pages and Vorsatz (2006), I do not find that normative social behavior is driving the overcommunication phenomenon. Subject's behavior is more analogous to hierarchical models of strategic thinking (Camerer et al. (2004); Stahl and Wilson (1994)). While less than half the subjects chose to use honest forecasting strategies, more than half chose to use trusting trading strategies. This pattern is consistent with hierarchical models with limited iterations of reasoning: when playing the role analyst, a subject assumes others are habitually trusting, and thus adopts a deceptive forecasting strategy. When playing the role of shareholder, the subject assumes others are habitually honest, and thus adopts a trusting strategy.

In further research I intend to explore the possibility that subjects adopt simple heuristic behavior as the game becomes more complex, but use predicted strategies in less complex games (Costa-Gomes et al. (2000); Costa-Gomes and Crawford (2004)). By reducing the shareholders' decision to a simple binary choice, or simplifying the information structure, will subjects be better able to coordinate on socially normative behavior to capture gains?

Appendix A

Gains To Trading Figures

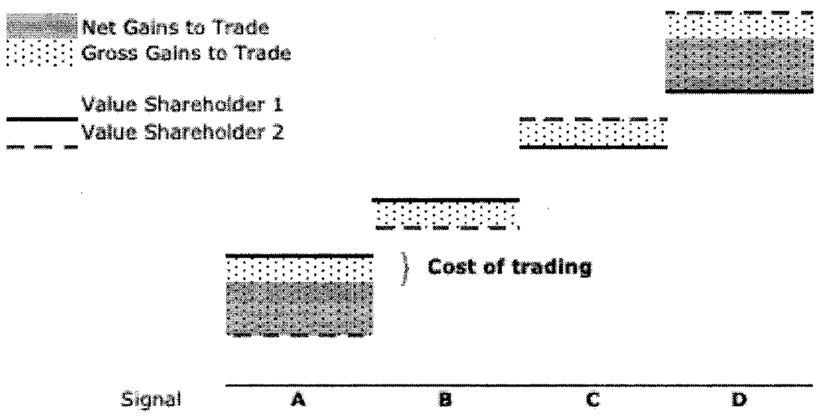

Figure A.1: Gains to Trade Given Honest Forecasting Strategy
In both the high and low cost treatments, the cost of trading was such that shareholders had ex-ante gains to trade when the analyst's private signal suggested the state of nature was likely A or likely D, but not when the signal suggested the state was likely B or likely C.

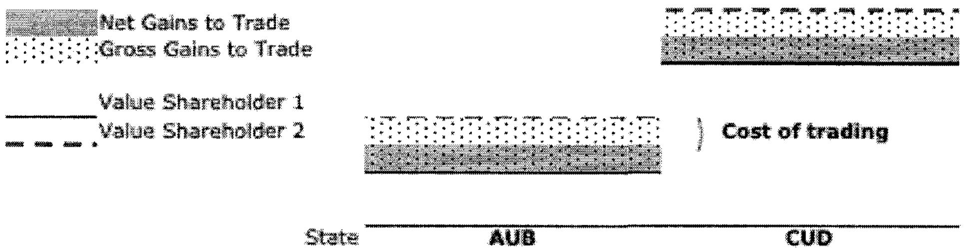

Figure A.2: Gains to Trade Given Deceptive Forecasting Strategy
In the low-cost treatment trading costs were such that shareholder had ex-ante expected gains to trade despite the analyst's use of a deceptive forecasting strategy, as the gain to trading in the extreme states, A or D, was greater then the loss of trading in the interior states B and C. This relationship is illustrated in the figure. In the high-cost treatment the loss from trading in the interior states B and C was greater than the gain from trading in the extreme states A and D, and as such, there are not net gains to trading.

Appendix B

Screen Captures of Experiment

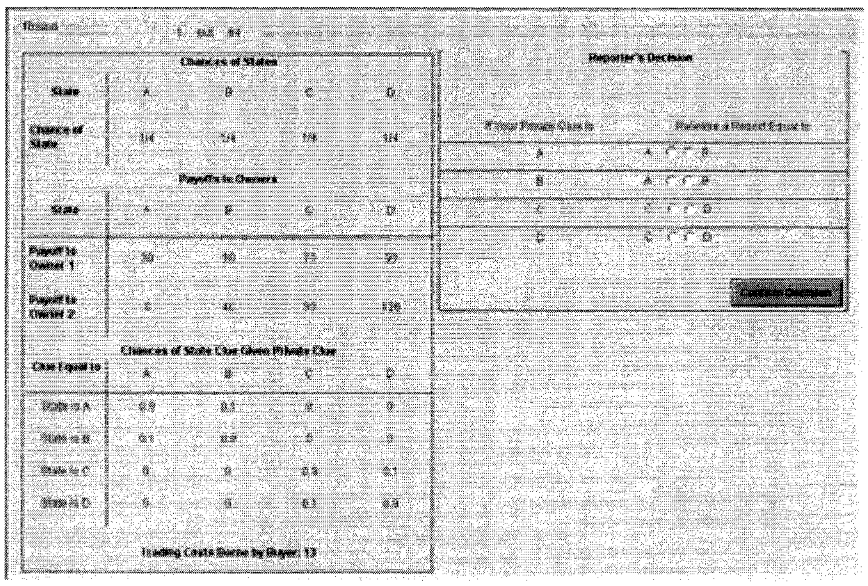

Figure B.1: Entry screen for Analyst in Low-Cost Treatment

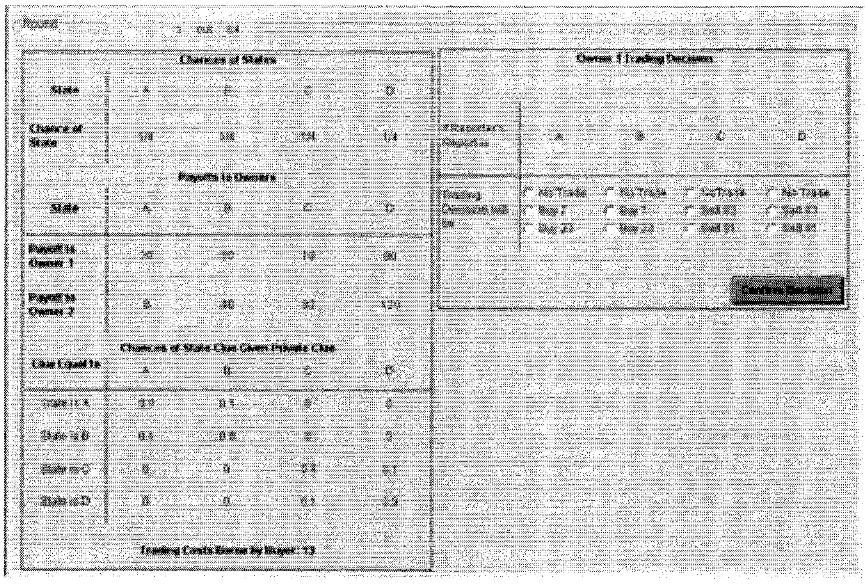

Figure B.2: Entry screen for Shareholder 1 in Low-Cost Treatment

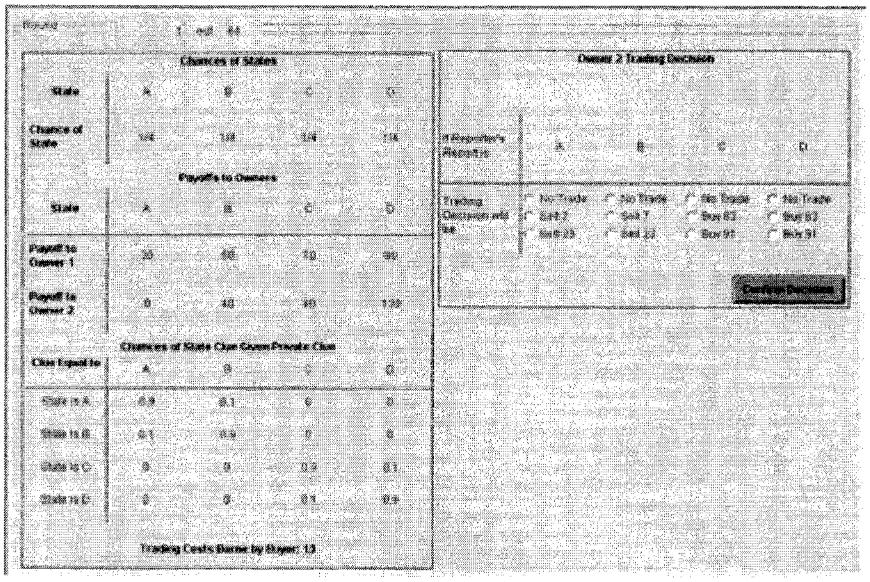

Figure B.3: Entry screen for Shareholder 2 in Low-Cost Treatment

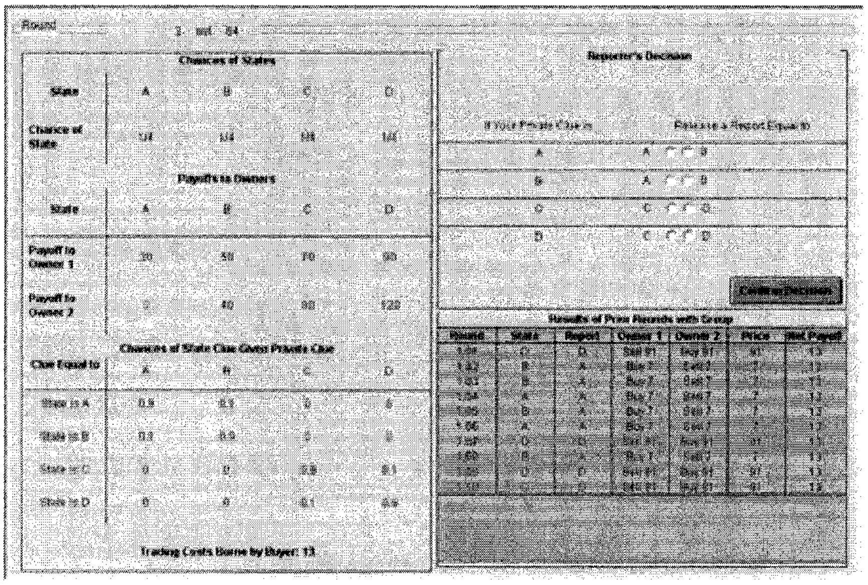

Figure B.4: Entry screen with Feedback for Analyst in Low-Cost Treatment

Appendix C

Statistical Plots

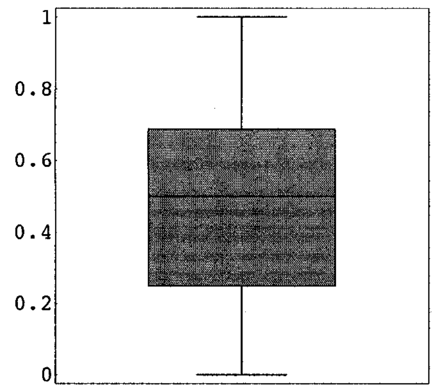

(a) Box-whisker Plot

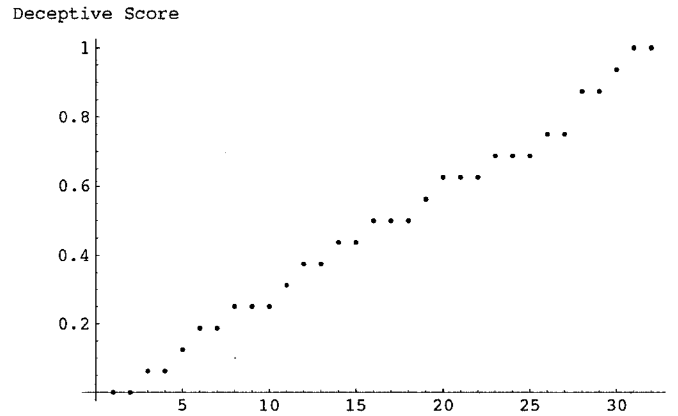

(b) List Plot

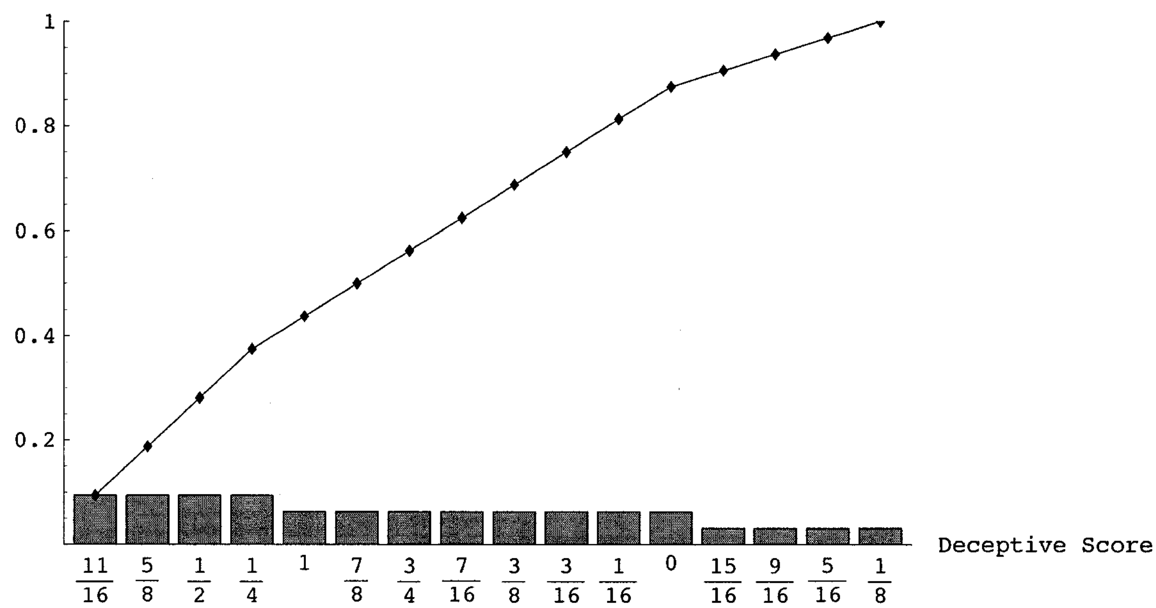

(c) Pareto Plot

Figure C.1: Analyst Set Scores in Low Cost Treatment
Within each set of eight rounds, each analyst's deceptive score was calculated as per Table 6.4. The average deception score over the set is reported for all 32 sets in the low-cost treatment. An analyst who was always deceptive in every round of a set would score 1 (or 0.5 if partially deceptive), whereas an analyst who was always honest in every round of a set would score 0. Plot (a) shows the quartile scores over 32 sets within the colored box (0.25, 0.50 and 0.688), and the minimum and maximum scores of 0 and 1 respectively. The 32 average deceptive scores are arranged sequentially in the plot (b), each point representing a set, while in plot (c) the scores are arranged by relative frequency. Plots (b) and (c) both show the majority of sets contain deceptive behavior.

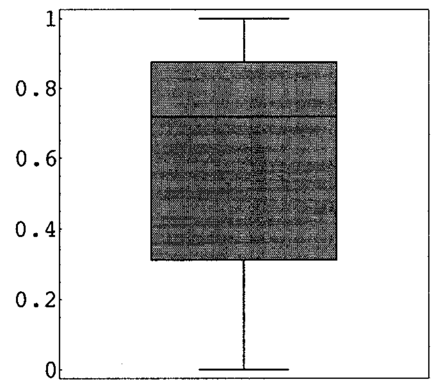

(a) Box-whisker Plot

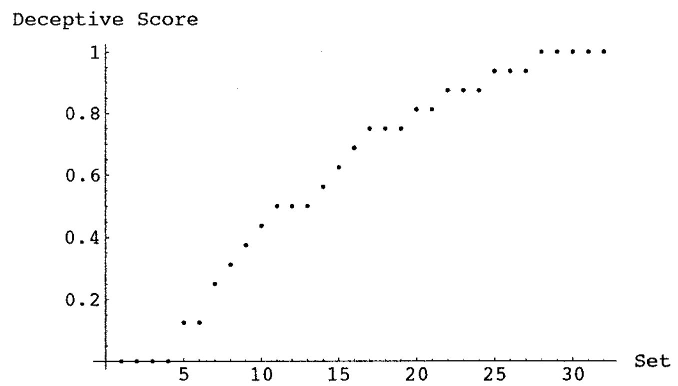

(b) List Plot

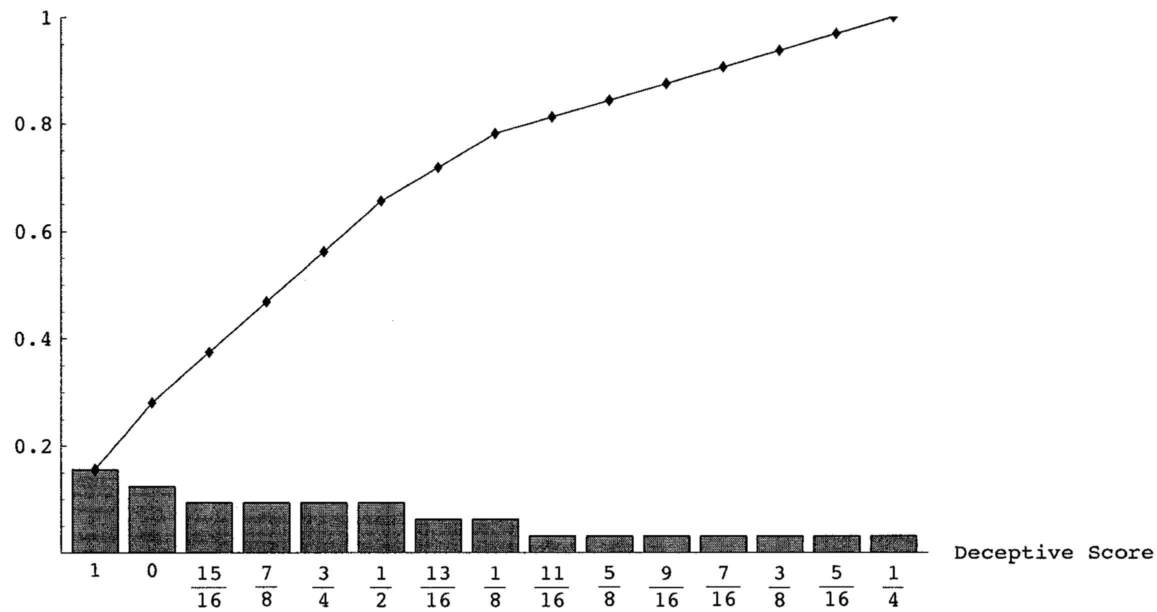

(c) Pareto Plot

Figure C.2: Analyst Set Scores in High Cost Treatment
Within each set of eight rounds, each analyst's deceptive score was calculated as per Table 6.4. The average deception score over the set is reported for all 32 sets in the low-cost treatment. An analyst who was always deceptive in every round of a set would score 1 (or 0.5 if partially deceptive), whereas an analyst who was always honest in every round of a set would score 0. Plot (a) shows the quartile scores over 32 sets within the colored box (0.359, 0.719, 0.891), and the minimum and maximum scores of 0 and 1 respectively. The 32 average deceptive scores are arranged sequentially in the plot (b), each point representing a set, while in plot (c) the scores are arranged by relative frequency. Plots (b) and (c) both show the majority of sets contain deceptive behavior.

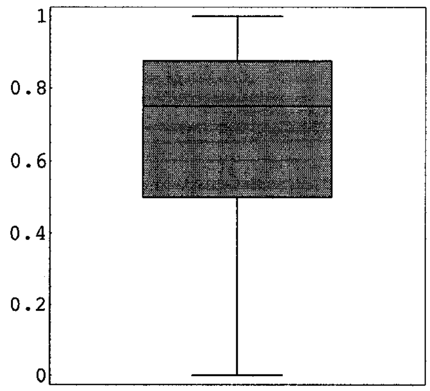

(a) Box-whisker Plot

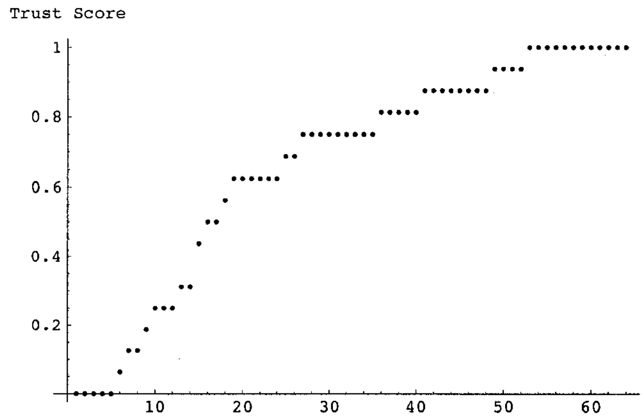

(b) List Plot

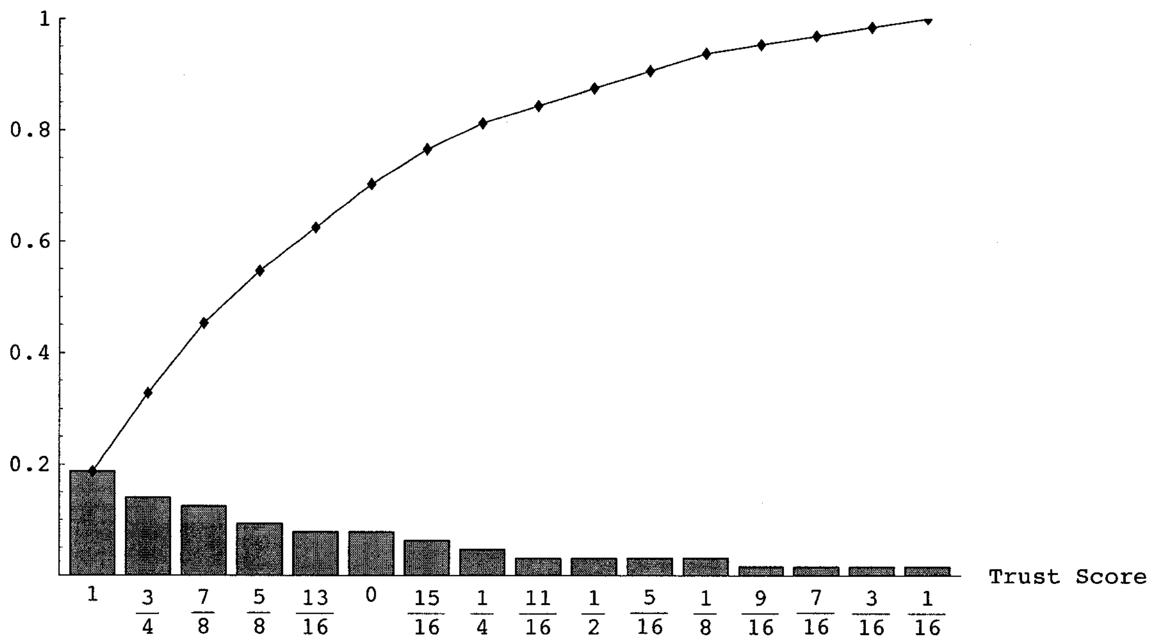

(c) Pareto Plot

Figure C.3: Shareholder Set Scores in Low Cost Treatment
Within each set of eight rounds, each shareholder's trusting score was calculated as per Table 6.7. The average trusting score over the set is reported for all 32 sets in the low-cost treatment. A shareholder who was always trusting in every round of a set would score 1, whereas a shareholder who was always skeptic in every round of a set would score 0. Plot (a) shows the quartile scores over 32 sets within the colored box (0.50, 0.75 and 0.891), and the minimum and maximum scores of 0 and 1 respectively. The 32 average trusting scores are arranged sequentially in the plot (b), each point representing a set, while in plot (c) the scores are arranged by relative frequency. Plots (b) and (c) both show the majority of sets contain trusting behavior.

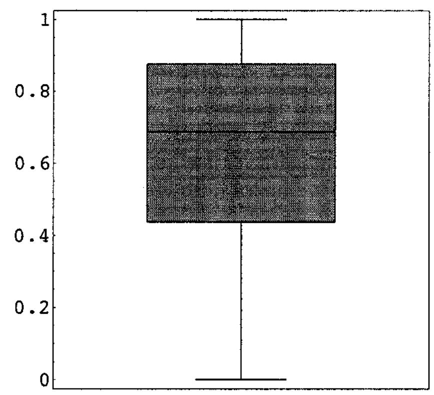

(a) Box-whisker Plot

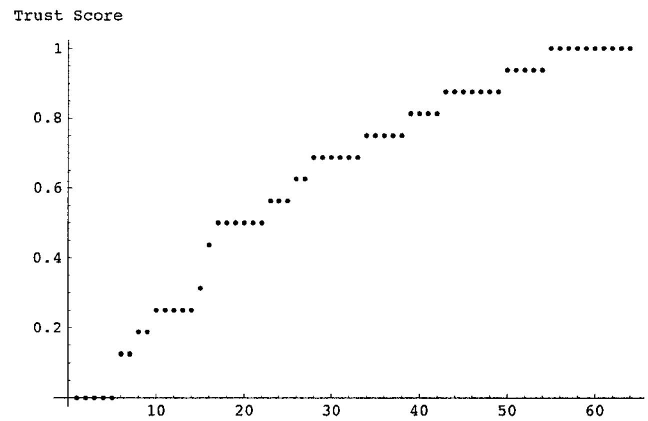

(b) List Plot

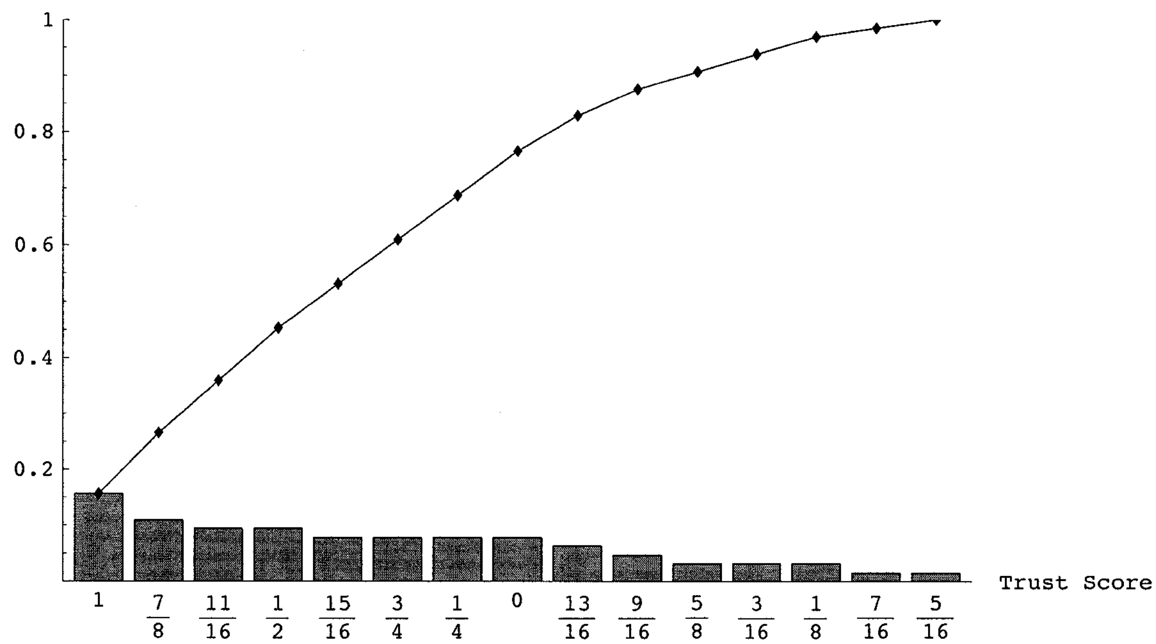

(c) Pareto Plot

Figure C.4: Shareholder Set Scores in High Cost Treatment
Within each set of eight rounds, each shareholder's trusting score was calculated as per Table 6.7. The average trusting score over the set is reported for all 32 sets in the high-cost treatment. A shareholder who was always trusting in every round of a set would score 1, whereas a shareholder who was always skeptic in every round of a set would score 0. Plot (a) shows the quartile scores over 32 sets within the colored box (0.484, 0.688, 0.875), and the minimum and maximum scores of 0 and 1 respectively. The 32 average trusting scores are arranged sequentially in the plot (b), each point representing a set, while in plot (c) the scores are arranged by relative frequency. Plots (b) and (c) both show the majority of sets contain trusting behavior.

Appendix D

Experimental Instructions

Instructions

This is a computerized experiment in the economics of decision-making. This experiment will last approximately two hours.

To make a profit, you will be trading a financial asset, which may lose or gain money. You are guaranteed to receive $10 for showing up on time. By following the instructions carefully and making good decisions, you may earn an additional amount of money besides the $10 participation fee. The actual amount of additional money that you may earn will depend on your decisions and the decisions of other participants. Your money will be paid to you in cash after the experiment ends. You will need to sign and date a compensation receipt form before you receive your payment

There are some rules you must follow:

(1) Do not talk to others at any time during the experiment.
(2) You will use your computer to select decisions during the experiment. Do not use your mouse or keyboard to play around with the software running on your computer. If you unintentionally or intentionally close the software program running on your computer, you will be asked to leave. If this happens, you will receive only your $10 fee for showing up.
(3) If you have any questions during the experiment, please raise your hand. An experimenter will come to your location and answer your questions.

You are free to withdraw from the experiment at any time, for any reason. If you choose to do so, please raise your hand. In this case, you will be paid your $10 participation fee as you leave.

Details of the Experiment

Grouping of Participants

Every eight rounds, you will be randomly placed into a group of three participants. Two participants in each group will be owners (Owner 1 and Owner 2) of an asset and the other a reporter. Eight rounds make up a set of rounds. The experiment consists of multiple sets. At the beginning of each set you are randomly regrouped and the roles of reporter and owners randomly reassigned. You are not allowed to communicate with other participants during the experiment. You are not told who is in your group.

Overview of a Round

Throughout this experiment there are three participants: a Reporter and two owners of a financial asset. Each owner has one unit of the financial asset. All values, payoffs, prices, and costs will be expressed in the monetary unit lira. At the end of each round, the computer determines the State of the economy, which in turn dictates the lira value of each owner's asset. The chances of a particular State are:

Chances of States

State	A	B	C	D
Chance of State equal to	1/4	1/4	1/4	1/4

The State of the economy will be either A or B or C or D. There is an equal chance that the computer will select A or B or C or D. So, on average the computer will select A one in four times, B one in four times, etc… A higher state means the payoff from the asset is higher.

At the end of the round, if you are Owner 1 or Owner 2, you will be paid off only if you still own the asset. However, your lira payoff is different depending on which owner you are, Owner 1 or Owner 2.

Payoffs to Owners for Single Asset

State of Economy	**A**	**B**	**C**	**D**
Payoff to Owner 1	30	50	70	90
Payoff to Owner 2	0	40	80	120
Difference	30	10	10	30

The **Payoff to Owners** table shows when the State of the economy is A or B, the Owner 1 receives a higher payoff than the Owner 2. If the state is C or D, then Owner 2 receives a higher payoff than Owner 1.

If the two owners knew in advance the State, one owner might sell their asset to the other before the end of the round so that both owners are better off. However, there are costs to trading. If the owners elect to trade, then the buyer is charged 21 liras. The seller incurs no costs. How do owners come to agree upon a selling price? During the experiment, we will present multiple selling prices. If an owner chooses to buy, we will let the owner borrow, interest free, lira from future payoffs.

If both owners agree upon a selling price, the seller has no unit of the asset at the end of the round. Instead, he or she keeps the lira received from the buyer. The buyer now holds two units of the asset, so receives twice the payoffs, and pays the price to the seller. The buyer is also charged 21 liras in trading costs.

At the beginning of each round the Reporter learns a private clue of the State of the economy. No one else learns this clue. There is an equal chance that the clue is either A or B or C or D. The chances of a particular State are related to this private clue:

Chances of State Given Private Clue

If Clue is:	**A**	**B**	**C**	**D**
Chance State is A	9/10	1/10	0	0
Chance State is B	1/10	9/10	0	0
Chance State is C	0	0	9/10	1/10
Chance State is D	0	0	1/10	9/10

For example, if the Reporter's private clue is B, there is a 9 in 10 the State is B and a 1 in 10 chance the State is A. If the clue is D, there is a 9 in 10 chance the State is D and a 1 in 10 chance the State is C. So sometimes the Reporter learns a clue equal to the State; other times the Reporter learns a clue different from the State. No one other than the Reporter ever discovers the value of the private clue at any time during or after the round.

After learning his or her private clue, the Reporter releases a report to the owners. If the Reporter learns the clue is A, the Reporter can elect to report either A or B. If the Reporter learns the clue is B, they can elect to Report A or B. If the Reporter learns the clue is C, they can elect to report C or D. If the report learns the clue is D, they can elect to report C or D.

If there is trade between Owner 1 and Owner 2, the Reporter receives 13 liras at the end of the round (out of the costs of trading). Otherwise, the Reporter receives nothing.

Below, please write down your answers to the following questions. In a few minutes, an experimenter will review the correct answers with you privately.

1. What is the chance the State is B?
2. What is the chance the State is D?
3. On average, how many times out of 10 will State be B if the Reporter sees a clue of B?
4. If the Reporter's report is A, what are the possible States?
5. If the Reporter's clue is A, whose payoff from the asset is higher: Owner 1's or Owner 2's?
6. If the Reporter's clue is C, whose payoff from the asset is higher: Owner 1's or Owner 2's?
7. If both owners knew the State was A or D, could they benefit from trade?
8. If both owners knew the State was B or C could they benefit from trade?
9. If the owners trade, what will the Reporter receive?

Reporter's Decision

If you are a Reporter, you will see the following screen at the beginning of each round:

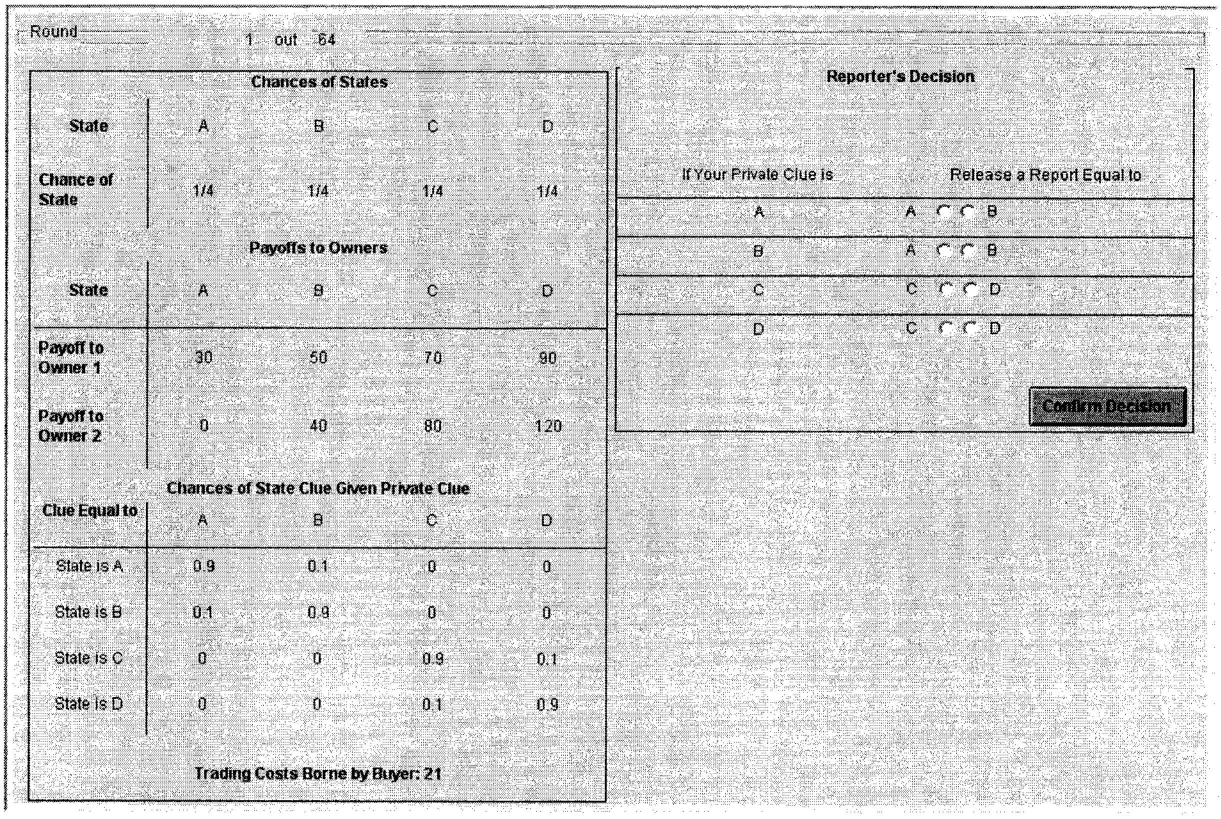

Your task as a Reporter is to fill in the entries in the **Reporter's Decision** table. For example, if your private clue were equal to B, would you decide to report A or B? For each possible clue you may receive, you need to tell the computer what you will report. After you make your decision, click on the red 'Confirm Decision' button. *You will have 60 seconds to make and confirm your decision; else you will not earn any payoff for this round. After 60 seconds expire, if you did not input and confirm a decision, the computer will randomly determine your decision for the first round with the group, or use the decision played in the prior round.*

As a reporter you may not see all the clues in a round. Perhaps only one of the four possible clues will materialize. However, we want you, as the Reporter, to say in advance what you would do in each of the four scenarios. The computer implements the decisions you make.

Remember the table **Chances of States** shows the chance of each of the State's possible values in each round. Note this table simply describes the rule computer *will* use in determining the State in each round. The table **Chances of State Given Private Clue** is simply the rule computer *will* use in generating the Reporter's private clue. It does not imply the computer has already calculated these clues before generating the above screen.

Below, please write down your answers to the following questions. In a few minutes, an experimenter will review the correct answers with you.

10. If the reporter's private clue is D, what values can the report take?
11. What is the chance that the State is D if the private clue is C?
12. If the reporter elects to report A when the clue is either A or B, what report will the owners see when the State is actually A? When the State is actually B?
13. Has the computer already generated the private clue before generating the above screen? (Yes / No)
14. Has the computer already determined the State before generating the above screen? (Yes / No)

Owners' Trading Decision

If you are Owner 1, you will see the following screen at the beginning of each round:

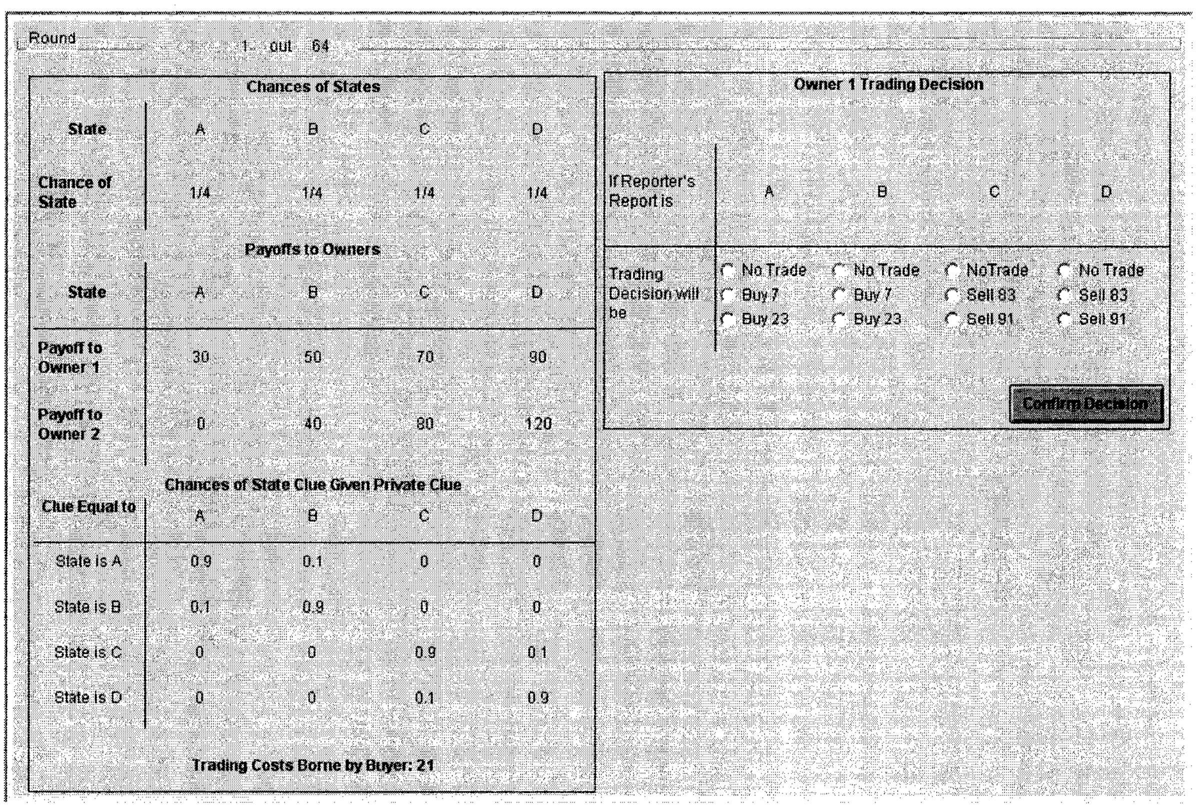

Your task as Owner 1 is to fill out the four entries in the table **Owner 1 Trading Decision**. For each of the four possible values of the Reporter's report, do you want to trade at the given lira prices? For example, if the Reporter's report were B, would you buy at 7 (but not at 23), or buy at 23 (and thus also at 7), or elect not to trade? Recall that if you buy, you also pay the trading costs of 21 liras. If you sell, you will not pay any trading costs. After you make your decision, click on the red 'Confirm Decision' button. *You have 60 seconds to make and confirm your decision. After 60 seconds expire, if you didn't input and confirm a decision, the computer will assume you do not want to trade.*

You may not see all reports in the round. For example, the Report may only choose to report B and D. So not all four reports will materialize in each round. But we want you to tell us what you would do for each possible report. Your decision will not affect which report materializes in a round.

If you are Owner 2, your screen is identical to Owner 1's, except for the decision table. Instead you will see the following screen:

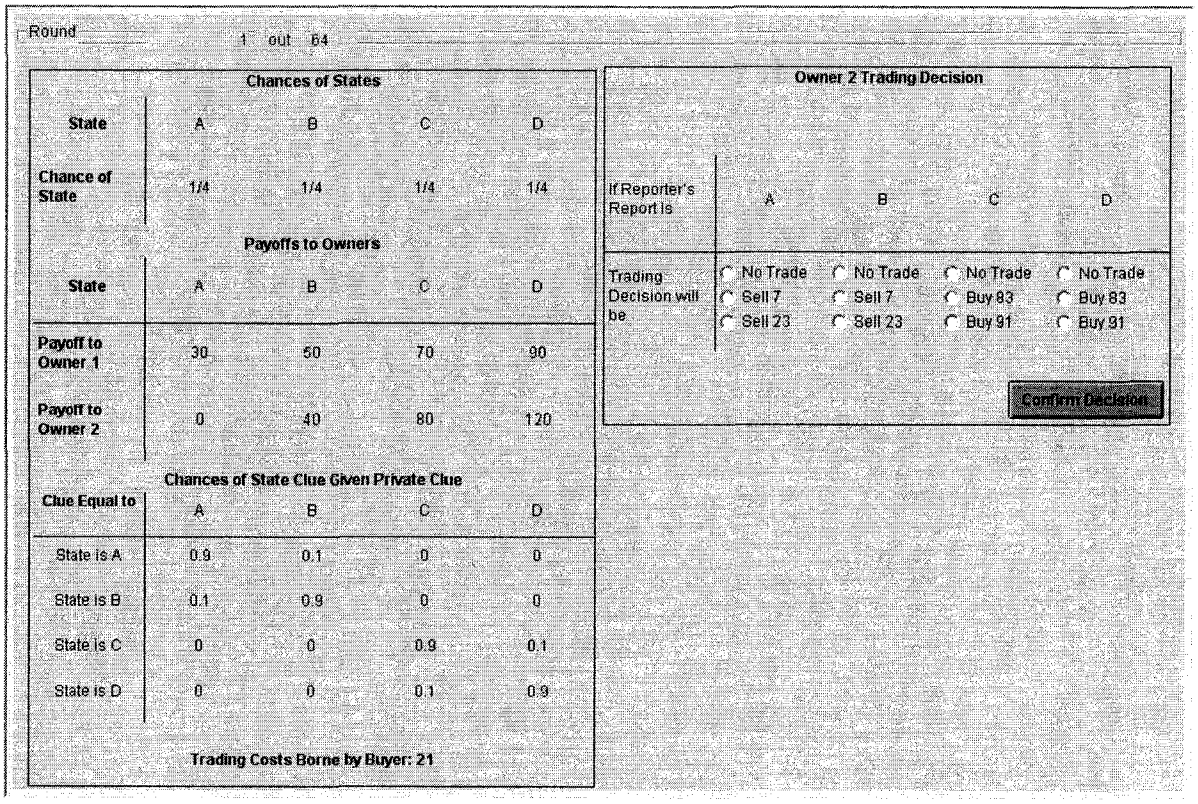

Your task as Owner 2 is to fill out the four entries in the table **Owner 2 Trading Decision**. For each of the four possible values of the Reporter's report, do you want to trade at the given lira prices? For example, if the Reporter's report is equal to B, would you sell at 7 (and therefore also at 23), or sell at 23 (but not at 7), or elect not to trade? Recall that if you buy, you pay the trading costs of 13 liras, but if you sell, you pay no trading costs.

Below, please write down your answers to the following questions. In a few minutes, an experimenter will review the correct answers with you.

15. If the State is A and there is no trade between the two owners, what is Owner 1's payoff?
16. If Owner 1 selects 'Sell at 83' when the Reporter's report is C, will owner 1 be willing to sell for 91 when the Reporter's report is C?
17. If Owner 2 elects to 'Sell at 23' when the Reporter's report is B, will Owner 2 have to sell at 7 when the Reporter's report is B?
18. Can Owner 1 elect to sell when the Reporter's report is A or B?
19. Can Owner 1 elect to sell when the Reporter's report is C or D?
20. Can the Owners elect not to trade?
21. Has the computer already generated the private clue before generating the above screen? (Yes / No)
22. Has the computer already determined the State before generating the above screen? (Yes / No)

Results of a Round

After all decisions by all participants have been input, you will see a new decision screen with a table entitled **Results of the Round with Group**. For example, if you where the Reporter, at the beginning of the 2nd round you would see:

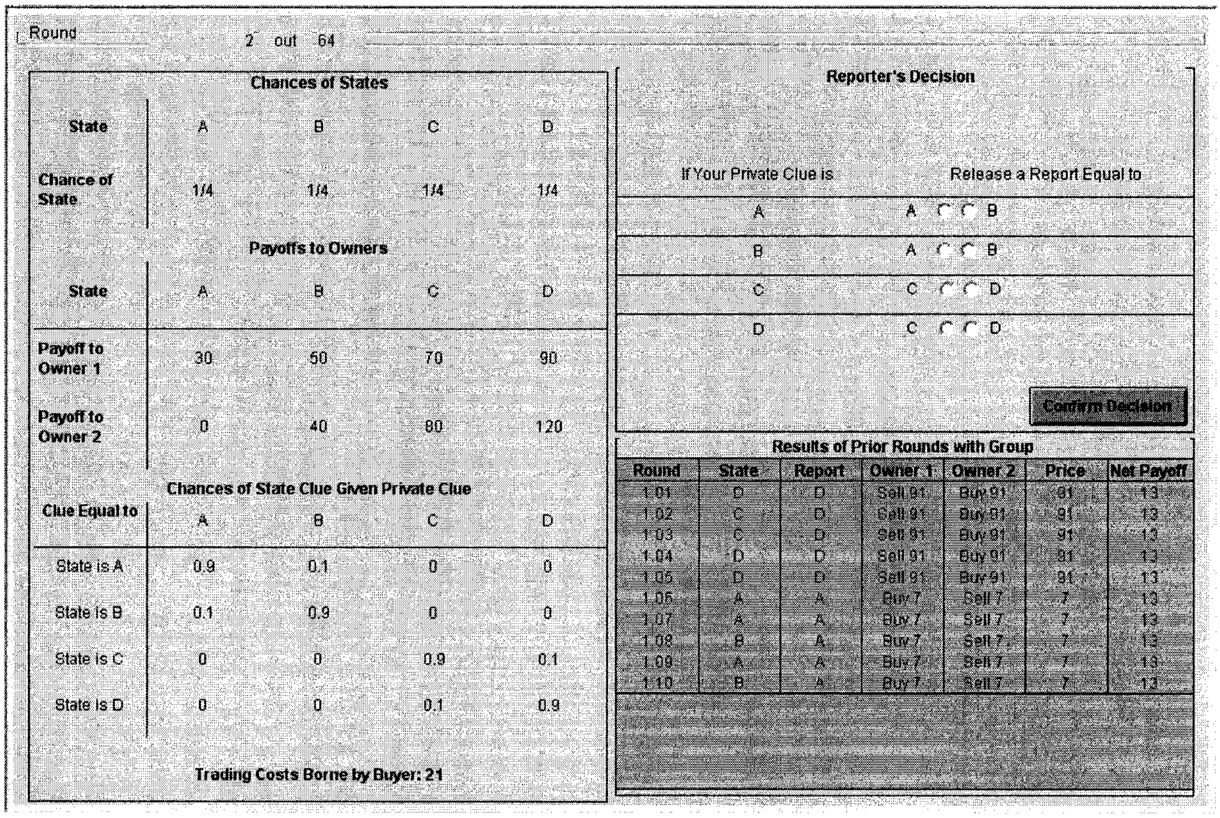

After the Reporter, Owner 1 and Owner 2 have input their decisions, the computer will draw ten realizations of the Clue and State. So you will see round 1.01 through 1.10 under the caption 'Round' for these ten realizations. The computer determines the State as per the **Chances of State** table. These appear under the caption 'State'.

Given the State, the computer then determines the Reporter's private Clue as per the **Chances of State of Economy Given Private Clue** table. Given the values input by the Reporter into **Reporter's Decision** table, the report is determined and appears under the caption 'Report'.

The decision input by Owner 1 into the **Owner 1's Decision** table in shown under the caption 'Owner 1', and likewise the decisions input by Owner 2 into the **Owner 2's Decision** table in shown under the caption 'Owner 2'.

The price is determined and appears under the caption 'Price':

- If either Owner 1 or Owner 2 elects not to trade, the price is 0.
- If one owner wants to sell at a price higher than the other is willing to buy at, the price is 0. For example, if one owner chooses 'buy at 7' and the other 'sell at 23', then the price is 0.
- If both owners agree to trade only at a single price, the single feasible price is used. For example, if one owner chooses 'buy at 23' and the other 'sell at 23', then the price is 23.
- If both owners agree to trade at more than a single price, then the lower price is used. For example, if one owner chooses 'buy at 23' and the other 'sell at 7', then the price is 7.

Under the caption 'Payoff' is your payoff. Payoffs to Owner 1 and Owner 2 will be calculated as follows:

- If there is no trade each owner has a single unit of the asset. The payoff is determined using the lira values shown in the appropriate row of the table **Payoffs to Owners**: the row 'Payoff to Owner 1' for Owner 1 and the row 'Payoff for Owner 2' for Owner 2.
- If the owner sells, then he or she has no units the asset. So, the lira payoff is equal to the price.
- If the owner is a buyer, the owner has two units of the asset. The payoff is twice the amount shown in appropriate row of the table **Payoffs to Owners**, less the price, less the cost of trading (21 liras).

If there is trade between Owner 1 and Owner 2, then the Reporter's Payoff is 13 liras. Otherwise, the Reporter Payoff is 0.

Note that the State and Price, not the Report, determine your net payoff. The possible payoffs for all States and Prices are show below. Recall that when the Price is 0, there is no trade so the payoffs are identical to those shown in the **Payoffs to Owners** table.

		Payoff to:		
State	**Price**	**Owner 1**	**Owner 2**	**Reporter**
A	*0*	*30*	*0*	*0*
A	7	**32**	**7**	**13**
A	23	16	**23**	**13**
B	*0*	*50*	*40*	*0*
B	7	**72**	7	**13**
B	23	**56**	23	**13**
C	*0*	*70*	*80*	*0*
C	83	**83**	56	**13**
C	91	**91**	48	**13**
D	*0*	*90*	*120*	*0*
D	83	83	**136**	**13**
D	91	**91**	**128**	**13**

The bold numbers represent net payoffs such that trade was beneficial to the applicable participant.

At the end of the 8 sets of 8 rounds, you will be paid some of your payoffs in addition to the $10 participation fee. The computer will randomly select ten realizations and pay the sum of your lira payoffs in those selected realizations. You will be paid $1 for every 15 liras in payoffs.

To summarize, the experiment consists of 8 sets of 8 rounds, where every set consist of:
1. Grouping of Participants
2. 8 rounds, where each round consists of:
   a. Reporter's Decision – Reporter enter their reporting decision
   b. Owners' Decisions – Both owners enter their buying/selling decision
   c. Payoffs – Computer calculates net payoffs for 10 iterations
   d. Results for the round – The Reporter and owners see the results for the rounds in the set upon the next round's decision screen.

Below, please write down your answers to the following questions. In a few minutes, an experimenter will review the correct answers with you.

23. Assume that in round 1.01:
    - i) The State is C,
    - ii) The Reporter's private clue is C,
    - iii) The Report elects to release a report of 'D' when seeing a clue of C,
    - iv) Owner 1 elects to 'Sell at 83' when seeing a report of D, and
    - v) Owner 2 elects to 'Buy at 91' when seeing a report of D.

24. Assume that in round 1.02:
    - i) The State is B,
    - ii) The Reporter's private clue is A,
    - iii) The Reporter elects to release a report of 'A' when seeing a clue of A,
    - iv) Owner 1 elects to 'Do not Buy' when seeing a report of A, and
    - v) Owner 2 elects to 'Sell at 7' when seeing a report of A.

What are the values that would be displayed in the **Results of the Round** table for rounds 1.01 and 1.02 if you are Owner 1? Owner 2? The Reporter?

Results of the Round with Group

Round	State	Report	Price	Net Payoff
1.01				Owner 1 Owner 2 Reporter
1.02				Owner 1 Owner 2 Reporter

Bibliography

Axelrod, R. (1981). The evolution of cooperation. *Science*, 211(4489):1390–1396.

Berg, J., Dickhaut, J., and McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, 10:122–142.

Blume, A., DeJong, D. V., Kim, Y.-G., and Sprinkle, G. B. (1998). Experimental evidence on the evolution of meaning of messages in sender-receiver games. *The American Economic Review*, 88(5):1323–1340.

Brennan, M. J. (2004). How did it happen? *Economic Notes*, 33(1):3–22.

Cai, H. and Wang, J. T.-Y. (2006). Overcommunication in strategic information transmission games. *Games and Economic Behavior*, 56(1):7–36.

Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898.

Chen, Q., Francis, J., and Jiang, W. (2005). The role of forecast patterns in conveying analysts' predictive ability. Working Paper.

Chen, Y. (2006). Perturbed communication games with honest senders and naive receivers. *SSRN eLibrary*.

Cosmides, L. and Tooby, J. (2005). Neurocognitive adaptations designed for social exchange. In *Evolutionary Psychology Handbook*, chapter 20. Wiley.

Costa-Gomes, M. and Crawford, V. (2004). Cognition and behavior in two-person guessing games: An experimental study. University of California at San Diego, Economics Working Paper Series 2004-11, Department of Economics, UC San Diego.

Costa-Gomes, M., Crawford, V. P., and Broseta, B. (2000). Cognition and behavior in normal-form games: An experimental study. University of California at San Diego, Economics Working Paper Series 2000, Department of Economics, UC San Diego.

Crawford, V. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50:1431–1451.

Crawford, V. P. (2003). Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions. *The American Economic Review*, 93(1):133–149.

Dawkins, R. (2006). *The God Delusion*. Houghton Mifflin.

Dickhaut, J., McCabe, K., and Mukherji, A. (1995). An experimental study of strategic information transmission. *Economic Theory*, 6:389–403.

Engle-Warnick, J. and Slonim, R. L. (2004). The evolution of strategies in a repeated trust game. *Journal of Econcomic Behavior and Organization*, 55:553–573.

Fehr, E. and Fischbacher, U. (2004). Third-party punishment and social norms. Experimental 0409002, EconWPA.

Fehr, E. and Gaechter, S. (2003). Altruistic punishment in humans. Microeconomics 0305006, EconWPA.

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

Fon, V. and Parisi, F. (2005). Revenge and retaliation. *The Law and Economics of Irrational Behavior*, George Mason Law & Economics Research Paper No. 02-31(02-31).

Guttman, I. (2005). The timing of analyst's earnings forecasts. Working Paper.

Hoffman, E., McCabe, K. A., and Smith, V. L. (1998). Behavioral foundations of reciprocity: Experimental economics and evolutionary psychology. *Economic Inquiry*, 36(3):335–52.

Hopfensitz, A. and Reuben, E. (2005). The importance of emotions for the effectiveness of social punishment. Tinbergen Institute Discussion Papers 05-075/1, Tinbergen Institute.

King-Casas, B., Tomlin, D., Anen, C., Camerer, C. F., Quartz, S. R., and Montague, P. R. (2005). Getting to Know You: Reputation and Trust in a Two-Person Economic Exchange. *Science*, 308(5718):78–83.

McCabe, K. A., Rigdon, M. L., and Smith, V. L. (2003). Positive reciprocity and intentions in trust games. *Journal of Economic Behavior & Organization*, 52(2):267–275.

Michaely, R. and Womack, K. (1999). Conflict of interest and the credibility of underwriter analyst recommendations. *Review of Financial Studies*, 12(4):653–686.

Morgan, J. and Stocken, P. C. (2003). An analysis of stock recommendations. *The RAND Journal of Economics*, 34(1):183–203.

Myerson, R. (1979). Incentive compatiblity and the bargaining problem. *Econometrica*, 47:61–73.

Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5):1313–26.

Ottaviani, M. and Sorensen, P. N. (2006). Reputational cheap talk. *RAND Journal of Economics*, 37(1):p155 - 175.

Ottaviani, M. and Squintani, F. (2006). Naive audience and communication bias. *International Journal of Game Theory*, V35(1):129–150.

Sanchez-Pages, S. and Vorsatz, M. (2006). An experimental study of truth-telling in a sender-receiver game. ESE Discussion Papers 128, Edinburgh School of Economics, University of Edinburgh.

Stahl, D. I. and Wilson, P. W. (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization*, 25(3):309–327.

Trueman, B. (1994). Analyst forecasts and herding behavior. *The Review of Financial Studies*, 7(1):97–124.

U.S. Securities and Exchange Commission (2005). Investor alert: Analyzing analyst recommendations. http://www.sec.gov/investor/pubs/analysts.htm.

Wang, J. T., Spezio, M., and Camerer, C. F. (2006). Pinocchio's pupil: Using eyetracking and pupil dilation to understand truth-telling and deception in games. Technical report, UCLA Department of Economics.