A Peek Behind the Curtain that Conceals Self-Dealing

Gregory Waymire
(Emory University)

Radhika Lunawat
(University of California-Irvine)

and

Baohua Xin
(University of Toronto)

April 9, 2014

We wish to thank the staff at CIRANO in Montreal who helped in running our experiment. We benefited from comments and suggestions received in workshops presented at the following universities: Chapman, Emory (both the Psychology Department and Goizueta Experimental Brownbag), Illinois (Champaign-Urbana), Ole Miss, Penn State, Reading (Classics), Temple, Toronto, and Yale. Sudipta Basu, Kendall Bowlin, Jeremy Douthit, Jessen Hobson, Kathryn Kadous and Bart Wilson also provided helpful suggestions. Financial support of this research was provided to authors by the Economic Sciences Institute at Chapman University, the Carlson School at the University of Minnesota, the Rotman School of the University of Toronto, and the Goizueta Business School at Emory University. Prem Prakash Singh provided valuable expertise in developing the chat program used in our experiment and Stephen Deason provided additional research assistance. We especially benefitted from the guidance and friendship of John Dickhaut, who began this paper as a co-author. John passed away shortly after a first draft was completed, but before the extensive changes that occurred later. We dedicate this paper to his memory.

ABSTRACT

Self-dealing is difficult to study because its success depends on its concealment from its victims. A by-product of this is that researchers using naturally occurring data cannot directly observe strategies that successfully conceal self-dealing. We circumvent this problem by using an experiment where self-dealing emerges endogenously along with concealment actions that are fashioned by the self-dealer. We find that a significant percentage of subjects profitably self-deal and use multiple actions that likely have the combined effect of concealing self-dealing. These actions include opacity-preserving resource allocations, misleading direct communication, false promises about future resource divisions, and providing high early benefits to partners that can be subsequently exploited for personal gain. Interestingly, these choices also include profit skimming, which likely reduces both partner suspicions and the private gains from self-dealing. The ultimate consequence is that self-dealing occurs frequently, but its impact on the distribution of wealth and economic efficiency is limited considerably. Economies with frequent self-dealing produce comparable social gains to ones where self-dealing is largely eliminated by the provision of hard information that makes behavior transparent.

INTRODUCTION

Though the end of the rules of justice be, to hinder us from hurting our neighbour, it may frequently be a crime to violate them, though we could pretend, with some pretext of reason, that this particular violation could do no hurt. A man often becomes a villain the moment he begins, even in his own heart, to chicane in this manner. The moment he thinks of departing from the most staunch and positive adherence to what those inviolable precepts prescribe to him, he is no longer to be trusted, and no man can say what degree of guilt he may not arrive at. The thief imagines he does no evil, when he steals from the rich, what he supposes they may easily want, and what possibly they may never even know has been stolen from them. The adulterer imagines he does no evil, when he corrupts the wife of his friend, provided he covers his intrigue from the suspicion of the husband, and does not disturb the peace of the family. When once we begin to give way to such refinements, there is no enormity so gross of which we may not be capable.
Adam Smith[1]

F—k my victims. I carried them for twenty years, and now I'm doing 150 years.
Bernard Madoff[2]

Self-dealing is profitable only when it is successfully concealed from its victims. Bernard Madoff sustained a massive Ponzi scheme for two decades through falsified financial records and an ability to build false trust through force of personality (Henriques 2012). Cases like the Madoff scheme raise significant questions about why such frauds occur. Researchers have studied both individuals' willingness to deceive for personal gain and their ability to spot deception by others. Evidence suggests that some may be averse to lying while others are willing to cheat others for personal gain (Charness and Dufwenberg 2006; Gneezy 2005). Other research suggests that detecting deception is not straightforward and the success of such efforts varies widely across differing contexts (Vrij 2008).

Extant research on deception sheds less light on how interaction between self-dealers and victims unfolds while successfully concealed self-dealing is occurring. A major reason for this is that successfully concealed self-dealing by definition allows neither the victim nor a third-party researcher to directly observe such conduct. Because we can more readily observe

[1] *The Theory of Moral Sentiments*, Part III, Chapter VI.
[2] Attributed to Madoff by a fellow inmate who had expressed concern for the investors he had defrauded; quoted in Fishman, S. "Bernie Madoff, Free at Last, *New York* magazine (June 10, 2010).

deception and fraud after the perpetrator has been caught, we are left with incomplete answers to questions such as: what devices do fraud perpetrators use to successfully conceal self-dealing, what magnitude of private gains can perpetrators extract and still be able to conceal their behavior, and what social losses occur when self-dealing may be suspected but is never confirmed? Our contribution is an experiment that provides evidence on these questions.[3]

Our experiment involves multi-period interactions between an "investor" and a "trustee" in an investment game with uncertain payoffs (Berg, Dickhaut, and McCabe 1995; King-Casas et al. 2005; Coricelli, Morales, and Mahlstedt 2006; Basu et al. 2009; Lunawat 2012; 2013). The investor is endowed each period with resources that can be invested to generate a gain that the trustee then divides. The wealth generated for a period depends on the investment level and the realization of an uncertain multiplier observed only by the trustee. The investor and trustee directly communicate through chat (i.e., "cheap talk") at the end of each period. We refer to this as a *NoInfo-Chat* economy.

Self-dealing and concealment strategies used by *NoInfo-Chat* trustees can arise endogenously during the experiment. More importantly, we can observe several aspects of the trustee's behavior that the investor cannot. This means that we can identify instances where a trustee returns a lower percentage to the investor and takes actions that could lead the investor to falsely believe that the trustee did not self-deal. We can also observe actions by a *NoInfo-Chat* trustee that may amplify false trust by the investor including the provision of high early returns or the use of communication to build a personal relationship that can be exploited for private gain.

Two experimental manipulations allow us to measure how opportunities for trustee information provision alter behavior in the *NoInfo-Chat* economy. First, we run an economy

[3] One purpose of an experiment is to provide controlled variation to examine behavior in counterfactual worlds that cannot be directly observed in the field (Smith 2004; Falk and Heckman 2009).

where the trustee can choose at the start of a period to have the investor told (at the end of the period) the total amount that the trustee had divided. This allows the investor to infer the realized multiplier with the result being that trustee return decisions become fully transparent in this *HardInfo-Chat* economy. Comparing the *NoInfo-Chat* and *HardInfo-Chat* economies allows us to identify how *NoInfo-Chat* trustee behavior is altered by information that can improve transparency and increase accountability. Second, we also run *NoInfo* and *HardInfo* economies that are otherwise identical to the *NoInfo-Chat* and *HardInfo-Chat* economies except that subjects cannot communicate through chat. The *NoInfo* and *HardInfo* economies allow us to gauge how trustee and investor behavior differ in settings where exchange cannot be personalized through direct communication (i.e., chat).

Our analysis first considers whether *NoInfo-Chat* trustees can convince their investor partners to invest despite the possibility that the trustee will self-deal. These tests show that sizable social gains from exchange can be generated by economies where concealed self-dealing is possible. Investment is lowest in the *NoInfo* economies (48% of endowment invested), *HardInfo* economies perform considerably better (65% invested), and adding direct communication further increases welfare gains for *HardInfo-Chat* economies (86% invested). Surprisingly, *NoInfo-Chat* investors invest even more than *HardInfo-Chat* investors (88% of endowment). This suggests that direct communication personalizes exchange with effects that swamp those associated with a formal reporting institution.

Our next tests establish that *NoInfo-Chat* trustees are more likely than *HardInfo-Chat* trustees to exploit their investor partners through self-dealing. Specifically, the percentage of total resources that *NoInfo-Chat* trustees send back to investors declines as the realized multiplier increases. This effect is less pronounced for *HardInfo-Chat* trustees, especially for later periods where *NoInfo-Chat* trustees continue self-dealing while *HardInfo-Chat* trustees make more egalitarian resource divisions.

Our final tests explore the nature of actions taken by *NoInfo-Chat* trustees that could conceal self-dealing. *NoInfo-Chat* trustees cluster into three groups based on their return decisions: some always behave honestly, others will behave dishonestly when their partners make low investments, and the remainder repeatedly self-deal even though their partners make large investments (Gneezy, Rockenbach, and Serra-Garcia 2013). Concealment actions by *NoInfo-Chat* trustees include opacity preserving returns, deceptive chat disclosures in periods when higher multipliers are realized, the making (and breaking) of explicit promises about future return decisions, and providing high early returns to investors that can build false investor trust. These strategies also involve profit "skimming" rather than taking the maximum amount possible. While profit skimming likely allows a *NoInfo-Chat* trustee to lessen investor suspicions of self-dealing, it also means *NoInfo-Chat* investors earn rates of return that are higher than they might otherwise be. The ultimate result is that *NoInfo-Chat* investors' rates of return on endowment are only marginally lower than those of *HardInfo-Chat* investors and remain substantially higher than those of *HardInfo* and *NoInfo* investors.

One implication of our results is that a formal system for reporting hard information *per se* may not reduce social losses from concealed self-dealing when direct communication is available. Direct communication allows self-dealers to implement concealment strategies that sustain *both* high social gains from exchange and modest private benefits from self-dealing. Our findings are generally consistent with arguments that manipulation may be benign or even favorable in some cases (Arya, Glover, and Sunder 2003).

Our experimental evidence also adds to the body of knowledge in two research areas. A long line of experiments has established that cheap talk increases cooperation (Dawes, McTavish, and Shaklee 1977; Crawford 1998; Ledyard 1995; Cai and Wang 2006) and that these gains accrue in part because of individual promises and apologies (Charness and Dufwenberg 2006; Schniter, Sheremeta, and Sznycer et al. 2013). Our experiment extends these findings to show that communication can add considerable economic value even when concealed self-dealing is frequent. A second set of experimental studies shows that deception can be either implicit, explicit, or involve telling strategic “truths” (Croson, Boles, and Murnighan 2003; Gneezy 2005; Sutter 2009; Kriss, Nagel, and Weber 2013). We add to this work by showing that deceivers can spontaneously craft complex strategies for deception that preserve both social gains from exchange and private gains from self-dealing.

Our study is subject to the standard caveat about the external validity of laboratory experiments. It is impossible to incorporate into a laboratory experiment every complexity present in naturally occurring settings where deception and self-dealing occur. Since an effective way to address this issue is through replication and extension, we offer several suggestions for useful future research in the final section of the paper. This research can aid our understanding of how information alters the boundary of feasible exchange in settings where self-dealing and deception can destroy gains from exchange.

The paper is organized as follows. We first describe the structure of our experiment and describe our approach to data analysis. Data collection is discussed next and results are then presented. A final section summarizes our findings and offers suggestions for future research.

EXPERIMENT STRUCTURE AND APPROACH TO DATA ANALYSIS

In this section we first describe the structure of our experiment and discuss how it differs from similar experiments in prior research. We then describe our approach to the analysis of data gathered from the lab.

Experiment Structure

Our experiment is a multi-period investment game derived from the one-period game in Berg, Dickhaut, and McCabe (1995), and is based primarily on the multi-period, multi-player game used by King-Casas *et al*. (2005) and Basu *et al*. (2009). The structure of our experiment differs from these studies because we incorporate payoff uncertainty as in Coricelli, Morales, and Mahlstadt (2006) and Lunawat (2012, 2013).

Each economy includes four persons; two play the role of "investor" and two play the role of "trustee" (see Panel A of Figure 1).[4] Each period begins when an investor receives a separate initial endowment of ten currency units (*lira*) for each trustee. All or part of this endowment can be sent to that trustee in whole numbers (see Panel B of Figure 1). The trustee receives a varying amount equal to the product of the investment and a multiplier that ranges from one to five in whole numbers with each being equally likely. Thus, if an investor invests all ten *lira*, the trustee will receive 10, 20, 30, 40, or 50 *lira*.

The trustee is told the amount sent by the investor and the amount she has received. The investor is not told the amount received by the trustee or the realized multiplier. Thus, the trustee knows the realized value of the multiplier but the investor does not. A separate multiplier is randomly drawn (with replacement) in every period for each of the four dyads in the economy. This precludes the possibility that the investor can infer the multiplier based on her interactions with the other trustee. After learning the investment level and total amount received, the trustee decides how much to send back to the investor.

[4] Each economy in Basu *et al*. (2009) included five investors and five trustees. We reduced that number to two investors and two trustees to preserve interactions with multiple partners while reducing the likelihood that the task would be excessively complex.

We implement a multi-period game by repeating this process ten times (King-Casas *et al*. 2005; Basu *et al*. 2009). To lessen end game behavior, subjects are not told how many periods the game will be repeated. Subject anonymity is maintained throughout. This simple case is referred to as the *NoInfo* economy.

Two experimental manipulations vary whether and how the trustee can provide information on the realized multiplier to the investor. The first allows the trustee to choose whether to inform the investor about the total amount received by the trustee. In this *HardInfo* economy, each trustee (B1 or B2) is asked prior to the investor's decision and the draw of the multiplier whether she wants the experimenter to inform a given investor (A1 or A2) of the actual amount received by the trustee. This institution provides hard information in that the actual amount received by the trustee is revealed to the investor. Because the investor will know the actual realized multiplier, the trustee's return decision becomes completely transparent when hard information is provided.

Our second manipulation allows the trustee and investor to communicate through a chat mechanism. Investors and trustees could send text messages via a text-based electronic notebook on their computer screens. Subjects could send any message they wished that excluded personal information during a three-minute interval at the end of each period prior to the final period.[5] Any specific message could be sent only to a single partner; a subject wanting to send identical messages to both partners could do so only if they were sent at different times. To avoid collusion, text messages could not be sent to the other subject playing the same role – i.e., one investor (trustee) could not send messages to the other investor (trustee). We refer to an economy where *only* the chat institution is available to transmit information as a *NoInfo-Chat*

[5] Communication via chat involves only text-based communication to avoid confounds due to facial expressions and vocal changes (Ekman 2006; Hobson et al. 2012).

economy. Economies where a trustee can provide hard information to their partners *and* communicate through chat are referred to as *HardInfo-Chat* economies.

Our experimental economies share several features with those examined in prior research. Prior research demonstrates that communication promotes cooperation and increases economic efficiency (Dawes, McTavish, Shaklee 1977; Crawford 1998; Ledyard 1995; Ochs 1995; Sutter and Strassmair 2009). Gains from communication arise in part from non-binding promises made by subjects that prefer to keep their word (Ben-Nera, Putterman, and Ren 2011; Bochet, Page, and Putterman 2006; Charness and Grosskopf 2004; Cohen, Wildschut, and Insko 2010; Duffy and Feltovich 2002; Ellingsen and Johannesson 2004; Vanberg 2008). Charness and Dufwenberg (2006) show how guilt aversion improves cooperation in a single-period trust game with uncertain payoffs and simple communication by the trustee.

While individuals prefer not to deceive, they will do so for personal gains when the opportunity is available. Gneezy (2005) documents that some individuals will deceive a trading partner for a private benefit when the other person's loss is small, but fewer will do so when the other's loss is large. Croson, Boles, and Murnighan (2003) document deception that affects bargaining behavior and outcomes in an ultimatum game experiment and Sanchez-Pages and Vorsatz (2007) demonstrate that not everyone prefers truth telling. Sutter (2009) demonstrates sophisticated deception in a sender-receiver experiment that includes truth telling when the message sender expects the receiver to not act upon the message. Kriss, Nagel, and Weber (2013) demonstrate that deceivers in an ultimatum game use both explicit deception (i.e., untruthful statements) and implicit deception (i.e., information revealing actions). Within the accounting literature, King (2002) shows that non-credible communication in an auditing context can be reduced through social pressure from other group members, and Hales, Kuang, and Venkataraman (2011) identifies conditions where individual investors' beliefs will be more sensitive to "hyped" information.

On the flip side of the deception question, there is some evidence that individuals can identify deception at rates better than chance, but these gains are not large (Belot and van de Ven 2013). This accords with evidence from numerous studies suggesting that detecting deception based on verbal and visual cues can be subject to frequent errors (Vrij 2008).

Research on individuals' willingness to deceive and their ability to detect deception provides less insight on how successfully concealed self-dealing is sustained. A key question from an economic standpoint is: how can a self-dealer successfully conceal cheating from a disadvantaged partner when the parties communicate spontaneously and opportunities for detection are plentiful? Even if the self-dealer makes no obvious mistakes that reveal deceit, a disadvantaged partner may still suspect that all is not right. Do suspicions that are not confirmed have material consequences for the distribution of wealth and economic efficiency? What kinds of communication and information disclosure strategies emerge as a part of a scheme to conceal material self-dealing?

Answering these questions with naturally occurring data is impossible since successfully concealed self-dealing is never directly observed in such data. We use an experiment to circumvent this problem. We incorporate several economic forces into our experiment that are likely present in naturally occurring settings where self-dealing and deception may occur but cannot be observed by a researcher. Experimental control thus allows us to see "close-up" how endogenous self-dealing and deception is shaped by institutions that enable information to be transferred from more informed trustees to less informed investors.

The contribution of our experiment is to allow us to directly observe data that would otherwise be concealed in naturally occurring settings because self-dealers act to intentionally conceal their actions from those harmed by their conduct. Specifically, a *NoInfo-Chat* trustee has superior information that makes self-dealing feasible, but she must deceive a trading partner in order to avoid a reputational penalty if self-dealing is suspected. Unstructured bilateral communication between the trustee and the investor allows for a variety of messages by both parties. In this way chat sessions in our experiment are analogous to a corporate conference call where a dishonest executive might successfully deceive analysts and investors or may face probing questions that reveal deception (Larcker and Zakolyukina 2012). In sum, our experiment is useful because it allows us to directly observe behavior that may exist but remain hidden behind the curtain that conceals self-dealing in naturally occurring settings.

Approach to Data Analysis

The issue motivating our experiment is how the potential for self-dealing by one party affects another's willingness to invest, and how that investment choice feeds back to influence the magnitude of self-dealing and related actions to conceal such conduct. We begin our data analysis by providing evidence on the following research question: How do total wealth and its distribution vary across the *NoInfo-Chat*, *HardInfo-Chat*, *HardInfo*, and *NoInfo* economies?

We expect that total investment will be lowest in *NoInfo* economies since trustees cannot transmit any information to investors, and that *HardInfo-Chat* investors will invest the most since hard information provides transparency and chat improves coordination. This implies that total investment in *HardInfo* and *NoInfo-Chat* economies will likely lie somewhere between a lower bound set by *NoInfo* economies' investment and an upper bound representing investment in *HardInfo-Chat* economies. The extent to which total investment in the *HardInfo* economies exceeds that of *NoInfo-Chat* economies (or vice versa) will depend on the relative effects of hard information and chat communication on investor trust.

The distribution of wealth within our experimental economies depends on both investment levels and trustee resource division decisions. Zero investment means that the investor gets all of the wealth in the economy and the trustee gets nothing. Positive investment can generate gains and lessen the disparity between investor and trustee wealth. We expect the distribution of wealth to be most egalitarian in *HardInfo-Chat* economies because investment is expected to be highest and a trustee can be held to account for unequal resource divisions. The relative ranking of the other economies will depend on differences in investment level as well as trustee resource division decisions.

We next consider how a *NoInfo-Chat* trustee will exploit her informational advantage to extract greater personal benefits when the realized multiplier is larger. To illustrate, assume a *NoInfo-Chat* investor invests the maximum of ten *lira* and the realized multiplier is 5 (i.e., the trustee receives 50 *lira*). If the trustee splits the total available equally with the investor and sends back 25 *lira*, it is likely that the investor will make a maximum investment in the next period. However, the investor might continue to invest the maximum if the trustee sends back only 20 *lira* and is convinced by the trustee that the realized multiplier was 4. The limit of this argument is reached when the trustee sends back five *lira* and claims (falsely) that the realized multiplier was 1. Of course, this will be difficult to do every period since a string of consistently low returns could arouse the investor's suspicion and lead her to withhold investment.

A *HardInfo-Chat* trustee is likely to have less flexibility in obtaining increased private benefits since an investor could condition her investment on whether the trustee provides hard information. Ijiri (1975, 36) describes hard information as resulting from "the processing of verifiable facts by justifiable rules" such that it is "difficult for people to disagree" whereas soft information "can easily be pushed in one direction or the other." Prior theoretical research suggests that verifiable accounting reports can favorably influence the quality of managers' other disclosures (e.g., Christensen 1981; Baiman and Evans 1983; Dye 1983; Gigler and Hemmer 1998; Arya *et al*. 2004).[6] Providing hard information allows a trustee to establish a reputation for equitably sharing gains with the investor (Basu and Waymire 2006; Basu *et al*. 2009).[7] Our second set of empirical tests focuses on the following alternative hypothesis: Are *NoInfo-Chat* trustees' returns to investors lower than those provided by *HardInfo-Chat* trustees when the realized multiplier is larger?

A *NoInfo-Chat* trustee's ability to secure greater private benefits depends critically on whether the investor's beliefs about the trustee's behavior can be manipulated. A concealment strategy used by a *NoInfo-Chat* trustee likely involves several interdependent actions that collectively manipulate investor beliefs to the trustee's advantage. First, the trustee might make a return that preserves a shroud of opacity over her resource division decision. This would entail returning a lower amount that could feasibly be attributed to a lower multiplier. An opacity preserving return might also involve "skimming" profits rather than trying to capture the maximum possible gains in order to lessen investor suspicions that the trustee divided resources unequally.[8]

A self-dealing *NoInfo-Chat* trustee might also try to mislead the investor by making a false disclosure of the realized multiplier or remaining silent when the realized multiplier is larger. A self-dealing *NoInfo-Chat* trustee who is particularly cunning might craft a concealment strategy that involved developing a personal relationship that can be exploited for personal gain in future periods. Our hypothesis is that *NoInfo-Chat* trustees will be more likely than *HardInfo-*

---

[6] Ball (2001, 135-6) remarks: "a reliable system of timely and accurate reporting… exerts a discipline on managers' privately held expectations and on their publicly stated expectations (in the form of plans and forecasts)."

[7] Harder information enables comparison of *ex post* performance and *ex ante* promises, which likely increases gains from exchange when information moves across time and place (Dickhaut *et al*. 2010).

[8] Trivers (1971, 46) uses the label "subtle cheating" to describe such behavior, which he suggests involves "reciprocating, but always attempting to give less than one was given, or more precisely, to give less than the partner would give if the situation were reversed."

Chat trustees to use complex strategies involving multiple actions to manipulate investor beliefs in order to conceal self-dealing.[9]

DATA COLLECTION

We conducted our experiments at the Center for Interuniversity Research and Analysis on Organization (CIRANO) in Montreal, Canada. CIRANO staff recruited subjects and ran 32 sessions (eight sessions for each of the four experimental conditions).[10] One hundred twenty-eight subjects from CIRANO's standard subject pool participated and remain completely anonymous to the authors.[11]

Each condition-specific experiment-session included four subjects. Each subject was randomly assigned to either an investor or a trustee role (with the restriction that every session contain two investors and two trustees). While subjects were not informed of the number of periods to mitigate end game effects, they were informed via recruiting materials that the experiment would last approximately two hours.

Subjects interacted anonymously over a local computer network facilitated by *z-Tree* software (Fischbacher 2007). The program kept track of all amounts sent and received by each subject in every period, provided feedback information to subjects, and tracked the time at which each subject confirmed investment or return decisions.

[9] A term to describe such behavior is "Machiavellian intelligence," which Byrne and Whiten (1997, 20) define as "social manipulation to achieve individual benefits at the expense of group members, but without causing such disruption that the individual's membership of the group is put in jeopardy." Dunbar and Shultz (2007) describe such behavior as resulting from a "social brain" shaped by Darwinian selection.

[10] In one session for the *HardInfo-Chat* economy, we experienced a computer malfunction that resulted in the loss of chat data after the fourth round for that specific session only. The other data from the session (e.g., investor investments, trustee returns, hard information provision) were not affected. Consequently, we include data from this specific session in all of our analyses except where we examine variables that require coded chat data. For variables based on coded chat data, we exclude data from this session.

[11] At the time of our experiment, the entire subject pool at CIRANO included 3,223 individuals, of which 52% were females. Over 90% are students at local universities like Concordia and McGill. Because our experimental materials were in English, this meant that persons in the subject pool who spoke only French (slightly over 40% of the CIRANO subject pool) could not serve as subjects in our experiment.

Subjects first received and read written experiment instructions, and then took a quiz to ensure sufficient understanding of experiment instructions. The experiment facilitator checked quiz answers and resolved discrepancies privately before the beginning of the first period. A copy of the experiment instructions for the *HardInfo-Chat* economies is shown in the Appendix.

A text-based electronic notebook was provided to subjects in *NoInfo-Chat* and *HardInfo-Chat* economies. The notebook is a blank textbox, situated at the right side of the computer screen. The *Textbox* program recorded the contents of every message sent by subjects during the experiment-session. These chat messages provide the data on the timing and content of chat communication between individual subjects. *Textbox* was the only method of communication available to subjects. Subjects did not have access to paper and pencils, nor could they use computer programs other than *z-Tree* and *Textbox* (the latter available for only those subjects in the *NoInfo-Chat* and *HardInfo-Chat* conditions). Our analysis of chat data is based on chat after rounds 1 through 9, inclusive.

Play within any given period of a *HardInfo* or *HardInfo-Chat* economy began when each trustee (B1 or B2) was asked whether she wishes to have the amount received from a given investor (A1 or A2) revealed to that investor at the end of the period. The investor was then informed of the trustee's choice to reveal the amount she receives from that investor. Each investor then decided how much of their ten *lira* endowment to invest with each trustee. All investors' investment decisions were required before trustees were notified of both the amount invested by the investor and the total amount they had received as a result of this investment. Trustees then decided how much to send back to each investor from the amount that had been received from that investor. All trustees' return decisions were required before investors learned of the amounts that had been sent back by trustees.

After receipt of returns by the investor, subjects paired in a *NoInfo-Chat* or a *HardInfo-Chat* economy could send online text messages to each of their two partners individually for a three-minute interval. The next trading period began after all subjects had completed the chat session and were finished reviewing feedback from the prior round. At the end of the tenth trading period, the *lira* earned for all periods was summed and converted to Canadian dollars at a rate of four cents per *lira*.

EXPERIMENTAL EVIDENCE

This section is organized as follows. We first describe subjects' use of the chat and hard information reporting mechanisms since our results likely depend on whether subjects actually use these institutions. We next describe investment levels and the distribution of wealth in our four experimental economies. Third, we report results from tests of whether *NoInfo-Chat* trustees exploit their informational advantage to capture greater personal benefits for larger multiplier realizations. Evidence on various features of *NoInfo-Chat* trustee concealment strategies is then reported.

Subjects' Chat Usage and Hard Information Provision

Table 1 suggests that *NoInfo-Chat* trustees send more messages than *NoInfo-Chat* investors, *HardInfo-Chat* trustees, and *HardInfo-Chat* investors.[12] *NoInfo-Chat* trustees also send somewhat longer messages than other types of subjects. *HardInfo-Chat* and *HardInfo* trustees both choose to frequently provide hard information, but this tendency is higher in the *HardInfo-*

[12] Differences in chat message frequency between the *NoInfo-Chat* and *HardInfo-Chat* economies apply across periods during the experiment. The average number of trustee and investor messages sent in the first chat period is higher in *HardInfo-Chat* economies than in *NoInfo-Chat* economies, but that reverses in the next chat period and holds until the final chat period.

Chat economies than in the *HardInfo* economies – 87.2% versus 73.4% of periods, respectively (see Panel C of Table 1).[13]

Investment Levels and the Distribution of Wealth in Different Economies

Table 2 provides summary statistics on the percentage of endowment invested (Panel A) and percentage of total wealth captured by the trustee for each type of economy (Panel B). These data indicate that (1) materially higher levels of investment are generated in economies where direct communication through chat is possible, and (2) the division of total wealth generally favors investors in all economies except the *NoInfo-Chat* economies.

Panel A of Table 2 show the mean and median of the economy-wide average per-period investment as a percent of endowment. We examine total investment since the wealth generated beyond endowment depends entirely on the investor's willingness to invest at least part of her endowment. To construct this table, we first computed the average percent of endowment invested for each of the 32 economies (four types of economies x eight replications per economy) using the 40 observations available for each economy (four dyads x ten periods). We then calculated the means and medians for each type of economy using the eight averages for that specific type of economy.

The mean (median) of the economy-wide percent invested equals 88.1% (92.9%) for the eight *NoInfo-Chat* economies. These amounts are comparable to the mean and median percent invested in *HardInfo-Chat* economies (mean = 85.5%; median = 87.8%). A Mann-Whitney test

[13] Both sets of trustees typically start reporting hard information in the first two periods. *HardInfo-Chat* trustees typically continue providing hard information after doing so initially – only six cases occur where the trustee does not supply hard information when she has done so in the prior period. In the *HardInfo* economies, 63% of the dyads have at least one period where the trustee did not provide hard information after having done so in period t-1.

cannot reject a null of identical distributions for the economy-wide average percent invested in the eight *NoInfo-Chat* and the eight *HardInfo-Chat* economies.

The economy-wide percent of endowment invested is substantially lower in the two economies where chat is not available. The mean (median) of the economy-wide average percent invested is 65.3% (65.0%) for the eight *HardInfo* economies and 47.7% (48.0%) for the eight *NoInfo* economies. The distributions of percent invested in the *HardInfo* and *NoInfo* economies are significantly different (one-tailed p-value of 0.014) using a one-tailed Mann-Whitney test.

The materially higher investments in *NoInfo-Chat* and *HardInfo-Chat* economies compared to the *NoInfo* and *HardInfo* economies are consistent with prior research suggesting that communication increases cooperation. The sizable difference in the percent of endowment invested in the *NoInfo-Chat* relative to the *HardInfo* economies (88.1% versus 65.3%) is surprising since it suggests that cheap talk communication is more important than reporting hard information in generating investor trust within the context of our experiment.

Figure 2 shows that the economically significant difference in investment levels between economies with chat versus economies without chat is sustained into later periods of the experiment. This figure shows that the rate of investment in the *NoInfo-Chat* and *HardInfo-Chat* economies increases dramatically in period 2-4. In contrast, investment in the *HardInfo* economies is comparable to that of the *NoInfo-Chat* and *HardInfo-Chat* economies in period 1, but (like the *NoInfo* economies) does not increase dramatically in later periods.

Panel B of Table 2 shows the economy-wide average percentage of total wealth captured by the trustee for each type of economy. Total wealth in a given period equals the sum of the amount retained by the investor from her endowment and the amount invested multiplied by the realized multiplier. Trustees capture more than half of the total wealth only in the *NoInfo-Chat* economy (mean 50.3%; median 50.7%). Trustees' share is lowest in *NoInfo* economies (mean 34.3%).

***NoInfo-Chat* Trustee Resource Division Decisions**

We now consider how *NoInfo-Chat* trustees' information advantage allows them to extract greater private benefits when the realized multiplier is larger. If this is the case, then the difference in percentage sent back by *NoInfo-Chat* trustees and *HardInfo-Chat* trustees will be negative for low multiplier realizations and become more negative as the multiplier increases.

The first two rows in panel A of Table 3 show the mean and median of the economy-wide average percentage of the total received that is sent back by *NoInfo-Chat* and *HardInfo-Chat* trustees. The economy-wide average percent returned equals the mean of the amount sent back divided by the amount received by the trustee calculated using all available dyad-period observations for a given economy.[14] The mean (median) of the eight economy-wide average percent returned equals 45.8% (46.0%) for the *NoInfo-Chat* economies compared to 47.5% (49.3%) for the *HardInfo-Chat* economies. The third row in panel A of Table 3 shows that the average percentage of cases where the trustee sends back less than half is greater for *NoInfo-Chat* trustees (34.1%) than for *HardInfo-Chat* trustees (24.9%). The evidence in Panel A of Table 3 suggests that *NoInfo-Chat* trustees return marginally lower amounts than *HardInfo-Chat* trustees across *all* multiplier levels.

To more directly identify how *NoInfo-Chat* trustee returns vary with the multiplier level, we estimate a regression model where the sensitivity of trustee returns to realized multipliers can vary between the *NoInfo-Chat* and *HardInfo-Chat* economies:

$\%RET_{ijt} = \alpha + \beta_1 \, INV_{ijt} + \beta_2 \, MULT_{ijt} + \beta_3 \, PERIOD_{ijt} + \beta_4 \, MULT_{ijt} \, PERIOD_{ijt} + \beta_5 \, \textit{NoInfo-Chat}_{ij} +$

[14] *NoInfo*-Chat investors make a zero investment three times and *HardInfo*-Chat investors do so six times. These observations are excluded when calculating the economy-wide percent returned.

β_6 INV_{ijt} *NoInfo-Chat*$_{ij}$ + β_7 $MULT_{ijt}$ *NoInfo-Chat*$_{ij}$ + β_8 $PERIOD_{ijt}$ *NoInfo-Chat*$_{ij}$ +
β_9 $MULT_{ijt}$ $PERIOD_{ijt}$ *NoInfo-Chat*$_{ij}$ + ε_{ijt} (1)

The dependent variable is the percentage of the total received by trustee j in period t that is returned to investor i ($\%RET_{ijt}$). Independent variables include the level of investment by investor i with trustee j in period t (INV_{ijt}), the multiplier realized in period t for the dyad of investor i and trustee j ($MULT_{ijt}$), the period of the experiment ($PERIOD_{ijt}$), and a 0-1 variable equal to 1 for *NoInfo-Chat* economies and 0 for *HardInfo-Chat* economies. We estimate (1) using Ordinary Least Squares with standard errors clustered by trustee. The sample used to estimate equation (1) includes all *NoInfo-Chat* and *HardInfo-Chat* returns except the nine cases where the trustee could not make a return decision because the investor had invested zero *lira*.

If both *NoInfo-Chat* and *HardInfo-Chat* trustees take greater personal gains when the realized multiplier is more positive, the coefficient on $MULT_{ijt}$ (β_2) will be negative. If the extent to which a trustee takes a greater share for larger multipliers is on average more pronounced in *NoInfo-Chat* economies at the start of the game, the coefficient on the interaction between MULT and *NoInfo-Chat* (β_7) will be negative. If *NoInfo-Chat* trustees' ability to extract private gains widens over time relative to *HardInfo-Chat* trustees, the coefficient on the interaction between MULT, *NoInfo-Chat*, and PERIOD (β_9) will be negative.

Panel B of Table 3 shows the estimated coefficients with corresponding t-statistics and Figure 3 plots the marginal impact of multiplier variation on the percent returned as a function of period for the *HardInfo-Chat* economies [β_2 + (β_4 * PERIOD)] and *NoInfo-Chat* economies [(β_2 + β_7 + ((β_4 + β_9) * PERIOD)]. Note first that the coefficient on MULT is reliably negative (β_2 = -0.076, t = -2.86). The positive coefficient on the interaction between MULT and *NoInfo-Chat* (β_7 = 0.061) suggests that the multiplier effect for *NoInfo-Chat* economies in the first period (equal to $\beta_2 + \beta_4 + \beta_7 + \beta_9$) is moderately negative (= -.019) while the multiplier effect for *HardInfo-Chat* economies in the first period ($\beta_2 + \beta_4$) is over three times larger (= -.068).

Most importantly, the coefficient on the three-way interaction between MULT, *NoInfo-Chat*, and PERIOD is reliably negative (β_7 = -0.012, $p < 0.001$), which suggests that the effect of the multiplier on percent returned varies with time differently across the two economies. Figure 3 shows that by the tenth period, variation in the realized multiplier level exerts no discernible impact on trustee returns in *HardInfo-Chat* economies (sum of relevant coefficients = 0.004) whereas the return percentage for *NoInfo-Chat* trustees is lower for higher multipliers (coefficient sum = -0.055). Overall, the evidence in Table 3 and Figure 3 suggests that *NoInfo-Chat* trustees secure greater private gains for larger realized multipliers in later periods whereas *HardInfo-Chat* trustees' return decisions become invariant to realized multipliers.[15]

***NoInfo-Chat* Trustees' Concealment Strategies**

We now consider elements of strategies used by *NoInfo-Chat* trustees to potentially conceal self-dealing. These include both "real" decisions that alter the distribution of resources as well as information manipulation that may induce false beliefs by the investor. An example of the former is a resource sharing decision that increases the trustee's share of the gains while preserving opacity from the investor's perspective. Information manipulation can involve either false disclosures or advantageous silence about realized multipliers during chat sessions. We first examine the incidence of opacity preserving returns by *NoInfo-Chat* and *HardInfo-Chat* trustees, and then present evidence on *NoInfo-Chat* and *HardInfo-Chat* trustees' communication choices. We conclude this sub-section by examining the consequences of *NoInfo-Chat* trustees' concealment strategies for investor rates of return.

[15] *HardInfo-Chat* trustees returns become increasingly egalitarian – e.g., they split the amount received equally in 28 of 32 cases in period 10.

An *opacity preserving return* (OPR) by a trustee is one with the potential to create false beliefs by the investor about the multiplier's realized value. This requires a return by the trustee that implies a multiplier that is both feasible (i.e., a whole number) and lower than that actually realized. To illustrate, assume an investment of ten *lira* and a realized multiplier of 5 in a situation where the investor expects an equal resource division by the trustee. If the trustee returns 22 *lira* to the investor, the multiplier implicit in the trustee's return equals 4.4 (22 x 2/10), which is lower than the actual multiplier of 5, but is infeasible since it is not a whole number. In contrast, a return of 20 *lira* has a feasible implicit multiplier of 4. For present purposes, we define an OPR specifically to be a return where the trustee returns an amount that implies a *feasible* multiplier that is *lower* than that actually realized under the assumption that the investor expects the trustee to send back half of the total amount she received.

We calculate the percentage of OPRs for a given economy as the number of trustee OPR returns divided by the number of non-zero investments for that economy. The percentage of OPRs for the eight *NoInfo-Chat* economies ranges from a low of 8.1% to a high of 47.5% with a mean of 22.6% (see Panel A of Table 4). The average OPR percentage for the eight *HardInfo-Chat* economies is 7.1% with a minimum of 0% and a maximum of 21.6%. The difference in OPR frequency across the *NoInfo-Chat* and *HardInfo-Chat* economies is significant at p = 0.005 based on a one-tailed Mann-Whitney test.[16]

The *NoInfo-Chat* and *HardInfo-Chat* economies also exhibit different patterns of OPR frequency for early and late periods (see Panel B of Table 4). The average of the mean OPR percentage for *NoInfo-Chat* trustees is higher in periods 6-10 (24.4%) than in periods 1-5 (20.7%). In contrast, the average mean OPR percentage for *HardInfo-Chat* trustees in periods 6-10 is only 2.2%, which suggests that OPRs are largely eliminated in later periods for *HardInfo-Chat*

[16] Comparable mean OPR frequencies for the *NoInfo* and *HardInfo* economies are 21.7% and 25.8%, respectively.

economies.[17] This is generally consistent with our prior finding that *HardInfo-Chat* trustees' resource division decisions are independent of realized multipliers by the final period of the experiment (see Figure 3).

Most OPRs by *NoInfo-Chat* trustees involve "skimming" of profits rather than taking the maximum gain possible. For example, suppose a *NoInfo-Chat* investor invests 10 *lira* and the realized multiplier equals 5. If the trustee seeks to skim off only a portion of the maximum available, she might return 15 *lira* to the investor and claim that the realized multiplier equaled 3. This lowers the trustee's private gains (compared to sending back 5 *lira* and claiming the multiplier was 1), but it may be beneficial if it lessens investor suspicions of trustee self-dealing.

The upper part of Panel C of Table 4 shows the frequency with which *NoInfo-Chat* trustees skim returns rather than take the maximum possible. These results suggest that skimming is the norm for *NoInfo-Chat* trustees when making an OPR – the trustee takes the maximum possible in only 12.5% of cases. Particularly noteworthy is the fact that *all* of the OPRs where the realized multiplier equals 5 are ones where less than the maximum is taken. We calculate the percentage skim as the difference between the actual and implicit multipliers divided by the difference between the actual multiplier and its minimum of 1. The mean/median/mode percentage skim equals 48.6%/50.0%/50.0%, which suggests that *NoInfo-Chat* trustees extract a private gain when making an OPR equal to about half of the maximum possible.

NoInfo-Chat trustees cluster into three groups based on OPR frequency, which is consistent with behavior documented in other studies (Gneezy, Rockenbach, and Serra-Garcia 2013) – see Panel D of Table 4. Some are routinely "dishonest" during the experiment – e.g.,

[17] Applying a two-tailed Mann-Whitney test to the changes in OPR percentage for the eight *NoInfo-Chat* economies (mean = +3.7%) and the eight *HardInfo-Chat* economies (mean = -9.7%) rejects the null of no difference at $p \leq 0.05$ based on a two-tailed test.

seven trustees make two or more OPRs in each of their dyads over the course of all ten periods. These instances likely reflect self-dealing since their investor partners generally make large investments (average of 95% of endowment invested). Three *NoInfo-Chat* trustees are "honest" in that they never make an OPR. Since these trustees also receive high investments from their partners (average of 98% of endowment), the main difference between these two groups is that the dishonest trustees exploit their informational advantage repeatedly (OPRs in 43% of all periods) to obtain a greater share of the profits. Six "other" trustees receive lower investments (mean of 75%), which suggests that these trustees' less frequent OPRs may in part reflect acts to punish low investment.

We now consider how communication through chat differs for *NoInfo-Chat* trustees. *NoInfo-Chat* trustees may use chat to make potentially misleading multiplier disclosures that could lead directly to false investor beliefs about the trustee's behavior in dividing resources. Examining this issue with a test comparing *NoInfo-Chat* and *HardInfo-Chat* trustees is uninformative since hard information provision renders the disclosure of specific multiplier realizations by *HardInfo-Chat* trustees during chat superfluous. Thus, we focus on how multiplier disclosures during chat sessions by *NoInfo-Chat* trustees vary with the level of the realized multiplier. We expect that a *NoInfo-Chat* trustee will be more likely to make no specific disclosure or provide an inaccurate disclosure when the realized multiplier is higher.

During any given chat session, a *NoInfo-Chat* trustee chooses between three types of multiplier disclosures: (1) no statement about the realized multiplier or a statement that includes no information about the specific value realized, (2) an accurate statement about the realized multiplier, or (3) an inaccurate statement about the realized multiplier.

The following messages apply to a case where a *NoInfo-Chat* trustee makes an accurate multiplier disclosure:

Trustee: I got exactly 7 from you. Nothing multiplied. But it was more generous!
Investor: what?
Investor: you got only 7?
Investor: ok so it got multiplied by one
Trustee: You sent 7 right?
Trustee: Yeah, A1 sent me 6 and it was also only multiplied by 1.
Investor: yes, i sent 7
Trustee: Yeah... times one.
Investor: I'll sent 7 this time again

In contrast, here is a case where the specific disclosure by a *NoInfo-Chat* trustee is inaccurate:

Trustee: no luck this time multiplier 1
Investor: It's OK
Investor: I trust you
Investor: 'And now we have a mutual trust exercise' - Ali G indahouse
Investor::)
Investor: I sent this to the other receiver too
Trustee: perfect
Trustee: trust is the most important thing

This disclosure is inaccurate in that the trustee explicitly states that the realized multiplier was one when in fact it was equal to three. The disclosure is even more misleading because the trustee returned 10 *lira*, which possibly led the investor to falsely believe that the trustee behaved altruistically by keeping nothing for herself.

For each of the eight *NoInfo-Chat* economies, we first calculated the percentage of cases where the trustee either made no multiplier disclosure or made an inaccurate multiplier disclosure when the realized multiplier equaled 1. The top row of Table 5 shows that the average percentage of non-disclosure for the eight *NoInfo-Chat* economies is 36.4% when the realized multiplier equals 1. We then repeated this for multipliers of 2, 3, 4, and 5 for each *NoInfo-Chat* economy. The next row in Table 5 shows analogous average percentages of inaccurate disclosures. The third row adds these together to show the combined frequency of non- and inaccurate disclosure for differing realized multipliers.

The evidence in Table 5 confirms that the combined frequency of non-disclosure and inaccurate disclosure by *NoInfo-Chat* trustees is generally increasing in the realized multiplier level.[18] When the realized multiplier equals 2, the mean economy-wide percentage of chat periods where no multiplier disclosure or an inaccurate disclosure is made equals 54.1% whereas the same calculation for multiplier=5 shows a frequency of 76.0%. To conduct a significance test, the final row aggregates the data across multiplier levels of 1 or 2 versus 4 or 5. A one-tailed Mann-Whitney test comparing the eight economy-wide mean frequencies of inaccurate and non-disclosure for multipliers of 1 or 2 compared to multipliers of 4 or 5 is significant at $p \leq 0.001$.

We next evaluate whether the findings in Table 5 can be linked more directly with *NoInfo-Chat* trustees' OPR decisions and other potential cues of deception in *NoInfo-Chat* trustees' chat messages (Vrij 2008, 101-114). We focus on the chat messages of "dishonest" *NoInfo-Chat* trustees during OPR periods. As benchmarks, we compare these messages with those sent by dishonest trustees in non-OPR periods as well as all messages sent by trustees classified as either "honest" or "other."

The first column with numbers in Table 6 shows various measures characterizing dishonest *NoInfo-Chat* trustees' chat messages in OPR periods. The second and third columns show the same measures for dishonest *NoInfo-Chat* trustees' chat messages in non-OPR periods and all messages sent by trustees classified as either "honest" or "other." We examine three broad constructs related to direct communication in chat messages. The first is the nature of the specific multiplier disclosures by *NoInfo-Chat* trustees during OPR periods. These measures are analogous to those reported in Table 5, and are shown in rows (1) and (2) in Panel A of Table 6.

[18] Consistent with the notion that inaccurate disclosure supports trustee decisions to extract greater private gains, every inaccurate disclosure except one was lower than the realized multiplier. This one exception appears to be a mistake by the trustee that generated multiple messages in the same period to correct.

A second construct is message complexity with relevant measures shown in rows (3) – (8) of Panel B. The third construct applies to the self-referential nature of chat messages as reflected in the use of first-person pronouns – see Panel C.

Three considerations influenced our choice of the specific measures reported in Table 6. First, we sought measures that did not require coding by third parties. The fact that some of our trustee subjects were likely trying to deceive their trading partners implied that the meaning of such messages would not be evident to a third-party coder who was blind to the study's hypotheses.[19] Second, many of the messages involved the use of text message language that makes interpretation of individual messages more difficult (e.g., some people may not know that "IMHO U R Gr8" means "in my humble opinion, you are great").[20] Third, there have been numerous studies on verbal cues of deception, but that literature has produced weak and sometimes conflicting findings (Vrij 2008, 112-114). In combination, these factors led us to use more objective and quantifiable measures that would likely be less subject to measurement error.

Panel A of Table 6 shows two measures of specific multiplier disclosure frequency for dishonest *NoInfo-Chat* trustees that correspond to those in Table 5. Rows (1) and (2) in Table 6 indicates that dishonest *NoInfo-Chat* trustees provide a specific multiplier disclosure in only 7.1% of their OPR periods and *all* of these disclosures are inaccurate. These same trustees make a specific multiplier disclosure in 22.2% of non-OPR periods and only 3.6% of these disclosures are inaccurate. The honest and other trustees provide multiplier disclosures in 57.1% of all chat periods. These results suggest that the sample-wide bias favoring non-disclosure and inaccurate disclosure for larger multipliers in Table 5 is mainly due to those trustees classified as dishonest.

[19] Using third-party coders who are uninformed of the study's hypotheses is used to code chat data in economic experiments (Cooper and Kagel 2005; Sutter and Strassmair 2009; Douthit, Kearney, and Stevens 2012; Bowlin, Hobson, and Piercey 2013).
[20] This example is taken from Farina and Lyddy (2011).

Panel B of Table 6 shows various measures that proxy for message complexity. Prior research hypothesizes that liars tend to communicate in ways that require less cognitive effort to lessen the odds that they will be tripped up in telling a lie (Newman et al. 2003). This suggests that *NoInfo-Chat* trustees will likely send fewer (and shorter) messages during OPR periods. Frequencies in row (3) suggest that *NoInfo-Chat* trustees send fewer messages in OPR periods than in non-OPR periods, but this difference is not large. Also, trustees classified as either "honest" or "other" send fewer messages than dishonest trustees. Row (4) suggests that *NoInfo-Chat* trustees do not send longer messages (as measured by mean character count) in OPR periods than in non-OPR periods.

The other measures of complexity attempt to measure avoidance of certain topics as a means for a lying trustee to simplify communication. The measure of complexity in row (5) is based on the notion that a trustee lying about the realized multiplier will avoid use of words that contain the term "mult." The data are strongly consistent with this conjecture. Dishonest *NoInfo-Chat* trustees use words containing "mult" in only 15.1% of OPR periods, which is less than half of that observed for the same trustees in non-OPR periods (35.6%) and only slightly over half the rate for honest and other trustees (29.6%).

Our next measure of topic avoidance is based on the idea that personalization of the investor-trustee relationship can allow the trustee to steer the conversation away from her resource division decisions. We expect that chat transforms impersonal market interactions into personalized exchange (e.g., McCloskey and Klamer 1995; Smith 2008), which might provide a partial explanation for our results indicating large investments in *NoInfo-Chat* economies despite high frequencies of trustee self-dealing. We expect that trustees can use personalized exchange as a device to distract investors with extraneous talk during chat sessions.

Here is an example from our sample that illustrates a case of extraneous conversation. This example takes place in the chat session after period 8. The investor has made the maximum investment in every period 2 – 8 inclusive, the multiplier draw was equal to 5 in all three periods 6 – 8, and the trustee made OPRs with returns of 15, 10, and 15 units in periods 6, 7, and 8, respectively. The following messages comprise the entire chat session following period 8:

> Trustee: do you think that internet dating is for real... i do know one co-worker who met his gf through lavalife
> Investor: i dunno, those eharmony commercials are pretty encouraging
> Investor: dr. phil can't be wrong
> Trustee: dr. phil is a schmuck
> Investor: whoa, back up
> Investor: okay i'm just kidding, he's a tool
> Investor: you know who's not a tool? oprah!!
> Trustee: uhoh, did i just reveal too much personal info by using the word schmuck
> Investor: haha
> Investor: you're in trouble now
> Trustee: well, minorities are cool
> Investor: yeah... high five
> Investor: (trying not to reveal to much personal info, haha)
> Investor: too much
> Investor: there
> Investor: im not dumb

One interesting aspect of this conversation is the use of the term "haha" to denote a response to humor, which is likely an indication of a personal (as opposed to business) conversation. We define a chat message to include a response to humor if one of the following terms appeared in the message: "lol," "hehe," "haha," or "fun."[21] Dishonest trustees make use of humor more often in OPR periods (34.0%) than in non-OPR periods (23.3%), both of which exceed the rate for trustees classified as either honest or other (19.1%).

Another device that might be indicative of a personalized relationship is use of an emoticon, which is demonstrated in **bold** in the following messages between the same two subjects in the chat session after period 5:

[21] "lol" is a text message meaning "laugh out loud."

Trustee: I actually was contenplating screwing myself that time and giving you 15 just to make you feel better
Trustee: should never have done psych... **;)**
Investor: haha that is so nice
Investor: what would make me feel better is a 5x... if you want

Rows (7) and (8) of Table 6 show that the frequency of positive emoticons (")") is much higher than the use of negative emoticons ("("). The frequency of positive emoticon use by dishonest trustees is the same across OPR and non-OPR periods (15.1%), but this frequency is dramatically higher than for trustees classified as either honest or other.

Row (9) shows the frequency of first-person singular pronoun usage in chat messages. As a benchmark, this panel also shows the frequencies with which other pronouns are used. Prior research finds that persons engaged in deception tend to use fewer self-references as they try to distance themselves from the false statements they make (Vrij 2008, 101-8; Newman et al. 2003; Larcker and Zakolyukina 2012). If this applies to our sample, then usage of first person singular pronouns by dishonest *NoInfo-Chat* trustees will be lower in OPR periods than in non-OPR periods. This in fact occurs: dishonest *NoInfo-Chat* trustees use either "I" or "me" in 54.7% of OPR periods compared to 64.4% of non-OPR periods and 72.2% by honest and other trustees. No similar pattern is apparent for the use of first-person plural pronouns ("we" or "us") and third-person plural pronouns ("they" or "them"), but a similar pattern holds for the use of second-person singular pronouns ("you" or "u").

Another way that trustees can build trust with their investor partners is use chat for making promises about their future resource division decisions. Prior research suggests it is plausible that trustees would make promises to investors to better coordinate behavior and build trust (Charness and Dufwenberg 2006). The following messages illustrate such a trustee promise:

Trustee: if you continue to send me the maximum, I'll continue to send you exactly half of the total earnings

Investor: good i think its better to send the maximum its good for both of us because of the multiplier between a and b
Trustee: agreed
Investor: good

Table 7 shows that *NoInfo-Chat* and *HardInfo-Chat* trustees frequently make return promises.[22] The typical trustee promise is one where the trustee agrees to return half *if* the investor makes a maximum investment ("invest all and split 50-50"). We define compliance as a trustee return that is greater than or equal to that which is promised. Not surprisingly, Table 7 indicates that compliance rates are higher in *HardInfo-Chat* economies (96.1%) than in *NoInfo-Chat* economies (74.7%) at $p < 0.001$ for a one-tailed Mann-Whitney test.

Our analyses in this sub-section have shown that *NoInfo-Chat* trustees engage in various actions that could cumulatively serve to conceal their self-dealing from their investor partners. These actions include returns that preserve opacity over how trustees divided resources as well as various forms of potentially manipulative communication during chat sessions. We conducted a final set of analyses to evaluate whether these actions cumulatively imply that *NoInfo-Chat* investors earn markedly lower rates of return earned than *HardInfo-Chat* investors.

We calculate the rate of return on endowment earned by investor i with trustee j in period t as the sum of the amount returned by the trustee plus any amount that the investor did not invest less the initial endowment of ten *lira* all divided by the initial endowment of ten *lira*. The economy-wide mean ROE for each of the eight *NoInfo-Chat* economies during all periods is calculated by taking the simple average of ROE across all 40 dyad-period observations (10 periods X 4 dyads). We then repeated these calculations for each *HardInfo-Chat*, *HardInfo*, and *NoInfo* economy. Finally, we also performed these calculations using only returns data for

[22] Four of the five cases where *HardInfo-Chat* trustee promises were not made occurred in a single session. This suggests that trustee promises are not independent across dyads.

periods 1-5 and 6-10 (20 dyad-period observations for each economy = five periods X four dyads).

Panel A of Table 8 indicates that the mean (median) of the economy-wide average per-period ROE across the eight *NoInfo-Chat* economies equals 30.3% (32.8%), which is somewhat lower than the mean (median) of 32.5% (40.0%) for the *HardInfo-Chat* economies. However, this difference in ROE across the two types of economies is not statistically distinguishable from zero. This result is likely related to our finding in Panel C of Table 4 that *NoInfo-Chat* trustees skim profits when making opacity-preserving returns. A similar result holds in a comparison between the *NoInfo* and *HardInfo* economies (means of 11.2% and 13.3%, respectively), even though investor ROEs in both the *NoInfo* and *HardInfo* economies are substantially lower than in the *NoInfo-Chat* and *HardInfo-Chat* economies.

An interesting pattern of results is evident when these data are partitioned by sub-period. Panel B of Table 8 shows that the average ROE earned by *NoInfo-Chat* investors in periods 1–5 (34.4%) is substantially higher than that earned in periods 6–10 (26.3%). This is the only case where ROE is higher in periods 1-5 than 6-10. *HardInfo-Chat*, *NoInfo*, and *HardInfo* investors earn greater returns in periods 6-10 than in periods 1-5 (increases of 9.3%, 15.0%, and 4.5%, respectively). These findings are consistent with *NoInfo-Chat* trustees attempting to build false trust early (by delivering higher returns) that can be later exploited for private gain.

Overall, the results in this sub-section provide strong support for the view that *NoInfo-Chat* trustees craft complex strategies that entail multiple actions consistent with an intent to conceal self-dealing. These actions include opacity-preserving returns, misleading communication during chat, and actions to build trust that can be subsequently exploited for personal gain (e.g., return promises and large early returns).

CONCLUSIONS AND SUGGESTIONS FOR FUTURE RESEARCH

We have reported results from an economic experiment that allows individuals to endogenously craft strategies for self-dealing and possibly conceal it from potential victims. Our experiment is built from the simple investment game presented by Berg, Dickhaut, and McCabe (1995) and incorporates multiple periods, multiple dyadic interactions, information asymmetry, and uncertainty as in King-Casas et al. (2005), Basu et al. (2009), and Lunawat (2012, 2013). We modify this game by allowing subjects to directly communicate through a chat institution at the end of each period.

Trustees in this *NoInfo-Chat* economy must persuade investors to commit resources to investment that may generate benefits that the two parties can share. We contrast the performance of this economy with others where the trustee has the option of revealing hard information that renders her behavior fully transparent. The experimental structure and manipulations in this paper circumvent a conundrum faced by researchers seeking to study self-dealing and deception with field data: how can we learn about behavior where success requires that it be concealed?

Two main conclusions present themselves based on analysis of our experimental data. The first is that the strategies fashioned by self-dealing trustees are complex and involve several features. These strategies include resource-sharing decisions that preserve the opacity that surrounds such behavior as well as potentially misleading communication and actions to build unwarranted trust that can be subsequently exploited for private gain.

A second, and more important, conclusion is that an economy can still generate sizable social gains even when self-dealing and deceptions lurk beneath the surface of seemingly cordial, mutually beneficial exchange interactions. This of course raises the question of: where precisely is the boundary between successful exchange and market failure due to suspicion of self-dealing or even fraud?

Additional research can profitably follow paths that go beyond the basic experiment we consider here. This work might explore the effects of changing the payoff structure to make it more realistic – e.g., by introducing bankruptcy due to negative outcomes that reduce or even eliminate a subject's wealth. How would the possibility of losing all wealth alter how investor subjects perceive the risks of self-dealing by trustees? Within such experiments, it would be interesting also to examine punishment institutions applied to individuals who go bankrupt.

Our view in this paper has mainly been from the perspective of a self-dealing trustee. But, we could have approached it more from the victim's perspective. Self-dealing and fraud are reciprocal problems in much the same way are pollution and trespass (Coase 1960). From this stance, incentives to protect against and deter fraud are as pertinent as incentives to commit fraud and conceal it. Useful experimental manipulations might lie in allowing investors to choose when and how the trustee provides information. This would enable an examination of the effect of competition among various trustees to secure investments from investors. Since this is likely an important force in constraining self-dealing in markets. Ultimately, such research could yield new insights into how investor protection institutions emerge and evolve in response to detected frauds.

References

Arya, A., J. Glover, and S. Sunder, 2003. Are unmanaged earnings always better for shareholders? *Accounting Horizons* **17** (Supplement): 111-116.

Arya, A., J. Glover, B. Mittendorf, and L. Zhang, 2004. The disciplining role of accounting in the long-run. *Review of Accounting Studies* **9**: 399-417.

Baiman, S. and J. H. Evans, 1983. Pre-decision information and participative management control systems. *Journal of Accounting Research* **21** (2): 371-395.

Ball, R., 2001. Infrastructure requirements for an economically efficient system of public financial reporting and disclosure. *Brookings-Wharton Papers on Financial Services*, 127-169.

Basu, S. and G. Waymire, 2006. Recordkeeping and human evolution. *Accounting Horizons* **20** (3): 1-29.

Basu, S., J. Dickhaut, G. Hecht, K. Towry, and G. Waymire, 2009. Recordkeeping alters economic history by promoting reciprocity. *Proceedings of the National Academy of Sciences of the USA* **106** (4): 1009–1014.

Belot, M. and J. van de Ven, 2013. How private is private information? The ability to spot deception in an economic game. Unpublished working paper, University of Edinburgh, December.

Ben-Nera, A., L. Putterman, and T. Ren, 2011. Lavish returns on cheap talk: two-way communication in trust games. *Journal of Socio-Economics* **40**: 1-13.

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

Bochet, O., T. Page, and L. Putterman, 2006. Communication and punishment in voluntary contribution experiments. *Journal of Economic Behavior and Organization* **60**: 11-26.

Bowlin, K., J. Hobson, and D. Piercey, 2013. The effects of auditor rotation, professional skepticism, and interactions with managers on audit quality. Unpublished working paper, University of Mississippi, February.

Byrne, R. and A. Whitten, 1997. Machiavellian intelligence. In A. Whiten and R. Byrne (Eds.) *Machiavellian Intelligence II* 1–23, Cambridge, UK: Cambridge University Press.

Cai, H. and J. Wang, 2006. Overcommunication in strategic information transmission games. *Games and Economic Behavior* **56**: 7-36.

Charness, G. and M. Dufwenberg, 2006. Promises and partnership. *Econometrica* **74** (6): 1579-1601.

Charness, G. and B. Grosskopf, 2004. What makes cheap talk effective? Experimental evidence. *Economics Letters* **83**: 383-389.

Christensen, J., 1981. Communication in agencies. *The Bell Journal of Economics* **12** (2): 661-674.

Coase, R., 1960. The problem of social cost. *Journal of Law and Economics* **3**: 1-44.

Cohen, T., T. Wildschut, and C. Insko, 2010 How communication increases interpersonal cooperation in mixed-motive situations. *Journal of Experimental Social Psychology* **46**: 39-50.

Cooper, D. and J. Kagel, 2005. Are two heads better than one? Team versus individual play in signaling games. *American Economic Review* **95** (3): 477-509.

Coricelli, G., L. Morales, and A. Mahlstedt, 2006. The investment game with asymmetric information. *Metroeconomica* **57** (1): 13-30.

Croson, R. T. Boles, and K. Murnighan, 2003. Cheap talk in bargaining experiments: lying and threats in ultimatum games. *Journal of Economic Behavior and Organization* **51**:143-159.

Crawford, V., 1998. A survey of experiments on communication via cheap talk. *Journal of Economic Theory* **78**: 286-298.

Dawes, R., J. McTavish, and H. Shaklee, 1977. Behavior, communication, and assumptions about other people's behavior in a commons dilemma situation. *Journal of Personality and Social Psychology* **35** (1): 1-11.

Dickhaut, J., S. Basu, K. McCabe, and G. Waymire, 2010. NeuroAccounting: consilience between the biologically-evolved brain and culturally-evolved accounting principles. *Accounting Horizons* **24** (2): 221-255.

Douthit, J., L. Kearney, and D. Stevens, 2012. Can agent cheap talk mitigate agency problems in the presence of a noisy performance measure? An experimental test in a single- and multi-period setting. *Journal of Management Accounting Research* **24**: 135-158.

Duffy, J. and N. Feltovich, 2002. Do actions speak louder than words? An experimental comparison of observation and cheap talk. *Games and Economic Behavior* **39**: 1-27.

Dunbar, R. and S. Shultz, 2007. Evolution in the social brain. *Science* **317**: 1344-1347.

Dye, R., 1983. Communication and post-decision information. *Journal of Accounting Research* **21** (2): 514-533.

Ekman, P. 2003 (Ed.). *Darwin and Facial Expressions*. Cambridge, MA: Malor Books.

Ellingsen, T. and M. Johannesson, 2004. Promises, threats, and fairness. *The Economic Journal* **114** (495): 397-420.

Falk, A. and J. Heckman, 2009. Lab experiments are a major source of knowledge in the social sciences. *Science* **326**: 535-538.

Farina, F. and Lyddy, F., 2011. The language of text messaging: "linguistic ruin" or resource? *The Irish Psychologist* **37** (6): 145-149. (available online at http://eprints.may.ie/view/creators/Lyddy=3AFiona=3A=3A.default.html)

Fischbacher. U., 2007. Z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* **10**: 171-178.

Gigler, F. and T. Hemmer, 1998. On the frequency, quality, and informational role of mandatory financial reports. *Journal of Accounting Research* **36** (Supplement): 117-147.

Gneezy, U., 2005. Deception: the role of consequences. *American Economic Review* **95** (1): 384-394.

Gneezy, U., B. Rockenbach, and M. Serra-Garcia, 2013. Measuring lying aversion. *Journal of Economic Behavior & Organization* **93**: 293-300.

Hales, J., J. Kuang, and S. Venkataraman, 2011. Who believe the hype? An experimental examination of how language affects investor judgments. *Journal of Accounting Research* **49** (1): 223-255.

Henriques, D., 2012. *The Wizard of Lies*. New York: St. Martin's Griffin.

Hobson, J., W. Mayew, and M. Venkatachalam, 2012. Analyzing speech to detect financial misreporting. *Journal of Accounting Research* **50** (2): 349-392.

Ijiri, Y., 1975, *Theory of Accounting Measurement: Studies in Accounting Research # 10*, Sarasota, FL: American Accounting Association.

King, R., 2002. An experimental investigation of self-serving biases in an auditing trust game: the effect of group affiliation. *The Accounting Review* **77** (2): 265-284.

King-Casas, B., D. Tomlin, C. Anen, C. Camerer, S. Quartz, and P. R. Montague, 2005. Getting to know you: reputation and trust in a two-person economic exchange. *Science* **308**: 78-83.

Kriss, P., R. Nagel, and R. Weber, 2013. Implicit vs. explicit deception in ultimatum games with incomplete information. *Journal of Economic Behavior & Organization* **93**: 337-346.

Larcker, D. and A. Zakolyukina, 2012. Detecting deceptive discussions in conference calls. *Journal of Accounting Research* **50** (2): 495-539.

Ledyard, J., 1995. Public goods: a survey of experimental research. In J. Kagel and A. Roth (Eds.), *The Handbook of Experimental Economics* (pp. 111-194). Princeton, NJ: Princeton University Press.

Lunawat, R., 2012. The role of information in building reputation in an investment-trust game. *European Accounting Review* **22** (3): 513-532.

Lunawat, R., 2013. An experimental investigation of reputation effects of disclosure in an investment/trust game. *Journal of Economic Behavior & Organization* **94**: 130-144.

McCloskey, D. and A. Klamer, 1995. One quarter of GDP is persuasion. *American Economic Review* **85** (2): 191-195.

Newman, M., J. Pennebaker, D. Berry, and J. Richards, 2003. Lying words: predicting deception from linguistic styles. *Personality and Social Psychology Bulletin* **29** (5): 665-675.

Ochs, J., 1995. Coordination problems. In *The Handbook of Experimental Economics* (eds J. Kagel & A.E. Roth), pp. 195-252. Princeton, NJ: Princeton University Press.

Sanchez-Pages and Vorsatz 2007. An experimental study of truth-telling in a sender-receiver game. *Games and Economic Behavior* **61**: 86-112.

Schniter, E., R. Shermata, and D. Sznycer, 2013. Building and rebuilding trust with promises and apologies. *Journal of Economic Behavior and Organization* **94**: 242-256.

Smith, V., 2004. Human nature: an economic perspective. *Daedalus* **133** (4): 67-76.

Smith, V., 2008. *Rationality in Economics*. Cambridge, UK: Cambridge University Press

Sutter, M., 2009. Deception through telling the truth?! Experimental evidence from individuals and teams. *The Economic Journal* **119** (4): 47-60.

Sutter, M. and C. Strassmair, 2009. Communication, cooperation and collusion in team tournaments – an experimental study. *Games and Economic Behavior* **66**: 506-525.

Trivers, R., 1971. The evolution of reciprocal altruism. *The Quarterly Review of Biology* 46: 35-57.

Vanberg, C., 2008. Why do people keep their promises? An experimental test of two explanations. *Econometrica* **76** (6): 1467-1480.

Vrij, A., 2008. *Detecting Lies and Deceit*. West Sussex, UK: John Wiley and Sons.

FIGURE 1
Multi-dyad investment-trust game with uncertainty

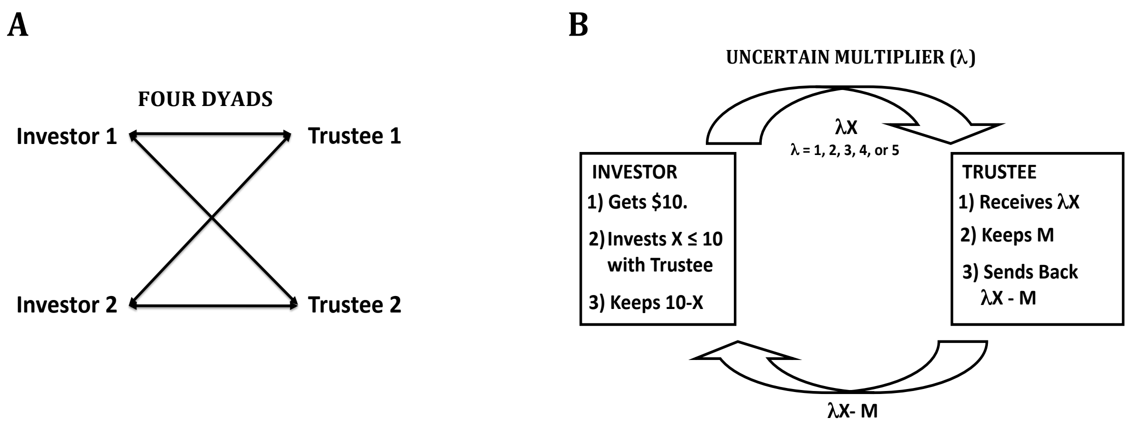

***Panel A*:** Each economy consists of four persons, two of whom are assigned to be investors and two to be trustees. Each investor receives an endowment of ten monetary units to invest with Trustee 1 and another ten units to invest with Trustee 2 each period.
***Panel B*:** Each period of play begins by endowing the investor with ten monetary units that can be sent to the trustee in whole numbers ranging from zero to ten units. Any amount invested is subject to multiplication at a rate λ that ranges from 1 to 5 in whole numbers. The trustee receives an amount equal to λ times investment (X), and then decides how much to keep (M) and how much to return to the investor (λX - M). (C) The investor does not know the multiplier or the amount retained by the trustee. The trustee likewise does not directly observe the realized multiplier or the amount invested by the investor although that can be inferred in many cases.

FIGURE 2
Percent of endowment invested in alternative economies by period

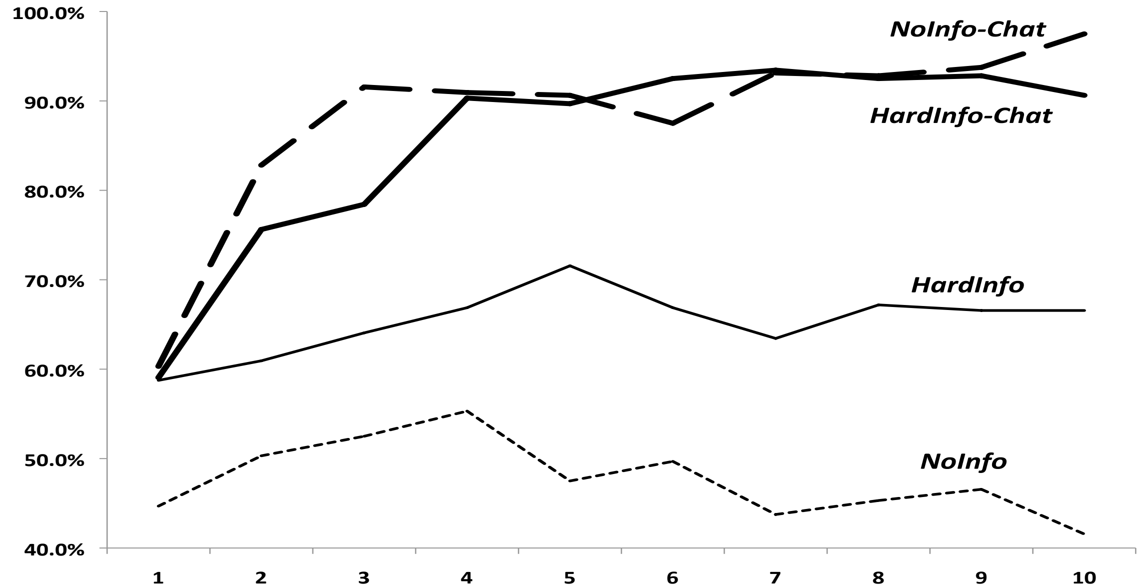

This figure plots the average percentage of endowment invested each period by investors in four different types of economies (*NoInfo-Chat*, *HardInfo-Chat*, *HardInfo*, and *NoInfo*). The average percentage of endowment invested in any period for a given type of economy is the mean of the ratio of *lira* invested divided by the endowment of ten *lira* calculated using all 32 available observations (two investors X two dyads X eight replications) for a given period (t=1, 2, 3…10).

FIGURE 3
Effect of differing realized multipliers on % sent back by trustee by period

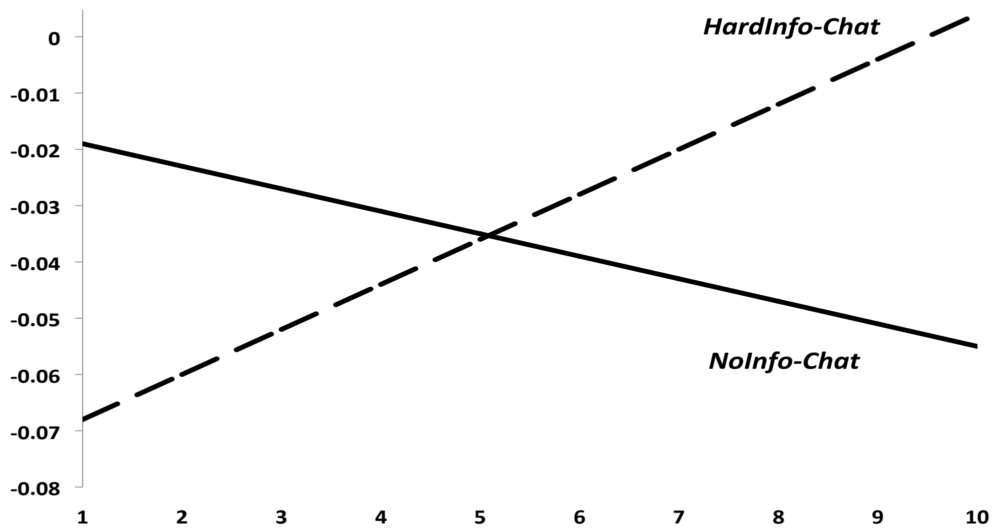

This figure plots how the realized multiplier affects the percent returned by the trustee as PERIOD increases from 1 to 10 for the *NoInfo-Chat* and *HardInfo-Chat* economies. These estimates are derived from coefficient estimates based on the following model:

$$\%RET_{ijt} = \alpha + \beta_1 INV_{ijt} + \beta_2 MULT_{ijt} + \beta_3 PERIOD_t + \beta_4 MULT_{ijt} PERIOD_t + \beta_5 \mathit{NoInfo\text{-}Chat}_{ij} + \beta_6 INV_{ijt} \mathit{NoInfo\text{-}Chat}_{ij} + \beta_7 MULT_{ijt} \mathit{NoInfo\text{-}Chat}_{ij} + \beta_8 PERIOD_t \mathit{NoInfo\text{-}Chat}_{ij} + \beta_9 MULT_{ijt} PERIOD_t \mathit{NoInfo\text{-}Chat}_{ij} + \varepsilon_{ijt}$$

$\%RET_{ijt}$ is the percentage returned to investor i by trustee j in period t, INV_{ijt} equals the level of investment by investor i with trustee j in period t, $MULT_{ijt}$ is realized multiplier in period t for the dyad of investor i and trustee j, $PERIOD_t$ is the period of the experiment, and *NoInfo-Chat*$_{ij}$ is a 0-1 variable equal to 1 for *NoInfo-Chat* economies and 0 for *HardInfo-Chat* economies. The model is estimated using OLS with standard errors clustered by trustee applied to the full sample excluding nine observations where the trustee did not make a return because the investor had invested zero *lira*. The specific estimates shown in the figure equal $[\beta_2 + (\beta_4 * PERIOD)]$ for *HardInfo-Chat* economies and $[(\beta_2 + \beta_7 + ((\beta_4 + \beta_9) * PERIOD)]$ for *NoInfo-Chat* economies.

TABLE 1
Frequency of Chat Messages and Provision of Hard Information for Different Economies

	NoInfo-Chat Economies	*HardInfo-Chat* Economies	*HardInfo* Economies
# Chat Messages Per Period			
Trustee – Mean (Median) # Messages	4.4 (4)	3.9 (3)	NA
Investor – Mean (Median) # Messages	3.8 (3)	3.4 (3)	NA
Chat Message Character Count			
Trustee – Mean (Median)	36.2 (27)	31.5 (24)	NA
Investor – Mean (Median)	35.7 (28)	33.7 (27)	NA
Hard Information Provision			
% Periods Where Trustee Provides Hard Information	NA	87.2%	73.4%

TABLE 2
Summary Statistics on Investment and the Distribution of Wealth in Different Economy Types

A: *Economy-wide average per-period % of endowment invested*

	NoInfo-Chat	*HardInfo-Chat*
Mean	88.1%	85.5%
Median	92.9%	87.8%
p-value Mann-Whitney (two-tailed)	ns at p ≤ 0.05	

	HardInfo	*NoInfo*
Mean	65.3%	47.7%
Median	65.0%	48.0%
p-value Mann-Whitney (two-tailed)	0.028	

B: *Economy-wide average per-period % of total wealth captured by the trustee*

	NoInfo-Chat	*HardInfo-Chat*
Mean	50.3%	47.4%
Median	50.7%	47.7%
p-value Mann-Whitney (two-tailed)	ns at p ≤ 0.05	

	HardInfo	*NoInfo*
Mean	43.9%	34.3%
Median	43.7%	31.8%
p-value Mann-Whitney (two-tailed)	ns at p ≤ 0.05	

TABLE 3
Characteristics of Trustee Resource Sharing Decisions
A: *Summary Statistics on % Sent Back by Trustee to Investor*

	HardInfo-Chat	*NoInfo-Chat*
Mean Economy-Wide Average % Sent Back	47.5%	45.8%
Median Economy-Wide Average % Sent Back	49.3%	46.0%
Total % of All Cases < 50%	24.9%	34.1%
Total % of All Cases = 50%	65.1%	57.8%
Total % of All Cases > 50%	10.0%	8.0%

B: *Regression Analysis of the Sensitivity of NoInfo-Chat Trustee Returns to Realized Multipliers*

Independent Variable	*Coefficient (t-statistic)*
Constant	0.590 (4.74)
Investment	0.013 (2.31)
Multiplier	-0.076 (-2.86)
Period	-0.024 (-2.29)
Multiplier*Period	0.008 (2.47)
NoInfoChat	-0.181 (-1.19)
Investment* NoInfoChat	-0.001 (-0.16)
Multiplier* NoInfoChat	0.061 (2.08)
Period* NoInfoChat	0.033 (2.50)
Multiplier*Period* NoInfoChat	-0.012 (-2.79)
N	631
Adjusted R^2	0.151

Panel A: The first two rows show the mean and median of the economy-wide average percent sent back by trustees in the *NoInfo-Chat* and the *HardInfo-Chat* economies. The percent sent back by trustee j in period t is the amount sent by the trustee to investor i divided by the total received by the trustee (equal to investment times the realized multiplier). The economy-wide average of this measure is the simple average across all dyad-period observations for a particular run of the economy. The three bottom rows show the economy-wide average percentage of cases where the trustee sends back an amount less than versus greater than or equal to half of what the trustee received.

Panel B: This panel provides coefficient estimates based on the following model:
$\%RET_{ijt} = \alpha + \beta_1 INV_{ijt} + \beta_2 MULT_{ijt} + \beta_3 PERIOD_t + \beta_4 MULT_{ijt} PERIOD_t + \beta_5 \textit{NoInfo-Chat}_{ij} + \beta_6 INV_{ijt} \textit{NoInfo-Chat}_{ij} + \beta_7 MULT_{ijt} \textit{NoInfo-Chat}_{ij} + \beta_8 PERIOD_t \textit{NoInfo-Chat}_{ij} + \beta_9 MULT_{ijt} PERIOD_t \textit{NoInfo-Chat}_{ij} + \varepsilon_{ijt}$
$\%RET_{ijt}$ is the percentage returned to investor i by trustee j in period t, INV_{ijt} equals the level of investment by investor i with trustee j in period t, $MULT_{ijt}$ is realized multiplier in period t for the dyad of investor i and trustee j, $PERIOD_t$ is the period of the experiment , and $\textit{NoInfo-Chat}_{ij}$ is a 0-1 variable equal to 1 for *NoInfo-Chat* economies and 0 for *HardInfo-Chat* economies. The model is estimated using OLS with standard errors clustered by trustee applied to the full sample excluding nine observations where the trustee did not make a return because the investor had invested zero *lira*.

Table 4
Incidence of Opacity Preserving Returns (OPR) in *NoInfo-Chat* and *HardInfo-Chat* Economies

A: *Mean of Economy-Wide Percentage of OPRs*

NoInfo-Chat (n=8)	22.6% (n = 72 total)
HardInfo-Chat (n=8)	7.1% (n = 22 total)
Mann-Whitney p-value	0.005

B: *Economy-Wide Average Percentage of OPRs by Sub-Periods for NoInfo-Chat and HardInfo-Chat Trustees*

	Periods 1-5	*Periods 6-10*	*Change*
NoInfo-Chat (n=8)	20.7% (n=33)	24.4% (n=39)	+3.7%
HardInfo-Chat (n=8)	11.9% (n=19)	2.2% (n=3)	-9.7%
Mann-Whitney p-value (two-tailed)			0.024

C: Extent of Skim in *NoInfo-Chat* Trustees' Opacity Preserving Returns (n = 72)

Actual Multiplier	*# (%) Cases Where Maximum Taken*
2	3 of 3 (100%)
3	1 of 14 (7.1%)
4	5 of 22 (22.7%)
5	0 of 33 (0%)
ALL	9 of 72 (12.5%)
Mean/Median/Mode % Skim (Max = 100%)	48.6%/50.0%/50.0%

D: *NoInfo-Chat Trustee Classified by the Frequency of OPRs*

	"Dishonest"	*"Honest"*	*Other*
# Trustees	7	3	6
# (%) OPR Periods	60 of 140 (43%)	0 of 60 (0%)	12 of 127 (9%)
# Zero Investment Periods	0	0	3
% OPR Periods	43%	0%	9%
Mean % of Endowment Invested	95%	98%	75%
Mean/Median % of Profits to Trustee	53.9/55.1	49.3/49.0	46.6/46.0

Panel A: The first two rows in this panel show the average economy-wide percentage of opacity preserving returns in *NoInfo-Chat* and the *HardInfo-Chat* economies. An opacity preserving return is one where the trustee returns an amount that implies a feasible multiplier that is lower than that actually realized under the assumption that the investor expects that the trustee to send back half of the total amount she received. The Mann-Whitney p-value applies to a two-tailed test.
Panel B: This panel is equivalent to Panel A except it is presented separately for sub-periods 1-5 vs. 6-10.
Panel C: "Skim" refers to a case where the trustee in making an opacity preserving return takes less than the maximum possible that would still allow her to plausibly deny having returned less than half the amount she received. Skim is calculated as the actual multiplier less the implicit multiplier divided by the difference between the actual multiplier and 1.
Panel D: This panel is based on a categorization of trustees by OPR frequency. "Dishonest" includes seven trustees that make two or more OPRs in both of their dyads and "Honest" includes three trustees that never make an OPR. The remaining six trustees are classified as "Other." Average investment and % of profits earned during the experiment are calculated for each of the 16 trustees separately.

Table 5
Economy-Wide Frequency of *NoInfo-Chat* Trustee Non-Disclosure and Inaccurate Disclosure of Realized Multipliers

	Mult = 1	*Mult = 2*	*Mult = 3*	*Mult = 4*	*Mult = 5*
(1) *Non-Disclosure*	36.4%	47.1%	61.1%	66.5%	63.7%
(2) *Inaccurate Disclosure*	0.0%	7.0%	6.6%	14.1%	12.2%
TOTAL (1) + (2)	36.4%	54.1%	67.7%	80.6%	76.0%

	Mult 1 or 2	*Mult 4 or 5*	*Mann-Whitney p-value*
TOTAL (1) + (2)	44.3%	77.1%	0.001

This table shows the incidence of both non-disclosure and inaccurate disclosure of multipliers during chat periods for *NoInfo-Chat* economies. Frequencies are calculated separately according to the realized multiplier during the period. The Mann-Whitney test is based on the eight specific frequencies for *NoInfo-Chat* economies during periods where the multiplier equaled 1 or 2 compared with the eight analogous frequencies for the same economies in periods where the realized multiplier was equal to 4 or 5.

TABLE 6

Characteristics of Dishonest *NoInfo-Chat* Trustees' Chat Messages in OPR Periods

	Dishonest Trustees in OPR Periods	*Dishonest Trustees in Non-OPR Periods*	*Honest & Other Trustees*
A. SPECIFIC MULTIPLIER DISCLOSURES			
(1) % periods with specific multiplier disclosure	7.1%	22.2%	57.1%
(2) % specific disclosures that are inaccurate	100.0%	3.6%	16.1%
B. COMMUNICATION COMPLEXITY			
(3) Mean # messages	4.13	4.47	4.01
(4) Mean character count per message	32.9	32.3	35.7
(5) % periods with "mult"	15.1%	35.6%	29.6%
(6) % periods with HUMOR ("lol," "hehe," "haha," or "fun")	34.0%	23.3%	19.1%
(7) % periods with positive emoticons – ")"	15.1%	15.1%	3.1%
(8) % periods with negative emoticons – "("	1.9%	4.1%	4.9%
C. SELF REFERENCES			
(9) % periods with "I" or "me" (*first-person singular*)	54.7%	64.4%	72.2%
% periods with "we" or "us" (*first-person plural*)	34.0%	34.2%	30.2%
% periods with "you" or "u" (*second-person singular*)	43.4%	50.7%	63.0%
% periods with "they" or "them" (third-person plural)	18.9%	8.2%	9.9%

Table 7
Frequency of Trustee Return Promises & Compliance Rates

	NoInfo-Chat	*HardInfo-Chat*
# Dyads Where Promise Made	31 of 32 (97%)	23 of 28 (82%)
# “Invest All & Split 50-50”	28 of 31 (90%)	16 of 23 (70%)
Average Economy-Wide % Compliance Rate	74.7%	96.1%
Mann-Whitney p-value (one-tailed)	0.001	

The upper section of this table shows the frequency with which trustees make future return promises to investors and the frequency with which these promises are of the form “invest all & split 50-50.” The lower part shows average economy-wide compliance rate by trustees for *NoInfo-Chat* and the *HardInfo-Chat* economies. The Mann-Whitney test is applied to eight specific values for the two types of economies.

TABLE 8
Summary Statistics on the Average Economy-Wide Return on Endowment for Different Types of Economies Overall and by Sub-Period

A: Economy-Wide Average Per-Period Investor ROE across All Periods

	NoInfo-Chat	*HardInfo-Chat*
Mean	30.3%	32.5%
Median	32.8%	40.0%
Mann-Whitney p-value (two-tailed)	ns at p ≤ 0.05	

	HardInfo	*NoInfo*
Mean	11.2%	13.3%
Median	6.6%	8.5%
Mann-Whitney p-value (two-tailed)	ns at p ≤ 0.05	

B: Economy-Wide Average ROE by Sub-Periods

	Periods 1-5	Periods 6-10	Change
NoInfo-Chat	34.4%	26.3%	-8.1%
HardInfo-Chat	29.9%	39.2%	+9.3%
Mann-Whitney p-value (two-tailed)			0.041
NoInfo	5.8%	20.8%	+15.0%
HardInfo	8.9%	13.4%	+4.5%
Mann-Whitney p-value (two-tailed)			ns at p < 0.10

Panel A: This panel shows the mean investor ROE across the eight *NoInfo-Chat*, *HardInfo-Chat*, *HardInfo-Chat*, and *NoInfo* economies separately. These numbers represent the economy-wide average based on 40 observations per economy (four dyads X 10 periods). The economy-wide average is the simple mean of these eight values for a given type of economy. tor ROE for a given dyad-period observation is the amount sent back by the trustee plus the amount not invested initially divided by the initial endowment less one. The Mann-Whitney test is applied to eight specific values for the two types of economies.
Panel B: This panel is equivalent to Panel A except it is presented separately for sub-periods 1-5 vs. 6-10. The Mann-Whitney test is applied to eight specific values of the change in ROE across sub-periods for the two types of economies.

APPENDIX
Experiment Instructions for *HardInfo-Chat* Economy

Introduction
You have been invited to participate in a decision making experiment. This experiment will last approximately two hours. During today's session, you will earn income in an experimental currency called Lira. At the end of the session, this currency will be converted to dollars at a rate of $0.04 (4 cents) per Lira, and you will be paid in cash. In addition to this income, you will also receive a show-up fee of $10.

Please read these instructions very carefully. You will be required to complete a quiz, in order to demonstrate that you have a complete and accurate understanding of these instructions. After you have completed the quiz, the administrator will check your answers and discuss with you any questions that have been answered incorrectly.

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

Session Overview
This session will be run entirely over the computer. Please do not talk with any of the other participants. If you have a question, you may raise your hand, and the administrator will answer the question privately.

Roles and Procedures
Every participant will be assigned to the role of either an A-player or a B-player. You will be organized into groups of 4 players, consisting of 2 A-players and 2 B-players. Each A-player in the group will be simultaneously paired with the 2 B-players in the group. Similarly, each B-player will be simultaneously paired with the 2 A-players. For example, if the 2 A-players in a group are designated A1-A2, and the 2 B-players are designated B1-B2, the following diagram demonstrates all pairings in the group.

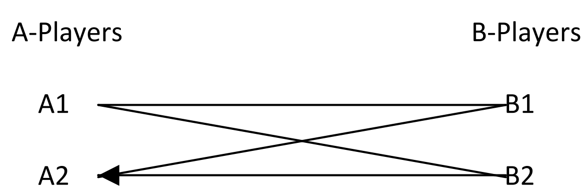

These roles will be completely anonymous. That is, you will know your own role, but you will not know the role of any other participant.

You will be asked to make decisions in a number of identical rounds. Each round proceeds through five stages. In Stage 1, each B-player decides whether private information will get revealed to each A-player. In Stage 2, each A-player receives an endowment of 10 liras per B-player. The A-player then decides how many of the 10 lira to send to each B-player. In Stage 3, the amount sent by each A-player is multiplied. These multiplied amounts are received by each of the B players. B players then decide how much to return to each A player. In Stage 4, the A and B players are told their payoffs. In Stage 5, the players can chat with each other.

Stage 1 – **B-Players' Disclosure Decisions**
In Stage 1, B-Players will have the choice of deciding whether they want to let the A players receive the same information that the B players will see. B players will see the following:

Screen 1

Each B-player may click either 'Yes' or 'No' for each of the 2 A-players he is paired with.

Stage 2 – **A-Players' Decisions**
Each A-player sees the decisions made in Stage 1 by the B-players he is paired with.

Each A-player also receives an initial endowment of 10 Lira for each paired B-player. That is, each A-player will receive 10 Lira for B1 and 10 Lira for B2, for a total of 20 Lira.

In the second stage, each A-player will be prompted by the computer to decide how much of the initial endowment to keep and how much to send to a paired B-player. For example, A1 will decide what amount (from 0 to 10 Lira) to send to B1. Similarly, A1 will decide what amount to send to B2. The amount sent will always be in whole Lira. The A-Player will keep any money s/he has not sent to each B-Player.

The A-Player's Stage 2 decision will be entered on the Screen 2, pictured below.

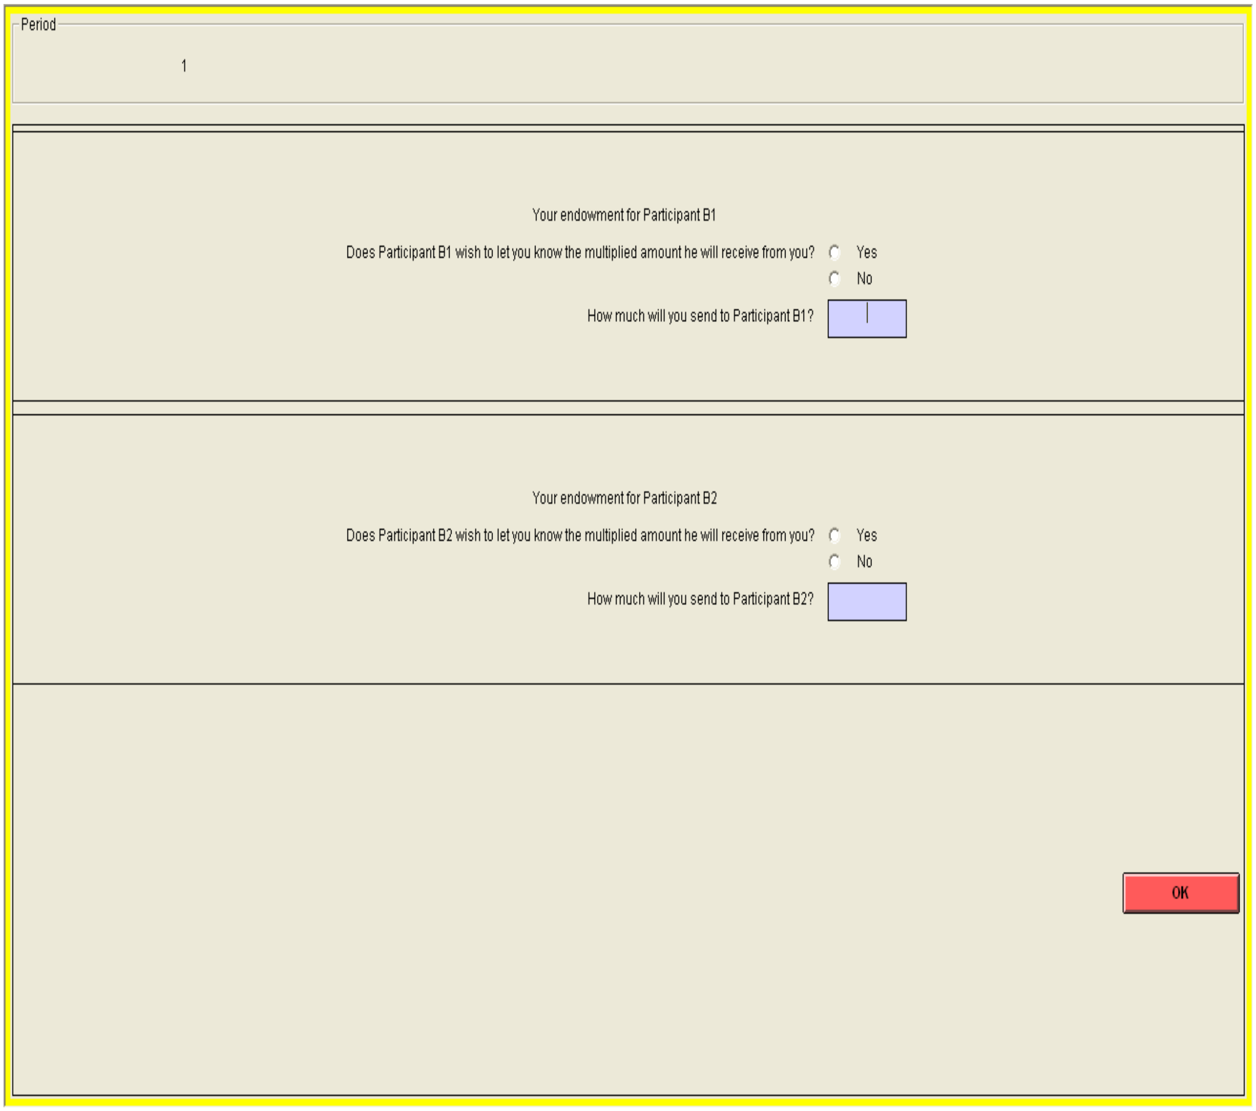

Screen 2

<u>*Stage 3*</u> – **B-players' Decisions**
For every pair of players, the amount sent by the A-Player is multiplied by 1 or 2 or 3 or 4 or 5 (referred to as 'the multiplier') before the B-Player receives it. Each B player will see the following screen (namely, Screen 3). Please note that different multipliers may be applied to different amounts sent. For each pair, every multiplier is equally likely in each round.

Screen 3

The B-player decides how much of the total amount to return to each A-player. Thus, each B-Player will send 2 amounts (from 0 to the total amount received) to the 2 different A-Players. The amount returned will always be in whole Lira.

The B-Player's Stage 3 decisions will be entered on Screen 3, pictured above.

Stage 4 - **Disclosure and Payoffs**
In each round, each A-Player's payoff will be the sum of the 2 total amounts that s/he did not send to the B-Players plus the sum of the 2 amounts returned by the B-Players.

In each round, each B-Player's payoff will be the sum of the 2 total amounts that s/he received minus the sum of the 2 amounts s/he returned to each of the A-Players.

Following each round, each A-Player will receive the information presented on Screen 4, pictured below.

Screen 4

Following each round, each B-Player will receive the information presented on Screen 5, pictured below.

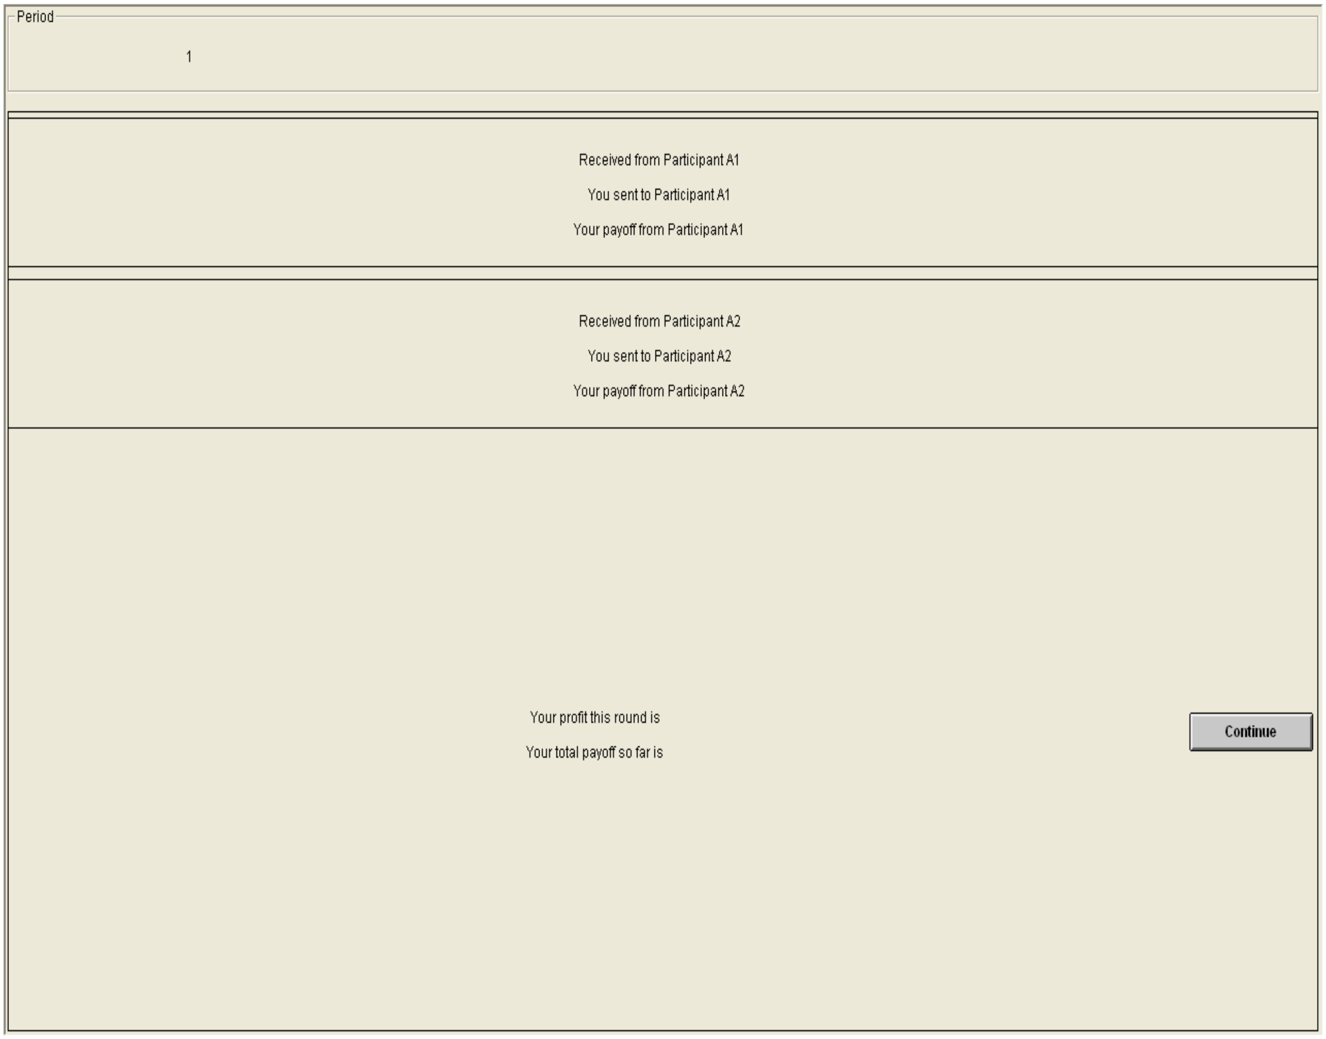

Screen 5

Stage 5 - **Chat Box**
After payoffs are shown, each A player will be able to chat privately with each of the 2 B players and each B player will be able to chat privately with each of the 2 A players. The chat will last for 3 minutes.

Do not use any profanities and do not disclose your name or any other personal information during the chat. You may type your message in the 'Type Message' box and then click send to send your message. The 'Chat Window' will show the messages send and received by you.

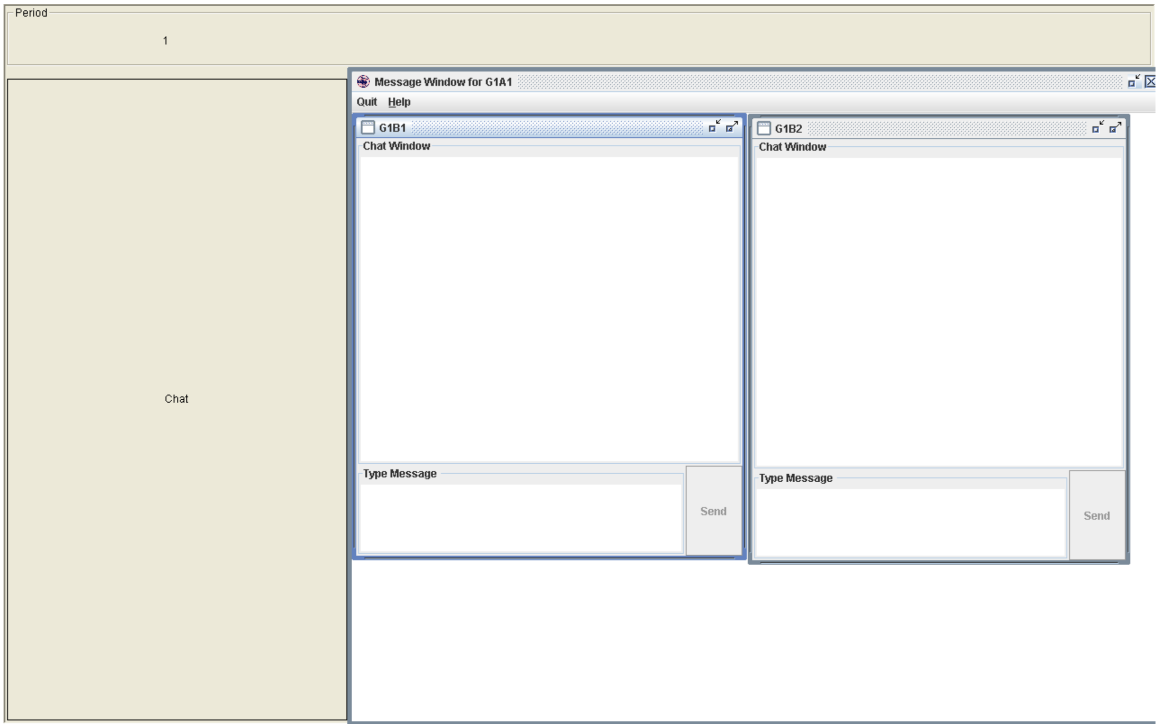

Screen 6

For instance, the chat screen of player A1 will appear as pictured in Screen 6 above. The 'Message Window' will be on your screen in Stages 1 through 4 also. However, you will be able to use this 'Message Window' for chatting only in Stage 5. (In Screen 6, 'G' indicates your group. For example, G1 means Group 1, G2 means Group 2, G3 means Group 3, and so on. You may ignore this.)

Completion of Rounds

After completing each round, the computer will proceed to the next round, which will be conducted identically to the previous round.

Once all rounds have been completed, you will be paid your cumulative income.

Summary
The following timeline summarizes what will happen in each round of the experiment:

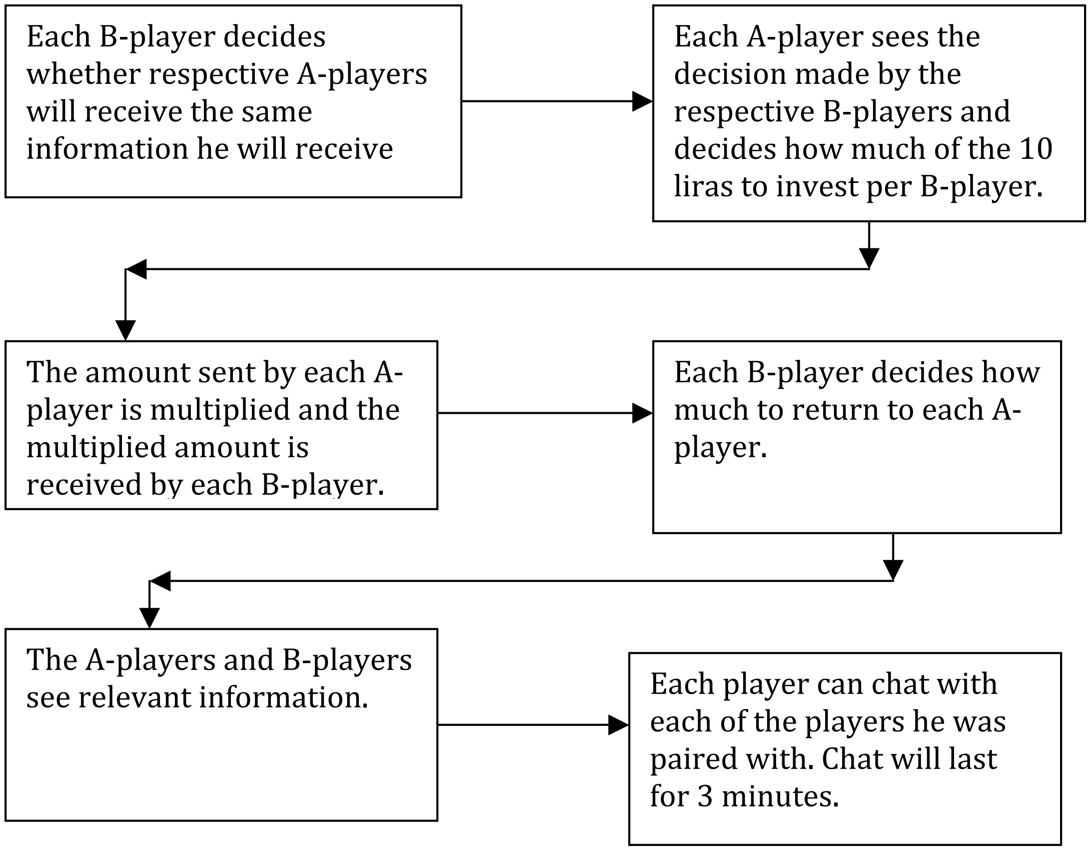

Please answer the following questions. In a few minutes, the administrator will check your answers and discuss with you any questions that have been answered incorrectly.

1. How many B-players will each A-player be paired with?

2. How many liras will each A-player receive for each B-player that s/he is paired with?

3. Will the amount sent by an A-player to a B-player be multiplied en route before it reaches the respective B-player (Yes / No)?

4. Is each multiplier equally likely for each pair in each round (Yes / No)?

5. Suppose player A1 sent to player B1 6 liras and then received from B1 10 liras. What will be player A1's profit from the pairing with B1?

6. Suppose player A2 sent 3 liras to player B2. Player B2 received 9 liras and sent back 2 liras to player A2. What will be Player B2's profit from pairing with A1?

7. The Message Window will be on your screen through Stages 1-5. What are the stages in which you can chat?