

2008s-11

Analysts, Incentives, and Exaggeration

Timothy Shields

Série Scientifique
Scientific Series

Montréal
Avril 2008

© 2008 *Timothy Shields*. Tous droits réservés. *All rights reserved*. Reproduction partielle permise avec citation du document source, incluant la notice ©.
Short sections may be quoted without explicit permission, if full credit, including © notice, is given to the source.



Centre interuniversitaire de recherche en analyse des organisations

CIRANO

Le CIRANO est un organisme sans but lucratif constitué en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du Ministère du Développement économique et régional et de la Recherche, de même que des subventions et mandats obtenus par ses équipes de recherche.

CIRANO is a private non-profit organization incorporated under the Québec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the Ministère du Développement économique et régional et de la Recherche, and grants and research mandates obtained by its research teams.

Les partenaires du CIRANO

Partenaire majeur

Ministère du Développement économique, de l'Innovation et de l'Exportation

Partenaires corporatifs

Alcan inc.
Banque de développement du Canada
Banque du Canada
Banque Laurentienne du Canada
Banque Nationale du Canada
Banque Royale du Canada
Banque Scotia
Bell Canada
BMO Groupe financier
Bourse de Montréal
Caisse de dépôt et placement du Québec
DMR Conseil
Fédération des caisses Desjardins du Québec
Gaz de France
Gaz Métro
Hydro-Québec
Industrie Canada
Investissements PSP
Ministère des Finances du Québec
Raymond Chabot Grant Thornton
State Street Global Advisors
Transat A.T.
Ville de Montréal

Partenaires universitaires

École Polytechnique de Montréal
HEC Montréal
McGill University
Université Concordia
Université de Montréal
Université de Sherbrooke
Université du Québec
Université du Québec à Montréal
Université Laval

Le CIRANO collabore avec de nombreux centres et chaires de recherche universitaires dont on peut consulter la liste sur son site web.

Les cahiers de la série scientifique (CS) visent à rendre accessibles des résultats de recherche effectuée au CIRANO afin de susciter échanges et commentaires. Ces cahiers sont écrits dans le style des publications scientifiques. Les idées et les opinions émises sont sous l'unique responsabilité des auteurs et ne représentent pas nécessairement les positions du CIRANO ou de ses partenaires.

This paper presents research carried out at CIRANO and aims at encouraging discussion and comment. The observations and viewpoints expressed are the sole responsibility of the authors. They do not necessarily represent positions of CIRANO or its partners.

ISSN 1198-8177

Partenaire financier

Développement
économique, Innovation
et Exportation

Québec 

Analysts, Incentives, and Exaggeration^{*}

Timothy Shields[†]

Résumé

Les analystes du « côté vendeur » sont rémunérés, du moins en partie, au moyen de commissions de courtage. Ces commissions représentent une incitation à fausser les prévisions dans le but d'accroître les activités du marché. Ainsi, les analystes font l'objet de stimulants économiques clairs qui les portent à décevoir, tandis que les négociateurs ont des stimulants économiques qui les incitent à détecter la déception. Les théories analytiques avancées antérieurement à partir des jeux de transmission de l'information prédisent, de façon catégorique, qu'il existera toujours une certaine déception (avec transaction) dans le meilleur des cas et des messages dénudés de renseignements (et sans transaction) dans le pire des cas, à moins que les stimulants de l'émetteur et du récepteur ne concordent. La preuve expérimentale qui a été faite dans le passé par les jeux de transmission de l'information a démontré que les émetteurs choisissent effectivement de décevoir, même si leurs messages sont plus informatifs que la théorie laisse entendre. De la même façon, les récepteurs comptent davantage sur les messages que ce qui est proposé par la théorie. Le comportement qui s'écarte des prédictions peut-il s'expliquer par un comportement social normatif, en l'occurrence la confiance et l'honnêteté? Par ailleurs, les sujets font-ils preuve de rationalité limitée et négligent-ils de considérer suffisamment les stimulants des autres joueurs au moment de prédire leurs décisions? Dans le but de répondre à ces questions, je mets au point, puis je conduis une expérience visant à observer les comportements liés aux prévisions et à la négociation et à établir si ces comportements s'expliquent mieux par la théorie analytique, la sophistication stratégique limitée ou les normes sociales. Les résultats de l'expérience confirment que, d'une part, une majorité de sujets adoptent des stratégies prévisionnelles malhonnêtes et, d'autre part, une majorité de sujets adoptent des stratégies commerciales fondées sur la confiance. De plus, les sujets ne semblent pas revoir leur comportement face au marché, malgré l'évidence de prévisions décevantes. Les résultats portent à croire qu'on peut mieux expliquer le comportement des sujets, dans un contexte donné, par la sophistication stratégique limitée que par le comportement social normatif ou la rationalité séquentielle.

Mots clés : prédictions des analystes, niveau de sophistication, norme sociale, rationalité limitée, confiance, honnêteté

^{*} This work is based upon my dissertation. I would like to thank my advisor, John Dickhaut, for the many conversations and suggestions that made this work possible. Thanks are also due to my committee members Frank Gigler, Robert King, Tom Issaevitch, and Ram Venkataraman for their insight and guidance. I have also benefited from the encouraging remarks of Poh Geok Chew, Chandra Kanodia and Jack Stecher. Any errors within are my own.

[†] University of Minnesota - Twin Cities - Carlson School of Management, email: tshields@umn.edu

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 show 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 majority of subjects adopt dishonest 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 limited strategic sophistication than by social normative behavior or sequential rationality.

Keywords: *analyst forecast, levels of sophistication, social norm, bounded rationality, trust, honesty*

1 Introduction

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

Wall Street security analysts garnered attention after the burst of the technology-stock bubble. Analysts were alleged to exaggerate the future prospects of firms in order to secure or maintain investment-banking relationships (Brennan, 2004; Michaely and Womack, 1999). During this time-frame, 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 U.S. Securities and Exchange Commission 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 securities—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.

The incentive created by brokerage commissions is an example of the economic tensions modeled in communication games, sometimes also called cheap talk games (Crawford and Sobel, 1982). In models of communication games, a sender with information sends a message to a receiver, who then takes an action that affects the payoffs of the both the sender and the receiver. If the incentives of the both the sender and receiver are not perfectly aligned—in that the sender would prefer the receiver take an action other than the action that is best for the receiver, then not all information can be credibly communicated. That is, in equilibrium the sender will not adopt a truth-telling strategy and the receiver will not be fully trusting. As a result, the equilibrium will, at best, be partially informative, and at worst, completely uninformative. The equilibrium is constructed assuming both the sender and receiver are sequentially rational. Sequential rationality requires the sender’s actions and beliefs are optimal and correct, respectively, given the receiver’s actions and beliefs.

Several experimental studies examining communication games 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). However, while subjects have shown a willingness to deceive others in strategic transmission games, a phenomenon of over-communication has been documented: senders reveal 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., 2001). Some of this work also postulates that some subjects are non-strategic in that the subjects appear to make limited use other players' incentives to predict their decisions. These subjects' behavior can be described as boundedly rational; instead of behaving as if sequentially rational, the subject adopts a heuristic. Using behavioral type analysis, these authors find better fit to subject behavior than game theoretic predictions for one-shot games (Costa-Gomes and Crawford, 2006; Camerer et al., 2004).

Related to behavioral type analysis are hierarchical models of sophistication, where an agent's beliefs of the others' level of sophistication dictate her actions. Unlike sequentially rational behavior, the agent's beliefs can be off-equilibrium. Stahl and Wilson (1994) characterize behavior over ten one-shot games with a hierarchical model of reasoning, where a level-0 type plays deterministically, a level-1 type acts if others were level-0, a level-2 type acts as if others were level-1 types, and so on. The authors find three fourths of the players were level-1 or level-2; the remaining fit the sequentially rational prediction. Camerer et al. (2004) perform a similar analysis of prior published results of coordination and market-entry games and find, on average, a type of level-1.5 predicts overall behavior.

In this work, I examine behavior in an institutional market setting to distinguish behavior from the prior experiments of communication games conducted in a context neutral setting. Within the stylized setting, I ask what best explains overall subject behavior: sequentially rational models, social norms of honesty and trust, or notions of bounded rationality?

In order to examine this question, 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 of nature. The grouping and game is repeated eight times. In the first treatment the trading costs are set such that all forecasts could induce trade if all players are sequentially rational. This is theoretically possible as expected gains to trade are available even if analysts use deceptive strategies. In the second treatment the trading costs are set so the analyst cannot adopt a forecast strategy that always induces trade if all players are sequentially rational. The game-theoretic prediction is that no forecast will induce trade—despite that half the time there are potential gains to trading.

In this repeated game, can subjects' overall behavior be explained by sequential rationality? If not, is the overall behavior boundedly rational, or can overall behavior be explained by social norms of honesty and trust? If, overall, subjects are boundedly rational, then I posit subjects use non-adaptive heuristics; that is, they play the same strategy repeatedly even if it results in economic detriment. Alternatively, if subjects are motivated 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 that subjects make limited use of others' incentives to predict decisions. A majority of subjects tend to use deceptive forecasting strategies when analysts. 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 over-communication phenomenon documented in the aforementioned communication game studies.

The remainder of the paper is organized as follows: in section 2, I describe the experimental design and introduce a hierarchical model of sophistication applicable to my experimental setting. In section 3, I predict the probability of trade and allocations for each of the three levels of hierarchical model, argue these predictions are consistent with behavior based on (i) social norms, (ii) bounded rationality, and (iii) sequentially rationality, and present competing hypotheses of subjects' behavior. Last, in section 4, I present experimental results and conclude.

2 Experimental Design and Procedures

The experimental design is based upon information transmission games with multiple receivers. Each independent economy consists of one player with private information who 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 with the same three players.

2.1 Trading Setting

This is a three-person sequential move game that consists of three stages illustrated in 1. 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. There is a 90% probability the realized state will equal the private signal. The signal has limited support in that the realized state will be an element of the set $\{A, B\}$ when the signal is either A or B , and an element of the set $\{C, D\}$ when the signal is either C or D (Table 1). The signal structure and all distributions are public knowledge, while the realization of the analyst's signal is private.

In the first stage, the analyst releases a forecast, $f \in \{A, B, C, D\}$. If the realization of the private signal, s , is an element of $\{A, B\}$ the analyst can release a 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 $\{C, D\}$.

<i>Stage 0</i>	<i>Stage 1</i>	<i>Stage 2</i>
Analyst releases forecast.	Market opens for trade.	State realized. All players receive payoffs.

Figure 1: Base Game Stages

Signal Realized	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>
Probability State is <i>A</i>	.9	.1		
Probability State is <i>B</i>	.1	.9		
Probability State is <i>C</i>			.9	.1
Probability State is <i>D</i>			.1	.9

Table 1: Chances of State Given Private Signal

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 , selected from a menu of prices, and pays trading costs λ including a commission paid to the analyst. Trading costs vary across two experiment treatments, hereafter referred to as the low-cost and high-cost treatments.

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 analyst's private signal, and shares are liquidated. The shareholders have different payoffs from the liquidated shares. The payoff for each share held by the shareholder at stage 2 is shown in Table 2.

State	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>
Payoff to Shareholder 1	30	50	70	90
Payoff to Shareholder 2	0	40	80	120

Table 2: Payoffs at Stage 2 Given State

2.2 Behaviors of Analysts and Shareholders

In this section I outline the combinations of behavior assessed to infer the level of strategic sophistication using a framework related to Camerer et al. (2004). Given a typical player in the experiment, there are 3 different levels of sophistication the typical player may have based on their observable behavior in both the role of analyst and shareholder. A typical player may be (a) naïve in both roles, (b) sophisticated in both roles, or (c) sophisticated as an analyst but naïve as a shareholder.

First, I construct the gains to trade and trading strategies given the typical player is naïve in both roles. That is, the analyst is honest and shareholders are trusting. Second, I assume the typical player is sophisticated both as an analyst and shareholder, and construct the gains to trading and

trading strategies given the analyst is dishonest and shareholders are non-trusting. Last, I assume the typical player is sophisticated as an analyst, but naïve as a shareholder, and construct the gains to trade and trading strategies given the analyst is dishonest but shareholders are trusting.

2.2.1 Level-0 Strategic Sophistication: Honest Analyst and Trusting Shareholders

Imagine the typical player employs strategies at the base level of sophistication. That is, the player’s behavior is naïve in both roles. As an analyst the player adopts an honest forecasting strategy and simply reveals her private signal via the forecast. As a trusting shareholder, the player believes the realized state in stage 2 is likely the forecasted value, but recognizes there is some chance the realized state will differ, as the analyst’s private signal was imperfect. Thus, risk-neutral shareholders have the expected payoffs shown in Table 3 after the release of the honest analyst’s forecast in the first stage. For example, if the forecast was A , both shareholders will believe the analyst’s private signal was A , and thus, believe the state is likely A , but realize the state could also be B . The expected payoff is: $\Pr(\text{State is } A \mid \text{signal}=A) (\text{Payoff from State } A) + \Pr(\text{State is } B \mid \text{Signal}=A) (\text{Payoff from State } B)$.

Forecast released	A	B	C	D
Shareholder 1’s Expected Payoff	32	48	72	88
Shareholder 2’s Expected Payoff	4	36	84	116
Difference (Potential Gains to Exchange)	28	12	12	28

Table 3: Expected Payoffs given Honest Forecasting Strategy

Trading costs are 13 and 21 in the low-cost and high-cost treatments respectively. In the setting with honest analysts and trusting shareholders, when shareholders believe the state is likely B or C , they will not trade as the difference in payoffs (12) are smaller than trading costs in either treatment. This is illustrated in Figure 2. Thus, there will not be trade if the analyst’s signal is B or C when players play honestly in the analyst role and trusting in the shareholder role.

2.2.2 Level-2 Strategic Sophistication: Dishonest Analyst and Non-trusting Shareholders

Conventional game theory assumes players are sequentially rational, and presupposes indefinitely iterative reasoning (“I think that you think that I think...”). Imagine a player who iterates at least two steps beyond the aforementioned level-0 behavior. In the first iteration of reasoning, the player adopts a deceptive forecasting strategy in order to increase commissions, as the player correctly assumes shareholders will not trade if the state is likely C or D . With a second iteration of reasoning, the player realizes all others will adopt deceptive forecasting strategies, and thus views all forecasts with skepticism. Consequently, after two or more steps of iterative reasoning, the player is sophisticated in both in the role of analyst and shareholder.

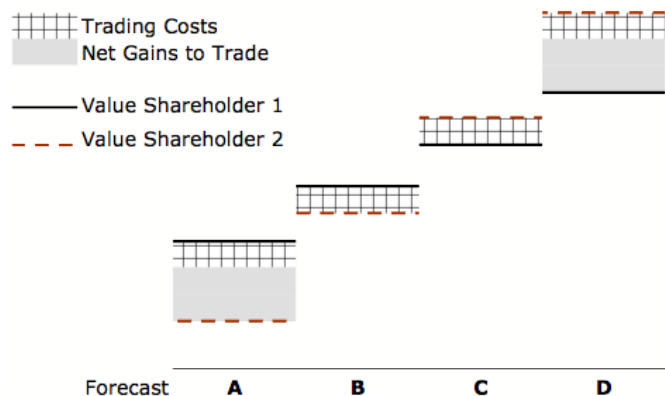


Figure 2: Gains to Trade Given Honest Forecasting Strategy

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*.

Forecast released	<i>A</i> or <i>B</i>	<i>C</i> or <i>D</i>
Shareholder 1's Expected Payoff	40	80
Shareholder 2's Expected Payoff	20	100
Difference (Potential Gains to Exchange)	20	20

Table 4: Expected Payoffs Given Deceptive Forecasting Strategy

As an analyst, the player adopts deceptive forecasting strategies. The same player, in the role of shareholder, anticipates forecasting deception, and after receiving a forecast of *A* or *B*, correctly surmises the state is equally likely to be *A* or *B*. Likewise, if the forecast was *C* or *D*, the state is equally likely to be *C* or *D*. The sophisticated risk-neutral shareholder has the expected payoffs shown in Table 4 after the release of the deceptive forecast.

In the low-cost treatment, shareholders at a level-2 degree of sophistication will trade, as the difference in expected gains to exchange (20) exceeds the trading costs (13). This relationship is illustrated in Figure 3 for the low-cost treatment. However, in the high-cost treatment, the costs of trading (21) exceeded the gains to exchange, and the shareholders will not trade. This suggests that if players are of level-2 or higher degree of strategic sophistication, every forecast will induce trade in low-cost treatment and no forecast will induce trade in the high-cost treatment. These arguments are presented formally in Appendix A.

2.2.3 Level-1 Strategic Sophistication: Dishonest Analyst and Trusting Shareholders

Suppose rather than being a Level-0 or Level 2 player, a typical player employs only one step of reasoning beyond the level-0 behavior. What strategies will this player adopt? In the role of

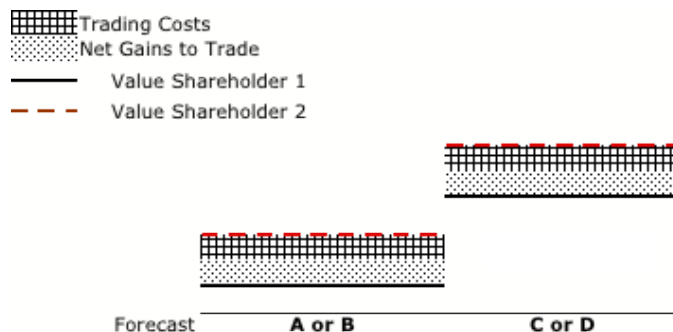


Figure 3: Gains to Trade Given Deceptive Forecasting Strategy

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 than 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.

analyst, she will anticipate that shareholders will not trade if they assess the state is likely B or C , and thus, will prefer to convince the shareholders that the state is likely A or D in order to earn a commission. Therefore, the player adopts a deceptive forecasting strategy.

Unlike the level-2 player, this player fails to consider that all other players will adopt deceptive forecasting strategies. As such, a player of this level of strategic sophistication adopts trusting trading strategies as a shareholder and trades whenever the forecast suggests the state is likely A or D . As a result, the naïve shareholder has the same expected payoffs as in Table 3, and trades for all forecasts released, erroneously expecting the gains to exchange to exceed the cost of trading in both treatments. However, the information content of the forecast is not as fine as the naïve shareholders assume. A forecast of A only suggests the state is equally likely to be A or B , and a forecast of D suggest the state is equally likely to be C or D .

2.3 Eliciting Subjects' Decisions

Within the experiment, an economy was defined as a grouping of one analyst and two shareholders. During each round all subjects within the economy played the base game. The economy remained intact for a set of eight rounds. Each economy was independent from another in the laboratory, since the results from one economy did not affect payoffs in other economies, and the information released within an economy was not available to subjects outside the economy. Each subject in the economy only witnessed information pertaining to his own economy.

Subjects who played 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 reveal her signal truthfully or attempt to deceive the shareholders. See Figure 5 in Appendix C.

Meanwhile subjects playing the role of shareholders entered their trading decision. Four trading prices were presented as per Table 5. If the shareholders believed the analyst was forecasting honestly, then P1 and P4 were the feasible prices when the forecast was A and D , respectively. In the low-cost treatment, if the shareholders believed the analyst was forecasting deceptively, then P2 and P3 were the feasible prices when the forecast was an element of $\{A, B\}$ or $\{C, D\}$, respectively. In the high-cost treatment, no price was feasible when shareholders believe the analyst forecasts deceptively.

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

Table 5: Prices for Both Treatments

For each possible forecast value two trading prices were made available to shareholders. For forecast values of A and B , prices P1 and P2 were presented, and for forecast values of C and D prices P3 and P4 were 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 and D . Shareholder 2 was symmetrically restricted.

Akin to the analyst, each 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 shareholders elected to trade at the high price, the low price, or elected not to trade. See Figures 6 and 7 in Appendix C.

2.3.1 Trading and Reporting History

After subjects entered their decisions, 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 used. 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 8 in Appendix C). This summary table displayed the realizations for all rounds played within the economy. Subjects were able to scroll down and see realizations for earlier rounds within the set.

Using the summary report of the ten realizations, a subject could partially infer other subjects' strategies 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 number of times the forecast released was not identical to the realized state. Knowing the analyst's signal structure (shown in Table 1), shareholders could, in theory, calculate the likelihood an analyst used an honest or deceptive forecast strategy given the probabilities shown in Table 6. If there is more than one mismatch of the forecast and realized state, it was more likely the analyst used a deceptive forecasting strategy.

Number of Times Forecast and State Not Equal	Probability of this Number of Times Given Analyst is Using a:	
	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%

Table 6: Probability of Forecast and State Inequality Given Forecast Strategies

2.4 Experimental Procedures

The experiments were conducted in Montreal, Canada by the Centre for Interuniversity Research and Analysis of Organizations (CIRANO) over two sessions during April and May 2007. Subjects were recruited by CIRANO from a standard subject pool and remain anonymous to the author. Subjects 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 subjects 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 subject followed along with their own copy of the instructions (available from author on request). The instructions explained the experimental procedures and the information structures used in the experiment. While going over the instructions, subjects were asked to write down their answers to several questions to ensure that they understood the instructions. Subjects' answers remained confidential.

The experimenter reviewed the correct answers. During and after the instructions were read, subjects were prompted to ask the experimenter any questions regarding the experiment procedures.

2. Each subject was randomly and anonymously grouped into an independent economy with two other subjects. The identities of members within the economy were not revealed to any subject. An analyst was randomly selected within each economy 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 state. Meanwhile, the shareholders entered their trading strategies for each possible value of the analyst's forecast. Each subject 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 analyst's private signal and nature for each economy. A summary screen showed each subject the results, including their own payoff, of these ten realizations.
5. The preceding two steps constitute a round. Each economy of subjects played eight rounds.
6. Subjects were regrouped into new economies, as described above, and played another set of eight rounds. The subjects played eight sets in total.
7. Each subject 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.

3 Theories of Behavior

In this analysis, I use a derivative of the cognitive hierarchy of Camerer et al. (2004) to frame the competing theories of behavior predicting overall subject behavior. The resulting hierarchical model of sophistication, applied to the base game described in section 2, supports three levels of sophistication, each dictating observable behavior and metrics that are consistent with (i) the influence of social norms of honesty and trust, (ii) bounded rationality, and (iii) sequential rationality, respectively.

3.1 Hierarchical Model of Sophistication

Behavioral type analysis is an approach used by several authors (Costa-Gomes et al., 2001; Stahl and Wilson, 1994; Camerer et al., 2004) to predict how subjects will behave. To apply the behavioral

type approach to communication games, Crawford (2003) cites early experiment evidence (Blume et al., 1998) and argues the system of types should be anchored on the honest type of sender (analyst) and the trusting type of receiver (shareholder). In this work, I also anchor my hierarchy upon norms of honesty and trust.

Entertaining the proposition that subjects may not possess indefinitely iterative reasoning yields a richer set of predictions than models assuming sequentially rational agents. Table 7 reports the predicted results if overall subjects' is analogous to one of the above levels of sophistication. The predictions are based on the experimental parameters discussed in Section 2.1. The calculations for the predictions contained in Table 7 are discussed thereafter.

	<u>Likelihood of Trade</u>		<u>Likelihood of Capturing Gain</u>		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
Deceptive Analyst and Trusting Shareholders (Level-1)	Dependent Upon Forecasting Consistency				Varies	Greater for Shareholder 1
Deceptive Analyst and Skeptical Shareholders (Level-2)	100%	0	100%	0	Varies	Equal Between Shareholders

Table 7: Predictions for Levels of Strategic Sophistication

3.1.1 Honest Analyst and Trusting Shareholders

Within the hierarchical framework, the combination of honest analysts 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 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 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. 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.

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 8.

	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

Table 8: Expected Gains to Trading Given Honest Reporting Strategy

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 first and last combinations of types shown in Table 7. 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 accordingly.

3.1.2 Deceptive Analysts and Trusting Shareholders

A combination of a deceptive analyst and trusting shareholders is equivalent to level-1 of the hierarchical framework. When the analyst is deceptive, she uses one of the forecast strategies described in section A.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.

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 a repeated game. In particular, do trusting types alter their beliefs in light of conflicting information, or do subjects fail to 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 trading strategies independent of the analyst's strategy.

3.1.3 Deceptive Analysts and Skeptical Shareholders

The last combination mirrors the predictions of the analytical model in Appendix A. The allocations are equivalent to level-2 of the hierarchical framework. 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 9

	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

Table 9: Expected Gains to Trading Given Deceptive Reporting Strategy

If subjects are sequentially rational, then equilibria behavior is found by backwards induction. Subjects will act as if they have reasoned backwards in time from last round in the finite set, and appropriately calculated the best course of action. Accordingly, all analysts will adopt deceptive forecasts, and all shareholders, believing analysts are deceptive, will be skeptical of the forecasts. These predictions are derived formally in Appendix A.

3.2 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 7.

3.2.1 Honest Analyst and Trusting Shareholders

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 row of Table 7. 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 trade is high and equal over treatments. Analysts' forecasts will be consistent, and the benefit of trading will be equal and positive for both shareholders.*

3.2.2 Deceptive Analysts and Trusting 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 second row of Table 7. 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.

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

3.2.3 Deceptive Analysts and Skeptical Shareholders

Last, if most subjects are able to employ more than one step of strategic reasoning, then allocations will be similar to the last row of Table 7. 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 sequentially rational, then the likelihood of 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.*

4 Findings and Conclusion

Within this section, I first describe the data sets analyzed, the methods used to classify subjects' behavior in each round of a set, and the method for classifying behavior over a set. Second, I report results that support that overall behavior is not consistent with theories of social norms or sequential rationality, but is best explained by theories of bounded rationality. Last, I conclude, discuss policy implications, and posit topics for further study.

4.1 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 Centre 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 2.1. Twenty-four subjects were randomly assigned to the two treatments from a subject pool consisting primarily of university students. The subjects' median age was 25 years, the youngest 20, and eldest 34.

Each treatment, low-cost and high-cost, was conducted with 12 subjects, each of whom played 8 sets. Each set had an analyst and two shareholders, randomly assigned to roles at the beginning of a set. Each set had 8 rounds of play. As such, there were 32 independent economies (4 per set) 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.

4.2 Strategic Behavior Classification

For each round, subjects were classified into one of three types based upon their input forecasting strategy or trading strategy. The forecasting strategy was classified as honest, deceptive, or partially deceptive (see section 4.2.1). The trading strategy was classified as trusting, skeptic, or non-trading (4.2.2). Based upon the eight round-based classifications, subjects were classified into three behavioral types for the every set.

4.2.1 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 choose 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 chose one forecast value for both signals C and D (A or B).

4.2.2 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.

¹This 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 capture subjects whom believed the forecast, rather than traded erratically.

The frequencies of the subjects' behavior in each round are shown in Tables 10 and 11. In both treatments analysts tended to display deceptive or partially deceptive forecasting behaviors. In both treatments, shareholders tended to display trusting behavior.

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

Table 10: Frequency of Round Behaviors: Low-Cost Treatment

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

Table 11: Frequency of Round Behaviors: High-Cost Treatment

4.2.3 Analyst Set Behaviors

Given the eight round strategies, a subject's forecasting set behavior was classified using a simple majority replicating the methods of Cai and Wang (2006) and Costa-Gomes et al. (2001). Borrowing from this work, the majority was defined as 5 out of 8 rounds.

Honest The subject used an honest round strategy a majority of the time.

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.

Inconstant 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.

The classified set behavior for each subject is shown in Table 12. As roles of analyst and shareholders were 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 non-honest (deceptive or inconstant). Most subjects had behavior that was classified into one type a majority of the time. The classification method yielded 5 subjects who were honest a majority of the sets². Only 1 subject used an honest forecasting strategy in every round of every set.

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

Subject	<u>Low-Cost Treatment</u>			Subject	<u>High-Cost Treatment</u>		
	Honest	Deceptive	Inconstant		Honest	Deceptive	Inconstant
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 12: Analyst Set Behavior by Subject
Prevalent strategy shown in bold.

4.2.4 Shareholder Set Behaviors

Similar to the classification of forecasting behavior, each subject's shareholder behavior was classified using 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. The average of the eight rounds' scores is used to determine the set behavior.

Trusting The subject used a trusting round strategy a majority of the time.

Skeptic The subject used a skeptic and non-trading round strategy a majority of the time.

Inconstant The subject used trusting and skeptic strategies an equal number of times.

The shareholder set behavior for all subjects are shown in Table 13. A majority of the subjects' set behavior (66%) was classified as trusting. Most subjects had trading behavior that was classified into one type a majority of the time. The classification yielded 17 subjects who where trusting in a majority of the sets³. Only two subjects used trusting strategies in all rounds of all sets.

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

Subject	<u>Low-Cost Treatment</u>			Subject	<u>High-Cost Treatment</u>		
	Honest	Deceptive	Inconstant		Honest	Deceptive	Inconstant
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 13: Shareholder Set Behavior by Subject
Prevalent strategy shown in bold

4.3 Subject Overall Behavior

Table 14 reports each subject's set behaviors classified using the majority of round behavior described in Section 4.2.3 and Section 4.2.4. In every row within the table, each subject's behavior in sets 1 through 8 is classified as either honest, deceptive or inconstant when playing the role of analyst, and trusting, skeptical, or inconstant 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.

Subject	Reporter			Shareholder 1			Shareholder 2		
	Honest	Deceptive	Inconstant	Trusting	Skeptical	Inconstant	Trusting	Skeptical	Inconstant
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 14: 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 hierarchal framework discussed in Section 3.1, 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. Examining 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). 4 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 the hierarchal framework; the remaining subject was honest as a reporter and skeptical as a shareholder. These results are reported in the first column of Table 15.

To gain insight as to whether subjects' behavior became more sophisticated over time, each subject's behavior was alternatively classified using only select sets within the experiment. Specifically, given the majority of analysts employ deceptive forecasting strategies, are shareholders becoming more skeptical over the course of the experiment? In the second column of Table 15, 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 Table 15, 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 the hierarchal framework using these modifications. Again, the remaining subject was honest as a reporter and skeptical as a shareholder. Using this later classification, 3 subjects' shareholder behavior changed from being classified as trusting to inconstant/skeptical. 2 of these 3 subjects used trusting strategies and inconstant/skeptical strategies an equal amount of sets. The other had only one observation of shareholder set behavior classified as an inconstant (using trusting and skeptical strategies an equal number of rounds).

Behavior Classification	Using All Sets	Using Sets 3 through 8	Using Post-Analyst Sets
Honest Analyst & Trusting Shareholder	4	3	4
Honest/Deceptive Analyst & Trusting Shareholder	3	1	1
Deceptive Analyst & Trusting Shareholder	10	12	9
Deceptive Analyst & Trusting/Skeptical Shareholder	2	1	3
Deceptive Analyst & Skeptic Shareholder	4	6	6

Table 15: 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 eight sets are used. In the last column, shareholder behavior was classified only using sets after the subject played the role of analyst. The number of subjects fitting each classification is reported. 23 of 24 subjects fit into the listed categories. The remaining subject was honest as an analyst but skeptical as a shareholder.

To determine whether playing the role of analyst altered shareholder strategies, each subject's

trading behavior was classified into two categories: trading behavior before playing the role of analyst, and trading behavior after. There were 15 out of 24 subjects for which there observed trading before and after playing the role of analyst. The average (median) proportion of which a skeptical strategy was used was 28.33% (0.0%) and 35.56% (25.0%) for sets observed before playing the role of analyst and after playing the role of analyst, respectively. These differences are statistically insignificant using the Wilcoxon Matched-Pairs Signed-Ranks Test ($W^+ = 25.50$, $W^- = 29.50$, $N = 10$, $p \leq 0.8457$).

4.4 Experimental Results

In this section, the predicted behavior of subject, metrics and allocations listed in Table 7 are analyzed. After reviewing the results, the surviving hypothesis is isolated from the competing hypotheses listed in Section 3.2.

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

Analysts adopted deceptive or inconstant forecasting set strategies 70% of the time. 15 subjects tended to adopt deceptive and/or inconstant forecasting strategies and 5 subjects tended to use honest forecast strategies. These results are consistent with prior work that documents 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 3.1 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 honest forecasting strategies ($X^2 = 7.04$, $DF = 1$, $p \leq 0.0079709$).

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

Shareholders adopted trusting trading strategies 66% of the time. 5 of 24 subjects tended to adopt skeptical and/or inconstant strategies and 17 subjects tended to adopt trusting trading strategies. The remaining subjects adopted trusting and skeptical/inconstant strategies an equal number of sets.

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 strategies. Using a Chi-square test for known distributions, I reject the null that half the shareholders used trusting strategies ($X^2 = 5.50$, $DF = 1$, $p \leq 0.019016$).

Result 3. *The likelihood of trade was greater for the high-cost treatment than the low-cost treatment*⁴.

To further examine the extent of the support for the conclusion that overall observed behavior is captured by a level-1 degree of strategic sophistication (deceptive and trusting), the difference between expected trade in the high-cost and low-cost treatments was computed for each hypothesis. Then the data were used to determine the likelihood of trade in the data. (Note because there is state uncertainty, the likelihood of trade will be a property subjects' input decisions—see Appendix B.1 for formulation). The data were only consistent with the hypothesis that subjects were deceptive as analysts and trusting as shareholders.

The mean likelihood of trade was 25.9% and 31.5% in the low-cost and high-cost treatments, respectively (see Tables 18 and 19 in Appendix B.5 for inter-quartile ranges). If overall subject behavior was best captured by a level-0 degree of sophistication, (honest and trusting), the likelihood of trade would be equal across treatments. I reject the null hypotheses that the medians are equal over the two treatments using a Wilcoxon Two Sample Test ($W = 61962.5$, $p \leq 0.02703$).

If overall subject behavior were captured by a level-2 or higher level of sophistication (deceptive and skeptical), then the likelihood of trade 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 likelihood being greater in the high-cost treatment than in the low-cost treatment, the prior Wilcoxon Two Sample Test is sufficient to reject the null hypothesis that the likelihood is 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 was correlated with analyst forecast consistency.*

Within hierarchical framework, forecast consistency would only affect trade when shareholders are trusting and analysts are deceptive. After two iterations of reasoning, shareholders are skeptical, understanding there is limited information content in the analyst's forecast, and thus forecast consistency would not affect the likelihood of trade.

The forecasting consistency metric captures how consistently the analyst maps private signals into a released forecast over the set of eight rounds. The metric depends upon the subject's input

⁴The likelihood of capturing gains to trade captures to what extent trade occurred when the state was A or D (see Appendix B.2). The likelihood of capturing gains is highly correlated with the likelihood of trade. The mean likelihood of capturing gains to trade was 29.1% and 34.8% in the low-cost and high-cost treatment respectively. Using a Wilcoxon Two Sample Test, I reject the null hypotheses the medians were equal over the two treatments ($W = 62010.5$, $p \leq 0.02908$).

forecasting decisions (see Appendix B.4 for formulation and Tables 18 and 19 in Appendix B.5 for inter-quartile ranges).

Table 4.4 reports the rank correlation coefficient of the constructed consistency metric against the likelihood 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 3, 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 examining honest analysts in isolation.

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, $W+ = 192.50$, $W- = 17.50$, $N = 20$, $p \leq 0.0003948$).

	Likelihood of Trade
Both Treatments (64 observations)	0.3368 $p \leq 0.007505$
Deceptive Analysts in Both Treatments (47 observations)	0.5572 $p \leq 0.000026$
Deceptive Analysts in low-cost treatment (22 observations)	0.5802 $p \leq 0.0046$
Deceptive Analysts in High-cost Treatment (25 observations)	0.5572 $p \leq 0.003793$

Table 16: Correlation Between Consistency and Trade
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.*

Within the hierarchical framework, trading benefits shareholders unequally only when analysts are deceptive and shareholders are trusting. While trade in states A or D would benefit the economy as a whole, it is possible the trade benefited one shareholder at the expense of another. The benefit of trade is calculated for each shareholder in an economy given the subjects' input forecasting and trading decisions (see Appendix B.3 for formulation). It captures to what extent the shareholder was better off electing to trade rather than keeping his endowment.

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 (see Tables 18 and 19 in Appendix B.5).

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 the differences between shareholders' gains in each round has a median value of zero using the Wilcoxon Matched-Pairs Signed-Ranks Test (low-cost treatment: $W_+ = 402$, $W_- = 63$, $N = 30$, $p \leq 0.000509$; high-cost treatment: $W_+ = 441$, $W_- = 87$, $N = 32$, $p \leq 0.0009657$).

Overall trade benefited shareholder 1 more so than shareholder 2. This resulted from trade at prices P1 and P4 when the analyst used a deceptive forecast strategy—shareholder 2 sold for too little and bought for too much. The rank correlation coefficient between the shareholders' average benefits of trading in economy is negative and significant in both treatments (low-cost treatment: $R = -0.9654$, $N = 32$, $p \leq 7.68e^{-8}$; high-cost treatment: $R = -0.9764$, $N = 32$, $p \leq 5.442e^{-8}$).

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, conditional upon trusting in the first round, I coded the subjects' subsequent set-strategy. If a subject used trusting strategies in a majority of the subsequent rounds, the subject 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 17. The sorting of analysts separated deceptive analysts who adopted honest and partially deceptive strategies an equal number of times within a set. These analysts were labeled as mimics in this analysis.

Table 17: Subsequent Strategies of Trusting Shareholders

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

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.

4.5 Hypotheses Selection

Taken collectively, the results support Hypothesis 2 from the competing hypotheses in Section 3.2. Subjects' overall behavior is best described as boundedly rational, as most subjects chose deceptive

forecasting strategies, but in turn, chose trusting trading strategies. This behavior is consistent with some level, 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 of the likelihood of trade with forecasting consistency is only consistent 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 only consistent with a limited level of strategic sophistication.

4.6 Conclusion

Departing from the homogenous reporting strategies predicted in sender-receiver games, Crawford (2003) posits the existence of honest sender types can induce behavior by strategic agents that is, at times, trusting. 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 a trading institution. 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 conclude that normative social behavior is driving the overcommunication phenomenon. Overall, subjects' behavior is more analogous to boundedly rational behavior within hierarchical models of sophistication (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.

This study suggests that brokerage commissions do create potential conflicts of interests between analysts and shareholders. Almost all subjects adopted deceptive forecasting strategies at one time or another, and the majority adopted deceptive forecasting strategies a majority of the time. Furthermore, disclosure of the conflicts of interest between analysts and shareholders was not sufficient to evoke shareholder skepticism. Again a majority of shareholders relied upon the forecasts more than warranted. Additionally, shareholders failed to consistently detect deception in analysts' forecasts. Taken together, the results suggest that subjects are inclined to trust forecasts despite deception and knowledge of potential conflicts of interests. If shareholders cannot discipline the analysts' forecast strategy with legitimate skepticism, then analysts have no incentive to forecast honestly. Thus, just as the SEC barred brokerage firms from tying analyst's compensation to investment banking transactions, the SEC might also consider eliminating brokerage commissions.

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., 2001; Costa-Gomes and Crawford, 2006). 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 to trading foregone in this study? Additionally, Wall Street analysts may have incentives to rank highly in one of the various financial presses' listings of 'Star Analysts', and thus issue accurate forecasts. Within this my institutional setting, would the introduction of published accuracy metric alter analysts' forecasting strategies or increase shareholder skepticism?

References

- Axelrod, R. (1981). The evolution of cooperation. *Science*, 211(4489):1390–1396.
- 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.
- 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. (2006). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review*, pages 1737–1768.
- Costa-Gomes, M., Crawford, V. P., and Bruno, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, pages 1193–1235.
- 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.
- Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Economic Experiments*, 10(2):171–178. forthcoming *Experimental Economics*.

- Michaely, R. and Womack, K. (1999). Conflict of interest and the credibility of underwriter analyst recommendations. *Review of Financial Studies*, 12(4):653–686.
- Myerson, R. (1979). Incentive compatibility 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.
- 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.
- 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.

A Sequentially Rational Analytical Model

The forecasting and trading game is similar, but not identical to, games of strategic information transmission. In this model, a commission is paid based on trading volume. The commission incentive 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 finitely repeated game, solving for the forecasting strategy of the analyst assuming all agents are sequentially rational. The usual 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 were ran.

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

A.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 that 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 $\Theta = \{\theta_1, \dots, \theta_n\}$, is distributed symmetrically around mean μ , and has probability $g(\theta)$ 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 $S = \{\varsigma_1, \dots, \varsigma_m\}$ where $m \geq n$. The signal is generated through a joint probability $f(\varsigma, \theta)$ satisfying first order stochastic dominance in the sense that $\sum_{k=1}^j f(\varsigma_k, \theta) \geq \sum_{k=1}^j f(\varsigma_k, \hat{\theta}) \quad \forall j \in \{1, \dots, m\}, \hat{\theta} > \theta$. 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.

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 faces different incremental tax rates. Specifically, I assume each shareholder incurs a tax expense (benefit) from the realization of θ if the realization is greater (less) than μ . I assume the tax rates are constant for all realizations of the

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

Figure 4: Base Game Timeline

liquidation value, even if a single owner owns the entire firm, and that at date 2 the taxable basis for all shares is μ^5 .

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

$$v_i(\theta) = (\theta - \mu)\beta_i + \mu \quad (1)$$

Without loss of generality, I normalize $\beta_1^{t=1} = \beta_2^{t=1} = \beta_2^{t=2} = 1$. Let $\beta_1^{t=2} \equiv \beta < 1$. 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, all costs of trading λ , both dead weight and the analyst's commission, are borne by the owner of the firm at date 2. 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

$$v_1(\theta) - \lambda > v_2(\theta) \Rightarrow \beta(\theta - \mu) + \mu - \lambda > \theta \Rightarrow \theta < \mu - \frac{\lambda}{1 - \beta} \quad (2)$$

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

$$v_2(\theta) - \lambda > v_1(\theta) \Rightarrow \theta - \lambda > \beta(\theta - \mu) + \mu \Rightarrow \theta > \mu + \frac{\lambda}{1 - \beta} \quad (3)$$

Defining the critical values of θ from equations (2) and (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 $\theta_1 < \theta_L < \theta_H < \theta_n$.

⁵The results are essentially unchanged (either every forecast induce trade, or no forecast induces trade) if the taxable basis in acquired shares is the price paid.

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.

A.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 a sequential reporting-bargaining game, or mechanism, to determine (i) what the analyst's report will reveal about the state of nature θ , (ii) whether there should be a redistribution of shares between the shareholders and, (iii) 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 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 4.

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. 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 the selling shareholder
by the buying shareholder satisfying budget balancing,
- $\phi(\hat{\theta})$ The probability of trade, and
- $\rho(\hat{\theta})$ 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 $I_m \equiv \{T(\hat{\theta}), \phi(\hat{\theta}), \rho(\hat{\theta})\}$ at date 0 to the shareholders. The shareholders do not see the analyst's private information ς , or her report $\hat{\theta}$, but only the function values in I_m . Recall that at date 1 the shareholders can commit to the mechanism, or choose to walk away from it (see figure 4).

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 $\hat{\theta}$.

$$v_1(\hat{\theta}) + \phi(\hat{\theta}) \left[(1 - \rho(\hat{\theta}))(v_1(\hat{\theta}) - \lambda) - \rho(\hat{\theta})v_1(\hat{\theta}) \right] + (2\rho(\hat{\theta}) - 1)T(\hat{\theta}) \quad (4)$$

$$v_2(\hat{\theta}) + \phi(\hat{\theta}) \left[\rho(\hat{\theta})(v_2(\hat{\theta}) - \lambda) - (1 - \rho(\hat{\theta}))v_2(\hat{\theta}) \right] + (1 - 2\rho(\hat{\theta}))T(\hat{\theta}) \quad (5)$$

The transfer payment in equations (4) and (5) is not multiplied by the probability of trade. If the transfer price is 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.

$$\begin{aligned} \underset{\phi(\hat{\theta}), \rho(\hat{\theta})}{Max} \quad & (1 - \phi(\hat{\theta})) \left[v_1(\hat{\theta}) + v_2(\hat{\theta}) \right] + \\ & \phi(\hat{\theta}) \left[\rho(\hat{\theta})2v_2(\hat{\theta}) + (1 - \rho(\hat{\theta}))2v_1(\hat{\theta}) - \lambda \right] \end{aligned} \quad (6)$$

subject to:

$$\text{IC} \quad \phi(\hat{\theta})\kappa \geq \phi(\tilde{\theta})\kappa \quad \forall \hat{\theta}, \tilde{\theta} \quad (7)$$

$$\begin{aligned} \text{IR 1} \quad & \phi(\hat{\theta}) \left[(1 - \rho(\hat{\theta}))(E[v_1(\theta)|I_m] - \lambda) - \rho(\hat{\theta})E[v_1(\theta)|I_m] \right] + \\ & (2\rho(\hat{\theta}) - 1)T(\hat{\theta}) \geq 0 \end{aligned} \quad (8)$$

$$\begin{aligned} \text{IR 2} \quad & \phi(\hat{\theta}) \left[\rho(\hat{\theta})(E[v_2(\theta)|I_m] - \lambda) - (1 - \rho(\hat{\theta}))E[v_2(\theta)|I_m] \right] + \\ & (1 - 2\rho(\hat{\theta}))T(\hat{\theta}) \geq 0 \end{aligned} \quad (9)$$

The programming problem (6) is the sum of the expected shareholders' utilities (4) and (5). The incentive compatible (IC) constraint (7) pertains to the analyst, who is paid only if there is trade. The individually rationality (IR) constraints (8) and (9) are the shareholders' expected net utilities if they choose to commit to the mechanism after learning the mechanism's triplet I_m .

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 $T(\hat{\theta})$.

A.2.1 Incentive Compatibility For The Analyst

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

$$\text{IC} \quad \phi(\hat{\theta})\kappa \geq \phi(\tilde{\theta})\kappa \quad \forall \hat{\theta}, \tilde{\theta} \quad (10)$$

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 (6), the programming problem is linear in ϕ , so ϕ will take on a value of either 1 or 0. \square

If the designer wanted to implement the first-best solution, 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 (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 (6) is maximized by*

$$\rho(\hat{\theta}) = \begin{cases} 1 & \text{if } \hat{\theta} > \mu \\ \varrho & \text{if } \hat{\theta} = \mu \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

where $0 \leq \varrho \leq 1$

Proof. From equation (1) the shareholders value realizations of θ differently at all values other than μ . Given trade, the programming problem (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. \square

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 more⁶. While the shareholders do not see $\hat{\theta}$, they can infer $\hat{\theta}$ is greater than (less than) μ when $\rho(\hat{\theta})$ is 1 (0). However, 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 (11), ϕ can be determined ex-ante. The mechanism designer determines if gains to unconditional trading exceed the costs of trading. ϕ maximizes (6) when

$$\phi = \begin{cases} 1 & \text{if } E[v_1(\theta) + v_2(\theta)] < \\ & E\left[2v_1(\theta)|\hat{\theta} < \mu\right] (G(\mu) - \frac{g(\mu)}{2}) + E\left[2v_2(\theta)|\hat{\theta} \geq \mu\right] (1 - G(\mu) + \frac{g(\mu)}{2}) - \lambda \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

⁶Consider 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 $2^*\theta$. In this 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 the 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 (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.

A.2.2 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 (11), I can rewrite equations (8) and (9) conditional upon the value of $\rho(\hat{\theta})$

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

$$T(\hat{\theta}) \geq E \left[v_2(\theta) | \hat{\theta} \rho(\hat{\theta}) = 0 \right] \quad (13)$$

$$E \left[v_1(\theta) | \rho(\hat{\theta}) = 0 \right] - \lambda \geq T(\hat{\theta}) \quad (14)$$

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

$$T(\hat{\theta}) < E \left[v_2(\theta) | \rho(\hat{\theta}) = 1 \right] - \lambda \quad (15)$$

$$E \left[v_1(\theta) | \rho(\hat{\theta}) = 1 \right] < T(\hat{\theta}) \quad (16)$$

Combining (13) and (14) yields

$$E \left[v_1(\theta) | \rho(\hat{\theta}) = 0 \right] - \lambda \geq T(\hat{\theta}) \geq E \left[v_2(\theta) | \rho(\hat{\theta}) = 0 \right] \quad (17)$$

Analogously, combining (15) and (16) yields

$$E \left[v_1(\theta) | \rho(\hat{\theta}) = 1 \right] \leq T(\hat{\theta}) \leq E \left[v_2(\theta) | \rho(\hat{\theta}) = 1 \right] - \lambda \quad (18)$$

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 satisfied and the shareholders will participate when directed to trade.*

Proof. Equation (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 (18) is true, and again there exists a transfer

price agreeable to both shareholders. Together (17) and (18) are sufficient conditions for ϕ to be 1 as per equation (12). As ϕ is constant, the analyst's IC requirement is satisfied. \square

Since θ is distributed symmetrically around μ , either (17) and (18) are true, or neither (17) or (18) is true⁷. Notice that if neither (17) or (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 (17) or (18) are true, then there will be no trade.*

A.2.3 Characterization Of The Transfer Price

If neither (17) nor (18) is true, then the mechanism dictates that the shareholders should not trade ($\phi = 0$). The IR constraints (8) and (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 to zero, which is accommodated by multiplying the transfer price by ϕ . All subsequent discussion in this section assumes that (17) and (18) are true, and thus $\phi = 1$.

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 (17) and (18) are true. The transfer price in this case can be characterized by

$$\alpha\{E[v_1(\theta)|\rho(\hat{\theta})] - (1 - \rho(\hat{\theta}))\lambda\} + (1 - \alpha)\{E[v_2(\theta)|\rho(\hat{\theta})] - \rho(\hat{\theta})\lambda\} \quad (19)$$

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 (19) is the coarsest partition.

Alternatively, a finer partition might be possible given equations (17) and (18) are true. Let $\underline{\theta} = \sup\{\theta \in \Theta : E[\theta|\underline{\theta} \leq \hat{\theta} \leq \mu] \leq \theta_L\}$. Note that by equation (17) $\underline{\theta}$ exists, although it may

⁷If θ is not symmetric around μ , it may be the case the only (17) or (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 $E[2v_1(\theta)|\hat{\theta} \leq \tilde{c}]G(\tilde{c}) + E[2v_2(\theta)|\hat{\theta} \geq \tilde{c}](1 - G(\tilde{c})) - \lambda > E[v_1(\theta) + v_2(\theta)]$, and the rationality constraints, $E[v_1(\theta)|\hat{\theta} \leq \tilde{c}] - \lambda \geq T(\hat{\theta}) \geq E[v_2(\theta)|\hat{\theta} \leq \tilde{c}]$, and $E[v_1(\theta)|\hat{\theta} \geq \tilde{c}] \leq T(\hat{\theta}) \leq E[v_2(\theta)|\hat{\theta} \geq \tilde{c}] - \lambda$ are true. Even in the case when θ is symmetric around μ , a different cut-off can be supported as an equilibrium.

be that $\underline{\theta} = \theta_1$. Denote the shareholder expectation in this region as $v_* \equiv E[\theta | \underline{\theta} \leq \hat{\theta} \leq \mu]$. Symmetrically let $\bar{\theta} = \inf\{\theta \in \Theta : E[\theta | \mu \leq \hat{\theta} \leq \bar{\theta}] \geq \theta_H\}$, and if equation (18) is true, $\exists \bar{\theta} \leq \theta_n$. Let v^* denote the shareholder expectation of θ in this region: $v^* \equiv E[\theta | \mu < \hat{\theta} \leq \bar{\theta}]$. Consider the transfer price

$$T(\hat{\theta}) = \begin{cases} \alpha\{v_1(\hat{\theta}) - \lambda\} + (1 - \alpha)\{v_2(\hat{\theta})\} & \text{if } \hat{\theta} < \underline{\theta} \\ \alpha\{v_1(v_*) - \lambda\} + (1 - \alpha)\{v_2(v_*)\} & \text{if } \underline{\theta} \leq \hat{\theta} \leq \mu \text{ and } \rho(\hat{\theta}) = 0 \\ \alpha\{v_1(v^*)\} + (1 - \alpha)\{v_2(v^*) - \lambda\} & \text{if } \mu \leq \hat{\theta} \leq \bar{\theta} \text{ and } \rho(\hat{\theta}) = 1 \\ \alpha\{v_1(\hat{\theta})\} + (1 - \alpha)\{v_2(\hat{\theta}) - \lambda\} & \text{if } \hat{\theta} > \bar{\theta} \end{cases} \quad (20)$$

For values of $\hat{\theta}$ in a neighborhood of μ , specifically $\underline{\theta} \leq \hat{\theta} \leq \bar{\theta}$, the transfer price given by (20) is similar to the price given by (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.

A.3 Allocations In the Original Game

Given the equilibrium allocations found in the last section, I examine the game shown in Figure 4 without a direct-revelation mechanism. The analyst will form a reporting strategy that yields the same allocations as the mechanism. Instead of being directed to trade, the shareholders will jointly decide to trade, or not. If there is trade, the shareholders will also determine whom will be the seller and negotiate a transfer price.

A.3.1 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 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. \square

Akin to the mechanism, let $\hat{\theta} \equiv E[\theta|\zeta]$, 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 than θ_L or greater than θ_H .

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

$$E[v_1(\theta)|f] - \lambda \geq E[v_2(\theta)|f], \quad \text{or} \quad (21)$$

$$E[v_1(\theta)|f] \leq E[v_2(\theta)|f] - \lambda \quad (22)$$

Proof. If the forecast f neither satisfies equations (21) or (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. \square

As in the mechanism, if equations (17) and (18) are true, then a reporting strategy that always induces trade is feasible. Consider a simple trigger strategy where $\hat{R}: \zeta \mapsto \{\hat{f}, \check{f}\}$. Let $\zeta = \sup\{\zeta : \hat{\theta} \leq \mu\}$ be the cutoff signal such that $\forall \zeta \leq \zeta$, the analyst releases the forecast \hat{f} , else \check{f} is released. Given this reporting strategy, both shareholders would be willing to trade. When the forecast \hat{f} is released shareholder 2 would sell, and when \check{f} is released shareholder 1 would sell. The transfer price is characterized by equation (19).

A reporting strategy need not be as coarse as the aforementioned simple trigger strategy to induce trade. Again, assume equations (17) and (18) are true, then the following reporting strategy will always induce trade.

$$R^*(\zeta) = \begin{cases} v_* & \text{if } \underline{\theta} \leq \hat{\theta} \leq \mu \\ v^* & \text{if } \mu < \hat{\theta} \leq \bar{\theta} \\ \hat{\theta} & \text{otherwise} \end{cases} \quad (23)$$

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 \hat{R} is the coarsest. In designing the mechanism, the direction of trade was dictated by (11). Clearly this method of determining the direction of trade is preferred by all agents given equations (17) and (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 A.1. Let R' be a simple trigger strategy, $R': \zeta \mapsto \{f', f''\}$ where there is a cutoff signal, ζ' , such that $\forall \zeta \leq \zeta'$, 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 (21) and f'' satisfies equation (22).

The reporting strategy \hat{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 $\zeta' \neq \zeta$ for all feasible R' .

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

Proof. By construction, for values of $\hat{\theta} > \mu$, the analyst releases \hat{f} , else \check{f} is released. A forecast of \hat{f} prompts shareholder 1 sell to shareholder 2, and \check{f} prompts shareholder 2 to sell to shareholder 1. This is equivalent to the mechanism trade parameter $\rho(\hat{\theta})$ in equation (11) which maximized the two shareholders' gains to trade. \square

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 (22) evaluated at f'' is true and/or a lower cutoff for f' so that equation (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. \square

A.3.2 Trade and Transfer Prices

As long as the forecast, given the forecast strategy, satisfies equation (21), shareholder 2 will sell to shareholder 1. If the forecast induces an expectation of θ such that equation (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*

$$\alpha\{E[v_1(\theta)|f] - (1 - I_f)\lambda\} + (1 - \alpha)\{E[v_2(\theta)|f] - I_f\lambda\} \quad (24)$$

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 (24) is simply a weighted average of the expected utilities, it characterizes any negotiated transfer price $T(f)$ that satisfies

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

\square

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

A.4 Representing Reputation in the Repeated Game

To recapitulate, I model three agents: an analyst with private signal on the state of nature, and two shareholders who 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. Therefore, 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 A.1. Sequentially rational shareholders would anticipate this in the last period, and trade accordingly. Anticipating this, the analyst now plays stages $t \in \{1, 2, \dots, T - 2\}$ 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.

B Metrics for Trade and Forecasting Consistency

Metrics for trade, capturing gains to trade, and benefits of trading are constructed rather than relying upon the random realizations generated during the experiment, as described in Section 2.3.1. These metrics are free of any stochastic noise inherent in the realizations. Parameters for the states, signals, and probabilities thereof are described in Section 2.1.

B.1 Likelihood of Trade

The likelihood of trade is derived from subjects' input forecasting and trading decisions. This metric captures to what extent subject's input decisions prompted trade between the shareholders. This metric was defined for a round in a single economy.

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

$I_f(signal)$ is an indicator function, defined in every round for the three subjects within a single economy. The function is 1 if there was trade for the released forecast f and 0 otherwise. There was 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.

B.2 Likelihood of Capturing Gains to Trade

The likelihood of capturing gains to trade captures to what extent trade occurred when the state was A or D . Not all trade benefited the economy as a whole, only trade in states where gains to trade exceed trading costs. This metric is also defined in every round in a single economy, but while the prior metric was summed over all states, this metric is summed over states A and D .

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

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

B.3 Benefit of Trading

While trade in states A or D benefited the economy as a whole, it was possible the trade benefited one shareholder at the expense of another. The benefit of trade was calculated for each shareholder in an economy given the subjects' input forecasting and trading decisions. It captures to what extent the shareholders were better electing to trade given simply keeping their original endowment. This metric was defined for a round in a single economy and, to avoid notational clutter, the metric is defined separately for each shareholder.

Shareholder 1

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

Shareholder 2

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

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 f . Both $v_1(\cdot)$ and $v_2(\cdot)$ 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 2.1.

B.4 Metric of Forecasting Consistency

The forecasting consistency metric captures how consistently the analyst maps private signals into a released forecast over the set of eight rounds. The metric depends upon the subject's input forecasting decisions.

I construct a consistency metric based on the analyst's forecasting strategy, using the specific mapping from the analyst's private signals to forecasts. The construction follows. In the first round of a set, the metric was zero. In subsequent rounds, if the analyst's mapping is identical to mapping used in the last round, the metric was increased by one for the current round. If the mapping was different, the metric was 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 was used to arrive at a consistency metric for the set. For the first example above, the set metric would be 28, while the in second example, the metric would be 12. The metric captured how consistently the analyst chooses to forecast, not necessarily how consistently the analyst's set decision was classified. Note that in both examples the analysts were deceptive in every round, but differ in how the deceptive strategy was executed.

B.5 Mean and Inter-quartile Ranges of Metrics

Statistics for constructed metrics are shown in Tables 18 and 19 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 4.4.

	Mean	Minimum	1st Quartile	Median	3rd Quartile	Maximum
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 18: Metric Statistics: Low-Cost Treatment

	Mean	Minimum	1st Quartile	Median	3rd Quartile	Maximum
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 19: Metric Statistics: High-Cost Treatment

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 4.4

C Figures

C.1 Screen Captures of Experiment

Round: _____ 1 out of 64 Remaining Time [sec]: 37

Chances of States				
State	A	B	C	D
Chance of State	1/4	1/4	1/4	1/4

Payoffs to Owners				
State	A	B	C	D
Payoff to Owner 1	30	50	70	90
Payoff to Owner 2	0	40	80	120

Chances of State Clue Given Private Clue				
Clue Equal to	A	B	C	D
State is A	0.9	0.1	0	0
State is B	0.1	0.9	0	0
State is C	0	0	0.9	0.1
State is D	0	0	0.1	0.9

Trading Costs Borne by Buyer: 21

Reporter's Decision	
If Your Private Clue is	Release a Report Equal to
A	A <input type="radio"/> B <input type="radio"/>
B	A <input type="radio"/> B <input type="radio"/>
C	C <input type="radio"/> D <input type="radio"/>
D	C <input type="radio"/> D <input type="radio"/>

Confirm Decision

Figure 5: Entry screen for Analyst in High-Cost Treatment

Round 1 out of 64 Remaining Time [sec]: 6

Chances of States				
State	A	B	C	D
Chance of State	1/4	1/4	1/4	1/4

Payoffs to Owners				
State	A	B	C	D
Payoff to Owner 1	30	50	70	90
Payoff to Owner 2	0	40	80	120

Chances of State Clue Given Private Clue				
Clue Equal to	A	B	C	D
State is A	0.9	0.1	0	0
State is B	0.1	0.9	0	0
State is C	0	0	0.9	0.1
State is D	0	0	0.1	0.9

Trading Costs Borne by Buyer: 21

Owner 1 Trading Decision

If Reporter's Report is

	A	B	C	D
Trading Decision will be	<input type="radio"/> No Trade <input type="radio"/> Buy 7 <input type="radio"/> Buy 23	<input type="radio"/> No Trade <input type="radio"/> Buy 7 <input type="radio"/> Buy 23	<input type="radio"/> No Trade <input type="radio"/> Sell 83 <input type="radio"/> Sell 91	<input type="radio"/> No Trade <input type="radio"/> Sell 83 <input type="radio"/> Sell 91

Confirm Decision

Figure 6: Entry screen for Shareholder 1 in High-Cost Treatment

Round 1 out of 64 Remaining Time [sec]: 0

Chances of States				
State	A	B	C	D
Chance of State	1/4	1/4	1/4	1/4

Payoffs to Owners				
State	A	B	C	D
Payoff to Owner 1	30	50	70	90
Payoff to Owner 2	0	40	80	120

Chances of State Clue Given Private Clue				
Clue Equal to	A	B	C	D
State is A	0.9	0.1	0	0
State is B	0.1	0.9	0	0
State is C	0	0	0.9	0.1
State is D	0	0	0.1	0.9

Trading Costs Borne by Buyer: 21

Owner 2 Trading Decision

If Reporter's Report is

	A	B	C	D
Trading Decision will be	<input type="radio"/> No Trade <input type="radio"/> Sell 7 <input type="radio"/> Sell 23	<input type="radio"/> No Trade <input type="radio"/> Sell 7 <input type="radio"/> Sell 23	<input type="radio"/> No Trade <input type="radio"/> Buy 83 <input type="radio"/> Buy 91	<input type="radio"/> No Trade <input type="radio"/> Buy 83 <input type="radio"/> Buy 91

Confirm Decision

Figure 7: Entry screen for Shareholder 2 in High-Cost Treatment

Round: 2 out of 64 Remaining Time [sec]: 39

Chances of States				
State	A	B	C	D
Chance of State	1/4	1/4	1/4	1/4

Payoffs to Owners				
State	A	B	C	D
Payoff to Owner 1	30	50	70	90
Payoff to Owner 2	0	40	80	120

Chances of State Clue Given Private Clue				
Clue Equal to	A	B	C	D
State is A	0.9	0.1	0	0
State is B	0.1	0.9	0	0
State is C	0	0	0.9	0.1
State is D	0	0	0.1	0.9

Trading Costs Borne by Buyer: 21

Reporter's Decision						
If Your Private Clue is			Release a Report Equal to			
A			A	<input type="radio"/>	B	<input type="radio"/>
B			A	<input type="radio"/>	B	<input type="radio"/>
C			C	<input type="radio"/>	D	<input type="radio"/>
D			C	<input type="radio"/>	D	<input type="radio"/>

Confirm Decision

Results of Prior Rounds with Group						
Round	State	Report	Owner 1	Owner 2	Price	Net Payoff
1.01	C	D	Sell 91	Buy 91	91	13
1.02	D	D	Sell 91	Buy 91	91	13
1.03	B	A	Buy 7	Sell 7	7	13
1.04	C	D	Sell 91	Buy 91	91	13
1.05	A	A	Buy 7	Sell 7	7	13
1.06	D	D	Sell 91	Buy 91	91	13
1.07	D	D	Sell 91	Buy 91	91	13
1.08	B	A	Buy 7	Sell 7	7	13
1.09	A	A	Buy 7	Sell 7	7	13
1.10	A	A	Buy 7	Sell 7	7	13

Figure 8: Entry screen with Feedback for Analyst in High-Cost Treatment