Do Actions Speak Louder Than Words? An Experimental Comparison of Observation and Cheap Talk

John Duffy*
Department of Economics
University of Pittsburgh
Pittsburgh, PA 15260, USA
jduffy+@pitt.edu

Nick Feltovich
Department of Economics
University of Houston
Houston, TX 77204, USA
nfelt@bayou.uh.edu

September 1999

Abstract

How do individuals achieve "good outcomes" in one–shot strategic situations? One much–explored possibility is that they engage in some kind of preplay communication—cheap talk—in which they endeavor to convince one another of the actions they intend to play. However, there may be no incentive for such communication to be truthful, or even informative. Another, less explored, possibility is that individuals take account of their knowledge of the past behavior of others when deciding which actions to play. While these two possibilities have been considered separately, there has been little research comparing the importance of these two devices as aids in achieving good outcomes. We design and run an experiment with human subjects that allows for such a comparison. The effects of cheap talk and observation of past actions are compared with each other, and with the standard (control) case where neither cheap talk nor observation is allowed. We consider three different 2×2 games and explain why cheap talk or observation is likely to be the more effective device for achieving good outcomes in each game. The experimental evidence suggests that both devices—cheap talk and observation—make cooperation and successful coordination more likely and increase payoffs relative to the control. The relative success of cheap talk versus observation in achieving such good outcomes depends on the game played, in accordance with our predictions. We also find that the signals players send are informative in the sense that they are cor related with their eventual actions, and that receivers of signals take this fact into account by conditioning their actions on the signal they receive. The results of this experiment can be used to extend game–theoretic models of how individuals make use of the different types of information available in strategic environments. As a first step in this direction, we construct a learning model in which individuals can condition their behavior on cheap talk or observed past actions, and we show that this model provides a good quantitative as well as qualitative fit to the experimental data.

---

*Funding for this project was provided by National Science Foundation grants SBR–9809622 and SBR–9813622. The raw data from the experiments we report in this paper are available from the authors upon request. We thank Nicole Bouchez, Catherine Eckel, Ann Gillette, Jack Ochs, Rick Wilson, and participants at several conferences and seminars for helpful comments and suggestions on earlier drafts. Any remaining errors will be useful in predicting our future actions.

1 Introduction

How do individuals achieve “good outcomes” in one–shot strategic situations where decisions must be made simultaneously?[1] One much–explored possibility is that they use “cheap talk”—costless nonbinding communication. Individuals communicate their intended actions to one another in the hope of swaying the expectations and actions of others. A second possibility, that has been largely overlooked, is that individuals turn to information regarding the past actions of others, when such information is available. Under this alternative scenario, individuals recognize that their current actions can be used by others to judge their likely behavior in future meetings. They may therefore choose to signal through their current choice of actions their intended behavior in future meetings.

The question we ask in this paper is *do actions speak louder than words?* We attempt to address this question by considering whether cheap talk or observation of past actions is the relatively better device for achieving good outcomes in simple one–shot games. We design and run an experiment in which human subjects play three 2×2 games—Prisoner’s Dilemma, Stag Hunt, and Chicken—under one of three information treatments: no information about opponents, cheap talk, or observation of opponents’ previous–round actions.

There is good reason to believe that in our current “postmodern era” actions *do* speak louder than words—that cheap talk may be the relatively weaker device for achieving good outcomes. The linguist John Haiman has recently observed (Haiman (1998)) that the old “talk is cheap” maxim is more relevant today than ever. Haiman notes that “unplain speaking,” which once carried a social stigma, is now commonplace. Individuals frequently say things that are very different from what they mean. Moreover, they recognize this same tendency in others. Haiman notes that the increasingly common use of such phrases as “yeah, right,” “whatever,” and “I couldn’t care less” reflect a pervasive and articulated cynicism with all human communications that was simply not present just a couple of decades ago.

Our experimental evidence suggests that both cheap talk and observation, while far from being perfect devices, do result in better outcomes—more cooperation, more coordination on pure–

---

[1] In Section 2, we will define “good outcomes” to be those in which cooperation and coordination on a pure–strategy Nash equilibrium are likely and payoffs are high.

strategy Nash equilibria and higher payoff efficiency relative to the standard treatment where players have no prior information about their opponent. Consistent with the work of Farrell and Rabin (1996), we find that the relative effectiveness of cheap talk or observation of past actions in achieving good outcomes varies systematically with the game. In particular, when the strategic structure of a game implies that cheap talk messages ought to be credible, cheap talk is found to be the relatively better device—words speak louder than actions. On the other hand, when the strategic structure of a game implies that cheap talk messages ought *not* to be credible, observation is relatively more effective than cheap talk—actions speak louder than words.

We also find that, at the individual level, subjects in the experiment take additional information—cheap talk messages or observed previous–round actions—into account when choosing their actions in each of the games. Indeed, we find that augmenting the reinforcement learning model of Roth and Erev (1995) to explicitly allow agents to condition their choice of action on all information, including cheap talk messages or observed previous–round actions when available, characterizes the experimental data better than either a standard reinforcement learning model, where individuals condition their choice of action on their past strategies and payoffs only, or a model that does not allow agents to change their behavior over time.

2 The Games

Figure 1 shows the games used in our experiment: Prisoner's Dilemma (which we will sometimes abbreviate PD), Stag Hunt (SH), and Chicken (CH). These games are alike in that they are well–known, symmetric, 2×2 games. However, they differ in their strategic structure (such as their best–reply correspondence), in the number and nature of Nash equilibria, and as we will see in Section 3, in the predicted effect of cheap talk on play relative to that of observation of prior actions. In each game, players choose between two strategies: Cooperate (C) and Defect (D).[2] Prisoner's Dilemma has a unique Nash equilibrium in which both players defect and earn 40 points. Stag Hunt has two

---

[2] Mehlmann (1997) notes that Prisoner's Dilemma, Stag Hunt and Chicken can be thought of as different parameterizations of a broad class of games, which he refers to as "Löwe–Lamm Spiele" (lion–lamb games). "Lion" corresponds to our Defect strategy, and "Lamb" corresponds to Cooperate.

pure–strategy Nash equilibria: one in which both players cooperate and earn 70 points and one in which both players defect and earn 55 points; the former is payoff dominant and the latter risk dominant. There is also a mixed–strategy Nash equilibrium in which both players cooperate with probability $\frac{3}{4}$ and earn 55 points. Chicken (CH) has three Nash equilibria: two in which one player cooperates and the other defects, the cooperating player earning 50 points and the defecting player earning 80, and one in which both players cooperate with probability $\frac{1}{2}$ and earn 60 points.

Figure 1: The Games

		Player 2	
		C	D
Player 1	C	70,70	10,80
	D	80,10	40,40

Prisoner's Dilemma

		Player 2	
		C	D
Player 1	C	70,70	10,55
	D	55,10	55,55

Stag Hunt

		Player 2	
		C	D
Player 1	C	70,70	50,80
	D	80,50	40,40

Chicken

We call the strategies "Cooperate" and "Defect" due to a characteristic shared by all three games: for a given opponent strategy (pure or mixed), playing Cooperate always weakly increases the *opponent's* payoff. This increase is strict in all cases except in Stag Hunt (when the opponent is choosing Defect with certainty). A high probability of cooperation, a high probability of coordination on a pure–strategy Nash equilibrium, and high payoffs are all desirable features of outcomes in the games we consider. By "good" outcomes, we shall mean outcomes in which as many as possible of these features are present. In Table 1 we show the extent to which these features are present in the equilibria of the three games we consider. Payoff efficiency is defined as the sum of row and column player payoffs, normalized so that the maximum possible joint payoff in a given game has an efficiency of one and the minimum possible joint payoff has an efficiency of zero. These features of the equilibria will be compared with the experimental results in Section 5.

Table 1: Characteristics of Equilibria of the Games

Game	Equilibrium	Prob(Cooperation)	Prob(Coordination)	Efficiency
PD	(0,0)	0	—	0
SH	(1,1)	1	1	1
	(0,0)	0	1	.6
	(.75,.75)	.75	.625	.6
CH	(1,0)	.5	1	.833
	(0,1)	.5	1	.833
	(.5,.5)	.5	.5	.667

Note: equilibria are of the form (Prob(row player cooperates),Prob(column player cooperates))

3 The Information Treatments: Theory and Hypotheses

In the experiment, players play each game ten times against changing opponents in one of three information treatments. In the first "control" treatment, players receive no information about their opponent before choosing actions. In the second "cheap talk" (or "communication") treatment, one member of each pair of players is randomly selected to send a costless, nonbinding message to the other player, prior to the choice of actions, indicating the action she intends to play when the two players meet in the subsequent round.[3] In the third "observation" treatment, one member of each pair of players is randomly selected to be informed of her opponent's action in the previous round (when the opponent was matched with a different player), prior to the choice of current–round actions.[4] We use the term "signal" to refer to a message in the cheap talk treatment or an observed action in the observation treatment; similarly, in both cheap talk and observation treatments, a "receiver" is a player who sees her opponent's signal, and a "sender" is a player whose signal is

---

[3] Note that in our design, cheap talk amounts to sending a signal of an intended action; no other communication is allowed. We chose to adopt this convention for cheap talk so that the information players receive would be comparable to the information received in the observation cells, where they observe an action played in the previous round.

[4] Remaining silent in the cheap talk sessions, or remaining unobserved in the observation sessions, was not an option available to subjects. Cooper et al. (1989) found that subjects rarely choose to remain silent in one–way cheap talk environments.

observed.[5] In our cheap talk and observation treatments, both players are equally likely to send or receive.

Theoretically, the additional information we allow in the observation and cheap talk treatments need not affect outcomes in the games we consider. For each game and each information treatment, any sequence of stage–game Nash equilibria is also a subgame perfect equilibrium of the corresponding finitely–repeated game. In the control treatment (where no extra information is available) and in Prisoner's Dilemma (which has a unique Nash equilibrium), these are the *only* subgame perfect equilibria of the finitely–repeated game. Since Chicken and Stag Hunt have multiple equilibria, allowing for either observation or cheap talk enlarges the set of equilibria; the additional information is known by both players, so they can condition the stage–game equilibrium they play on the information they send or receive. Additionally, observation allows subgame perfect equilibria of the finitely–repeated game that contain pairs of stage–game actions that are not equilibria of the stage game.[6]

Although the set of equilibria may be enlarged by the presence of observation or of cheap talk, standard game theory makes no prediction of which equilibrium should be expected, or even whether players take the additional information into account. Recently, however, theorists have suggested conditions under which cheap talk might lead to successful coordination in games with multiple equilibria, such as Battle of the Sexes (BoS).[7] Farrell (1987) studies repeated two–way cheap talk in the context of an entry game that is qualitatively simila r to BoS; Farrell argues that the presence of cheap talk allows for equilibria in which cheap talk messages are taken seriously if it

---

[5] Our cheap talk and observation treatments involve *one–way* signaling, in which (in each round) exactly one of the two players receives a signal sent by the other, as opposed to *two–way* signaling, in which both players receive signals from each other.

[6] For example, (C,C) in any stage but the last of a finitely–repeated Chicken game can be supported by the (correct) belief that in the next stage, signal receivers will play the strategy: "if I observe C, I will choose C; if I observe D, I will choose D," and suitable beliefs about later stages, in which case players will optimally play C in the current stage with an eye toward playing D in the next stage. This leads to equilibrium as long as players are sufficiently patient, as in our design, which would induce a discount factor of unity among hypothetical expected–utility maximizing agents with time–separable preferences.

[7] One example of a Battle of the Sexes game is our Chicken game, with the (70,70) outcome replaced by (40,40). Like Chicken, these games have multiple Nash equilibria, but none are payoff dominant. Unlike Chicken, both players prefer coordination on either pure–strategy Nash equilibrium over play of the mixed–strategy Nash equilibrium.

is optimal for senders of such messages to keep their promises and if senders believe that receivers believe those messages. Arvan, Cabral, and Santos (1999) go further, showing that in games like BoS, if a few additional conditions are satisfied, payoffs *will* be improved by two–way cheap talk. Rabin (1994) shows that allowing a large number of rounds of pre–play two–way communication in similar games enables players to ensure at least their minimum payoff in any Pareto efficient Nash equilibrium. Crawford and Sobel (1982) consider a class of signaling games with continuous type and message spaces and show that the more closely players' interests are aligned, the more likely one–way cheap talk is to be informative (in a sense which they define).

Aumann (1990) and Farrell and Rabin (1996) propose two conditions for cheap talk to facilitate coordination in situations where messages have literal meanings (that is, some convention exists for translating each message into a unique intended action) as in our design. One condition, *self–commitment*, is satisfied when the sender's message, if believed, binds the sender to playing the action she has signaled; that is, the sender's best response to the receiver's best response to the sender's message is the signaled action. For example, in Chicken, C is a self–committing message because the sender's best response to the receiver's best response (D) is the action C. By contrast, in Prisoner's Dilemma, C is not self–committing because the sender's best response to the receiver's best response (D) is not C, but D. The other condition, *self–signaling*, is satisfied when the sender prefers the receiver to play the best response to a given message if and only if she (the sender) truly intends to play the signaled action.[8] For example, in Stag Hunt, C is a self–signaling message because the receiver's best response C gives the *sender* a higher payoff than she would have received had the receiver chosen D, and had the sender intended to play D instead, she would not have preferred the receiver to choose C. By contrast, in Chicken, C is not a self–signaling message because the receiver's best response, D, is *not* preferred by the sender; a C response by the receiver would have given the sender a higher payoff.

According to Farrell and Rabin (1996), "a message that is both self–signaling and self–committing seems highly credible." Indeed, one expects that when messages in a given game are self–signaling

---

[8] Aumann (1990) used the term *self–enforcing* to refer to a message that was both self–committing and self–signaling.

and self–committing, they should be both truthful and believed. In our Stag Hunt game, both C and D messages are self–signaling and self–committing; in Chicken, both C and D messages are self–committing, but not self–signaling; and in Prisoner's Dilemma, C messages are neither self–signaling nor self–committing, while D messages are self–committing, but not self–signaling.[9] We can therefore hypothesize that messages (C or D) in the Stag Hunt are most often truthful and believed, messages in Chicken and D messages in Prisoner's Dilemma are less often truthful and believed, and C messages in Prisoner's Dilemma are least often truthful and believed.

There have been several experimental studies of cheap talk as a coordination device in games with multiple equilibria. Cooper et al. (1989) showed that when *one–way* communication is used in a two–player BoS game, coordination becomes much more likely—more so than when two–way communication is used. Cooper et al. (1992) find that two–way communication can be more useful than one–way communication in certain types of coordination games, such as in a game similar to the Stag Hunt game that we consider. However, they also find that one–way communication always improves coordination relative to the case of no communication in all varieties of coordination games that they examine. Swensson (1967) comes to a similar conclusion in his analysis of one–way versus no communication in the Prisoner's Dilemma game. Wilson and Rhodes (1997) provide further evidence that one–way communication is useful in achieving high–payoff equilibria in pure coordination games in which players' payoffs are common knowledge. In his survey of cheap–talk experiments, Crawford (1998) argues that one–way communication is preferable in coordination games such as BoS that require players to engage in "symmetry–breaking," that is, to play different actions, even though their action set and payoff functions are identical. (The two pure strategy Nash equilibria in our Chicken game require such symmetry–breaking.) Crawford further notes that in some games that do not require symmetry–breaking, such as the Stag Hunt game we consider, cheap talk may play an important "reassurance role," allowing the sender to signal that she understands the structure of the game and the existence of the payoff dominant equilibrium. Cooper (1999) provides a survey of coordination game experiments.[10]

---

[9] It can be shown that in a symmetric 2 × 2 game, if a message is self–signaling, it is self–committing.

[10] In addition to cheap talk, other devices that have been examined as aids in solving experimental coordination games are (1) an outside option in which one player first chooses between a sure payoff and playing the coordination

Surprisingly, cheap talk has also been found to induce greater cooperation in games, like Prisoner's Dilemma, in which messages are neither self–committing nor self–signaling: specifically, collusion in oligopoly games (for a survey, see Holt (1995), pp. 409–411) and public–good provision in public good games (for a survey, see Ledyard (1995), pp. 156–158). See also Dawes et al. (1977), Orbell et al. (1988), Ostrom et al. (1994), among others.

By contrast with cheap talk, observed previous–round actions are credible by their very nature. However, observed actions differ from cheap talk in the extent to which they can be considered signals of the sender's likely action. In the case of cheap talk, there is no question that a message is a signal; that is its only function. While a previous–round action may be thought of (by the receiver) as a signal, that was not necessarily its intent. Even if seen (which happens with probability only one–half in our design), previous–round actions play a dual role of signal and current–round action (in the previous round), the latter of which directly affects the sender's payoff. Thus, the receiver of such a signal must bear in mind that though credible, it is not a perfect forecast of the sender's current–round intended action. Since this is always the case in the three games of our experiment, we hypothesize that the extent to which observed previous–round actions correlate with current–round actions will not vary with the game. Therefore, the relative efficacy of observation versus cheap talk in facilitating good outcomes (in the sense that we've discussed previously) should depend on how credible cheap talk messages are in our games. When cheap talk is relatively more credible, it should be more effective than observation; when cheap talk is relatively less credible, it should be less effective than observation. Of course, information can always be ignored, so allowing more information should never be worse than, and may well be better than, allowing no information (as in our control).

While many researchers have looked at the effects of cheap talk, there has been little research

---

game (Cooper et al. (1993)), and (2) a preplay auction, in which the rights to play a coordination game are auctioned off to the highest bidders (Van Huyck et al. (1993), Crawford and Broseta (1998)). These devices, which are equivalent to giving some players an opportunity to choose a costly signal, typically do not reduce the number of Nash equilibria of these games, but "forward induction" refinements such as the Cho–Kreps (1987) intuitive criterion may eliminate all equilibria except those in which the sender sends a costly signal (bypassing the outside option in (1) or paying a high price for the right to play in (2)), and the Pareto efficient equilibrium of the stage game is then played.

into other potential coordination devices, such as observed past actions. The experimental literature contains some evidence that players do take this type of information into account, when available. Kahneman, Knetsch, and Thaler (1986), Eckel and Grossman (1996), and Fehr, Gächter, and Kirchsteiger (1997), for instance, show that players play more cooperatively toward opponents believed to have behaved cooperatively in the past, and less cooperatively toward opponents believed to have behaved uncooperatively in the past. This "reciprocity" can occur even when it is costly, and in some cases even when the opponent's past behavior was directed at someone else. Duffy and Feltovich (1999a) provide evidence that in a simple bargaining game, observation of other players' actions (along with their payoffs) can change the frequency of rejections of unfavorable offers, possibly by altering players' perceptions of what a fair outcome is.

We are aware of only one paper, that of Wilson and Sell (1997), in which *both* cheap talk and observation of past actions are considered. The Wilson–Sell study involved groups of subjects repeatedly deciding how much to contribute to a public good. They found that the combination of cheap talk and observation of past contributions resulted in a level of cooperation that is approximately the same as when both cheap talk and observation were absent, but that adding either cheap talk or observation by itself actually *decreased* the amount of cooperation.[11]

In addition to comparing the effects of observation to those of cheap talk (and the control), our experiment differs from previous studies in that we examine three different games, which yield quite different predictions regarding the relative efficacy of cheap talk versus observation as devices for achieving good outcomes. In Prisoner's Dilemma, cheap talk should be incredible (C messages are neither self–committing nor self–signaling, and even D messages are self–committing, but not self–

---

[11] There are several differences between the Wilson–Sell experimental setup and our own. One important difference is that their subjects play a repeated game, with the same opponents in every round. Consequently, their notion of "observation" necessarily includes not only opponents' previous–round actions, but also opponents' current–round actions (which will of course become the previous–round actions next round), and the subject's own payoffs in each round (which can be used to infer opponents' actions), so that their cells without observation (their "no–information" and "announcement" treatments, which correspond to our control and cheap talk treatments) present subjects with no opportunity to learn anything from one round to the next. By contrast, our experimental design involves a round–robin sequence of one–shot games, in which players always observe their payoff from every round, regardless of the information treatment. Thus, even in our control treatment, there is the possibility that subjects will learn on the basis of their own past history of play.

signaling); thus observation is hypothesized to be the relatively better device. In Stag Hunt, cheap talk should be credible (messages are both self–committing and self–signaling); thus cheap talk is hypothesized to be the relatively better device. Finally, in Chicken, cheap talk is less credible than in Stag Hunt, but somewhat more than in Prisoner's Dilemma (in Chicken, both types of messages are self–committing, but not self–signaling); thus cheap talk's efficacy relative to observation is hypothesized to lie somewhere in between what is found in Stag Hunt and Prisoner's Dilemma.

4 Experimental Procedures

We use a 3×3 experimental design in which we vary the game—PD, SH, or CH—and the information condition—cheap talk, observation, or control (neither cheap talk nor observation).[12] We have conducted 3 sessions of each information treatment (9 sessions total) where 20 subjects with no prior experience played 10 rounds of each of the 3 games (30 rounds total) under a single information condition (cheap talk, observation or control).[13] For some cells we have results from additional sessions, also involving 20 inexperienced subjects (in some of our experimental sessions, time constraints limited us to playing just two of the three games). We announced, prior to the tenth round of each game, that the tenth round would be the last round of that game. Subjects were primarily University of Pittsburgh undergraduate students. In each game, one–half of the subjects were assigned the role of "row player" and the other half were assigned the role of "column players." Subjects were randomly assigned one of these roles and remained in the same role throughout a game. We used a round–robin matching format, so that each row player faced each column player exactly once in a ten–round game. Hence, each round of each game was a "one–shot" encounter

---

[12] More precisely, we use a $3 \times 3 \times 2$ design, because we also considered two different orders in which subjects played the three games: PD–SH–CH and CH–SH–PD. We found no systematic differences in play due to the ordering of the games, so in our results, we pooled the data from both orderings. For a similar reason, we pooled the row– and column–player data.

[13] In another paper (Duffy and Feltovich (1999c)), we combine the two information conditions, cheap talk and observation, into a single treatment, so that receivers of cheap talk messages also observe the previous–round action played by the message sender.

for each pair of players.[14]

The experimental sessions were conducted in the Department of Economics computer laboratory at the University of Pittsburgh, using networked personal computers. Each subject was seated at a computer and given written instructions. These instructions were also read aloud in an effort to make the rules of the experiment common knowledge. A sample copy of the instructions can be found in the Appendix. The computer screen displayed the payoff matrix for each game, the results of the player's previous rounds of play of that game, and additional information depending on the treatment. The payoff matrix was also drawn on a blackboard for all subjects to see. Subjects input their actions by choosing which row or column of the current game payoff matrix they wanted to play. Row players' actions were labeled R1 and R2 on their computer screens, and Column players' actions were labeled C1 and C2, in both cases corresponding to C and D, respectively. In describing the actions available to subjects we avoided any reference to the labels "cooperate" or "defect," and we referred to each player's opponent as his or her "partner."

In the first round of every treatment, each row player was randomly paired with a column player. No cheap talk or observation of past actions occurred in this first round of play.[15] Beginning with the second round of each game of the cheap talk and observation treatments, and continuing with every round thereafter, cheap talk or observation took place before subjects chose their current–round actions. In both treatments, one member of each pair of players was chosen with probability one–half to be the sender of a cheap talk message or to have his previous–round action revealed to the other player.[16] In the cheap talk sessions, subjects who were selected to be message senders

---

[14] Kamecke (1997) shows that this matching technique ensures that the ten–round game maintains the one–shot character of the stage game, and does so efficiently in the sense that there is no way to increase the number of rounds, while keeping the same number of players and continuing to maintain the one–shot nature of th e game.

[15] In the first round of each game, observation of past actions is not possible. For comparison purposes, we chose to suspend cheap talk in the first round as well.

[16] In some one–way cheap–talk experiments, such as those of Cooper et al. (1989), subjects alternated between the roles of sender and receiver in a known, deterministic manner. Such a design may work well for studying cheap talk, but it is less well suited for the study of the effect of observation. It is not unreasonable to imagine that subjects behave differently when they know they're being observed than when they know they're not being observed. So, if a subject's action was observed in a particular round, and she knew she was not being observed in the next round, the earlier action may not be useful for forecasting the later action. Our design avoids this problem.

were limited to sending either the message C1 or C2 (C or D) if they were a column player or R1 or R2 (C or D) if they were a row player. The instructions asked that senders send a message "indicating the action they intend to choose in the next round." The instructions further explained that these messages were "not binding" on the action subsequently chosen by the message sender; the sender could choose either of the two available actions regardless of the message sent. The message choices were entered using the computer keyboard (there was no verbal communication) and then revealed to receivers on their computer screens prior to the play of the round. In the observation treatment, the information on previous–round actions was also revealed to receivers via their computer screen. In both the cheap talk and observation treatments, senders were asked to record, on record sheets, the signals their opponents would receive from them in the next round of play. Similarly, receivers were asked to record the cheap talk message they received or their opponent's previous–round action. Our aim was to call subjects' attention to the information that they had either provided or received. After cheap talk or observation had taken place and the various signals had been recorded, or following the pairing of subjects in all rounds of control games and in the first round of the cheap talk and observation games, each player chose an action—either R1 or R2 for row players, and either C1 or C2 for column players. After all choices were made, current round payoffs in points were revealed to all subjects.

Subjects were informed that each point in the payoff matrix represented a 1% chance of winning $1.00. At the end of each round, an integer was randomly drawn from the interval [1,100]. Subjects whose stage–game payoff was greater than or equal to the chosen number earned $1.00 for the round; those whose payoff was less than that number earned nothing for the round.[17] At the end of the session, subjects received in cash their total earnings from every round as well as a $5.00 participation payment. Sessions lasted between 60 and 75 minutes. Subjects earned an average of $25.00 for participating in each session.

---

[17] This binary lottery procedure is intended to induce risk neutral behavior among hypothetical expected–utility maximizing individuals. See, e.g., Roth and Malouf (1979) for a discussion.

5 Results

Thirteen experimental sessions were conducted. Our experimental findings include a number of interesting results, which we now describe.

Result 1 *Subjects' actions are influenced to some extent, but not completely, by strategic considerations.*

The third column of Table 2 shows the aggregated (all subjects, all rounds) relative frequencies of cooperation in each cell of the experiment. Notice that differences in the level of cooperation are much more stark between *games* than between *information conditions*; cooperation is generally highest in the SH cells and lowest in the PD cells. The only exception to this rule is that cooperation is less frequent in the SH–control cell than in the CH–observation cell. Recall that the unique Nash equilibrium in Prisoner's Dilemma involves no cooperation at all, while all three Nash equilibria in Chicken involve cooperation 50% of the time (differing only in how much each player cooperated), and two of the three Nash equilibria in Stag Hunt involve cooperation between 75% and 100% of the time. The ordering of the observed frequencies of cooperation across the three games is found to be very close to the ordering of the equilibrium levels of cooperation across the three games. (Admittedly, the third Nash equilibrium of Stag Hunt has cooperation never occurring.)

The Nash equilibrium point predictions fare well in the SH cells, provided we choose the equilibrium that is closest to the observed data; the mixed–strategy equilibrium has cooperation occurring 75% of the time, and the actual relative frequency of cooperation in the three SH cells ranges from 0.607 to 0.835. In the CH and PD cells, there are substantial deviations from equilibrium, all in the direction of higher cooperation. The actual relative frequency of cooperation in the three CH cells ranges from 0.537 to 0.634, while the equilibrium prediction is exactly 0.5, and the actual relative frequency of cooperation in the three PD cells ranges from 0.222 to 0.404, while the equilibrium prediction is zero. It thus appears that play is influenced, at least to some extent, by nonpecuniary considerations.[18]

---

[18] It is possible that the degree of cooperation observed reflects some of the features of our experimental design. Framing may have played a part. As noted earlier, we used the friendly word "partner" instead of the hostile "opponent" or the neutral (though cumbersome) "counterpart" in the instructions given to subjects in the experiment.

Table 2: Experimental Relative Frequencies and Efficiency (All rounds)

Game	Treatment	Cooperation		Coordination		Efficiency
PD	Control	.222	(222/1000)	—		.113
	Cheap Talk	.400	(320/800)	—		.260
	Observation	.404	(323/800)	—		.266
SH	Control	.607	(364/600)	.513	(308/600)	.453
	Cheap Talk	.835	(501/600)	.840	(504/600)	.803
	Observation	.757	(454/600)	.667	(400/600)	.636
CH	Control	.537	(430/800)	.475	(380/800)	.696
	Cheap Talk	.564	(451/800)	.532	(426/800)	.741
	Observation	.634	(507/800)	.438	(350/800)	.780

The next few results concern the effect of additional information on aggregate play.

Result 2 *The addition of either observation or communication increases the frequency of cooperation relative to the control sessions, though the relative efficacy of observation and communication depends on the game.*

Supporting evidence for this second result is provided in Table 2. Consider first the PD cells. We see that adding either observation or cheap talk greatly increases the frequency of cooperation. There is little difference in the degree of cooperation between the observation and cheap talk cells; in both, cooperation occurs roughly twice as often as in the control. Nonparametric statistical tests confirm this; according to robust rank–order tests on the session–level data, cooperation is significantly more likely both in cheap talk sessions ($p < 0.01$) and in observation sessions ($p < 0.05$)

(On the other hand, the level of cooperation may have been even higher had we named the strategies "Cooperate" and "Defect" rather than "1" and "2.") Also, the amount that is at stake in any single decision is small—never more than thirty cents, and sometimes as little as ten cents (in expected value). Thus, even in a game such as the Prisoners' Dilemma where cooperation is strictly dominated, it is never terribly costly. We will see shortly, though, that the degree of cooperation observed is not merely the effect of framing and nonsalient payoffs. In fact, it varies systematically not only with the game played, but also with the information treatment and the particular piece of information sent or received.

than in control sessions, but there is no significant difference in the level of cooperation between the observation and cheap talk sessions.[19]

In the SH cells, there are no significant differences between treatments. This lack of significance may seem surprising, since Table 2 suggests that adding either observation or cheap talk *increases* the overall frequency of cooperation substantially. The explanation for this discrepancy is that there is a lot of variance in the control sessions; in two of the three control sessions, the frequency of cooperation is close to 50%, while in the other one, it is much higher (81%) and comparable to that in the non–control cells. Since we have only three control sessions, the high variance across these sessions accounts for the insignificant results. In the CH cells, there is little difference in the level of cooperation between the control and cheap talk treatments. However, observation increases cooperation significantly, relative to both the cheap talk treatment ($p < 0.10$) and the control ($p < 0.05$). The differences in the cells can also be seen when the data are disaggregated by round. The round–by–round observed relative frequencies of cooperation are shown in Figure 2. Also shown are the average frequency of cooperation, at the far right of each box, and the frequency of cooperation in each Nash equilibrium, as a horizontal line marked by an "x" (see also Table 1).

Consistent with many repeated Prisoner's Dilemma experiments, cooperation decreases over time (toward the equilibrium level of zero) in all three PD cells. The level of cooperation in the PD–control cell is lower in every round than in the other two cells, and there is no visible difference between the level of cooperation in the cheap talk and observation cells. In the SH–control cell, cooperation decreases over time, while in the other two cells, cooperation increases (particularly in the SH–cheap talk cell). While it appears that the level of cooperation in the SH–observation cell approaches the mixed–strategy equilibrium, this is not the case in the control and cheap talk cells. In all rounds but the first, the level of cooperation is higher in the SH–cheap talk cell than in the observation cell, and in all rounds but the third, the level of cooperation is higher in the observation cell than in the control cell. In the CH–control and CH–observation cells, the level of

---

[19] We use the robust rank–order test instead of the more commonly used Wilcoxon–Mann–Whitney test because the latter assumes that the two samples being compared come from distributions with identical second– and higher–order moments, which we have no reason to believe a priori. See Siegel and Castellan (1988) for a discussion of this issue, as well as more thorough descriptions of the nonparametric statistical tests used in this paper.

Figure 2: Average Relative Frequency of Cooperation (All Sessions)

Rounds 1–10, 10–Round Averages, and Equilibrium Predictions

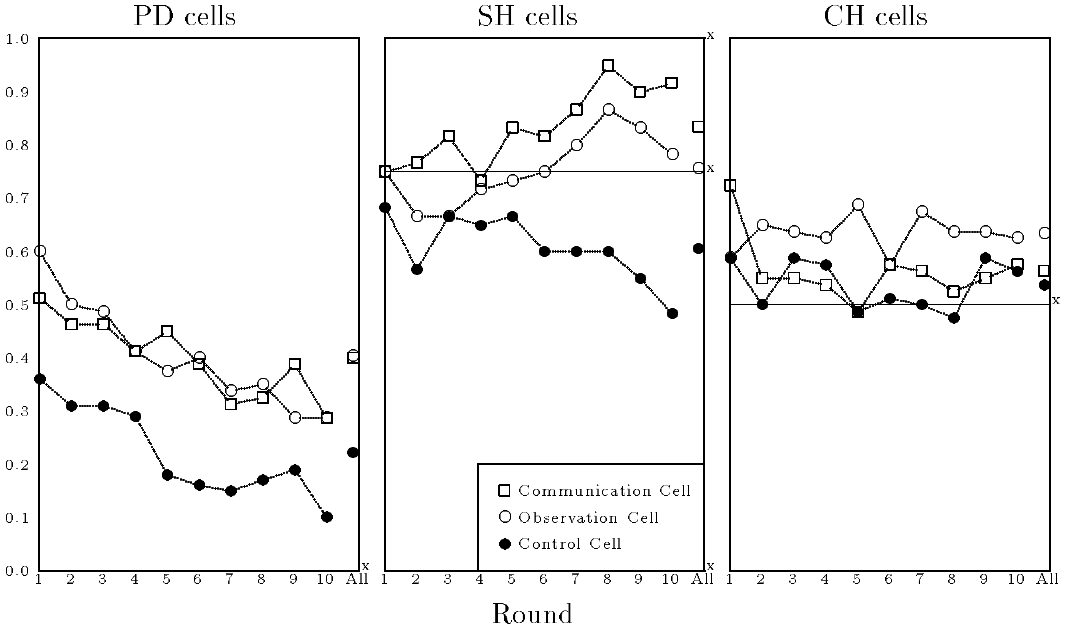

cooperation stays roughly constant over time, while in the CH–cheap talk cell, it drops between the first and second rounds and stays roughly constant thereafter. The level of cooperation starts above the equilibrium level and stays above it in almost every round (every round of the CH–observation cell). There is no visible difference between the level of cooperation in the control and cheap talk cells, but there is more cooperation in the observation cell than the control in all rounds but the first, and more in the observation cell than the cheap talk cell in all rounds but the first and sixth.

There are two other ways in which the information treatment affects aggregate behavior, both of which are suggested by Table 2.

Result 3 *In the games with multiple equilibria, cheap talk aids successful coordination on a pure-strategy Nash equilibrium.*

According to Table 2 (fourth column), coordination in the SH game is most frequent in the cheap talk cell; it is significantly more likely than in the control at the 5% level (robust rank–order test), and it is significantly more likely than in the observation treatment at the 10% level. In the

CH game, coordination is again more frequent in the cheap talk cell than in either of the other two cells; again the differences are significant (significant at the 5% level vs. the control, and at the 2.5% level vs. the observation cell).

An even more interesting feature of the cheap talk sessions is the high level of successful coordination relative to that which would be expected *given the observed frequency of cooperation.* In Stag Hunt, if q denotes the observed frequency of cooperation, and if there is no correlation between the actions of Row and Column players, then the probability of coordination ought to be $q^2 + (1-q)^2$. In the SH–cheap talk cell, the observed frequency of cooperation would imply a "predicted" frequency of coordination of 0.724. However, the *actual* frequency of coordination, 0.840, was considerably higher. In contrast, the actual frequency of coordination in the SH–observation cell was only slightly higher than the predicted frequency (0.667 vs. 0.632), and in the SH–control cell, the actual frequency was slightly lower than the predicted frequency (0.513 vs. 0.523). Similarly, in Chicken, if q denotes the observed probability of cooperation and there is no correlation between the actions of Row and Column players, then the probability of coordination ought to be $2q(1-q)$. In the CH–cheap talk cell, the actual frequency of cooperation would imply a predicted frequency of coordination of 0.492. However, the *actual* frequency of coordination was higher, at 0.532. In the other two cells, the predicted and actual frequencies of coordination were roughly the same—0.497 and 0.475 in the control cell, and 0.464 and 0.438 in the observation cell).

To further illustrate the relationship between predicted and actual frequencies of coordination, we disaggregate the data into individual sessions. Figure 3 shows, for the Stag Hunt and Chicken games and for all three information treatments, the predicted and actual frequencies of coordination *by session.* If there is zero correlation between the actions of Row and Column players in a session, the predicted and actual frequencies of coordination in that session will be equal, and the circle representing that session will lie on the 45° line in the box corresponding to that session's information treatment. Therefore, if there is no relationship between the actions of Row and Column players under a particular information treatment, the circles representing to the sessions under that treatment will be as likely to lie below the 45° line as above it. The circles corresponding to control sessions fit this pattern; four circles are below the 45° line and three are above it. The circles corresponding to observation sessions also fit this pattern; four circles are below the 45° line and three are above it. In contrast, all seven circles corresponding to cheap talk sessions are above the 45° line.

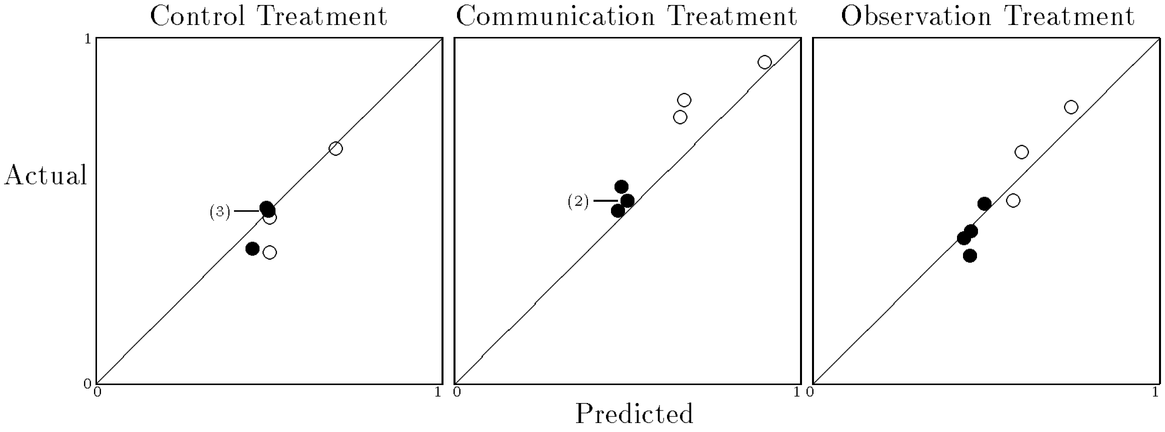

Nash Equilibrium (All Treatments)—Open Circle=Stag Hunt; Closed Circle=Chicken

Figure 3: Predicted and Actual Probability of Coordination on a Pure Strategy

We can use the nonparametric sign test to confirm whether or not there are significant differences between predicted and actual frequencies of coordination in each treatment. If there are no systematic differences between the predicted and actual frequencies in a treatment, the set of predicted frequencies and the set of actual frequencies will have the same distribution; if actual frequencies of coordination are higher than predicted frequencies, the distribution of actual frequencies will be higher than that of predicted frequencies. The test statistics confirm what is clearly visible in Figure 3. In the cheap talk (communication) treatment, actual coordination frequencies are significantly higher than predicted coordination frequencies ($p < .008$); in the other two treatments, the null hypothesis of no difference between predicted and actual frequencies of coordination cannot be rejected. These findings do not change if we use the stronger Wilcoxon summed–ranks test, which looks not only at the signs of the differences between predicted and actual probabilities, but also at their magnitudes.

A final way in which the information treatment affects aggregate behavior is its effect on joint payoff efficiency. (See Table 2, fifth column.)

Result 4 *In all three games, the provision of additional information increases payoff efficiency, but the increase is not always significant.*

In Prisoner's Dilemma, both communication and observation increase average payoff efficiency (over all sessions) from around 11% to over 25%. According to robust rank–order tests on the session–level data, payoffs in the PD–communication cell are significantly higher than those in the control at the 1% level, and payoffs in the PD–observation cell are significantly higher than those in the control at the 5% level. In Stag Hunt, efficiency increases somewhat in the observation cell (63.6% vs. 56.2% in the control cell), and still more in the communication cell (80.3%), but only the difference between the SH–communication cell and the SH–control is significant ($p < 0.05$). In Chicken, efficiency in the communication cell is slightly higher than that in the control (74.1% vs. 69.6%), but it is highest of all in the observation cell (78.0%); even though these subjects often fail to coordinate, their failures are generally high–payoff (C,C) pairs, rather than low–payoff (D,D) pairs. The difference between payoffs in the CH–cheap talk cell and in either of the other two CH cells is not significant at the 10% level, but payoffs in the CH–observation cell are significantly higher than in the control ($p < 0.05$).

As before, the differences across the cells can also be seen when the data are disaggregated by round. The round–by–round observed payoff efficiencies are shown in Figure 4. Also shown are the ten–round average payoff efficiencies and the payoff efficiency in each Nash equilibrium. As was true for the level of cooperation, payoff efficiency in all three PD cells decreases over time toward the equilibrium level of zero, with payoffs in the cheap talk and observation cells higher in every round relative to the control. The differences across the SH cells in the level of cooperation become magnified when payoff efficiency is considered. In every round, payoffs are higher in the cheap talk and observation cells than in the control. In addition, while the frequency of cooperation in the cheap talk cell wasn't much higher than in the observation cell, payoff efficiency is much higher; in the first round, payoffs are equal in the two cells, but payoffs rise over time in the cheap talk cell while remaining roughly constant over time in the observation cell. The differences in payoff efficiency across the CH cells are small; payoffs rise slightly in the observation cell and fall slightly in the other two cells, so that payoffs in the first round are lowest in the observation cell, but in

Figure 4: Average Payoff Efficiency (All Sessions)

Rounds 1–10, 10–Round Averages, and Equilibrium Predictions

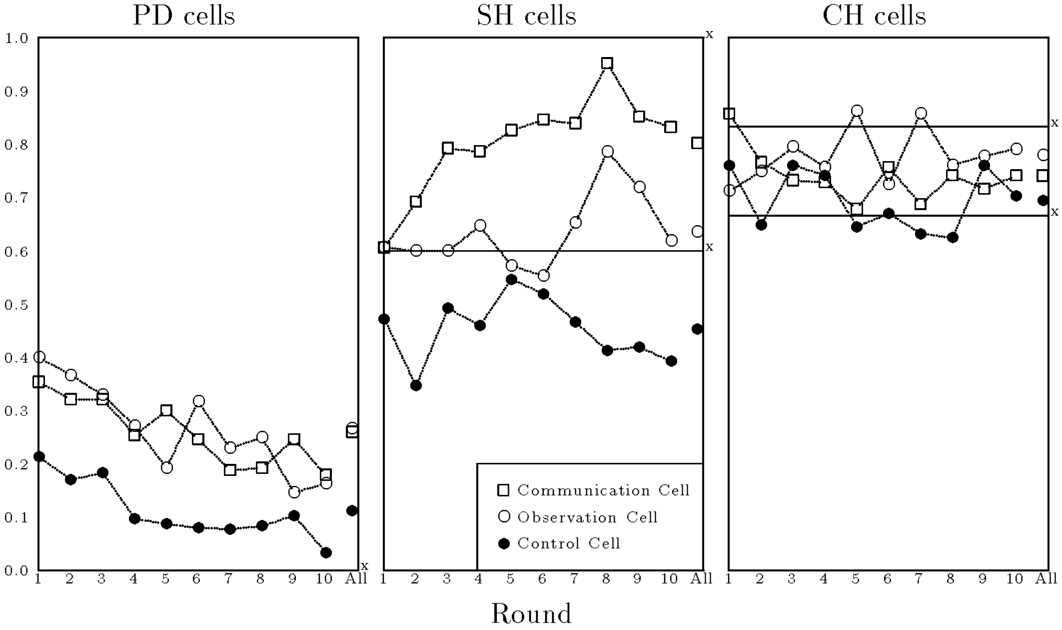

seven of the last eight rounds, payoffs are highest in the observation cell.

We now turn to the effect of signals on the behavior of individual receivers.

Result 5 *Subjects in the communication and observation cells who receive additional information condition their actions on this information.*

Supporting evidence is provided in Tables 3, 4, and 5, which show how the information given or received by subjects in the cheap talk and observation cells affected play. We report the relative frequency of cooperation by players in response to (a) receiving a C signal, (b) receiving a D signal, (c) sending a C signal, or (d) sending a D signal.

Notice that in every cell, in both communication and observation treatments, receivers of C signals cooperate more than receivers of D signals. Consider first the SH cells (see Table 4). In the SH–communication treatment, players almost always (over 97% of the time) choose C after receiving a C signal, and they usually (over 80% of the time) choose D after receiving a D signal. In the observation treatment, players choose C over 80% of the time after receiving a C signal,

Table 3: Conditional Relative Frequencies of Cooperation—Prisoner's Dilemma (All rounds)

Cheap Talk Cell			Observation Cell		
receive C	.504	(119/236)	observe C	.597	(89/149)
receive D	.161	(20/124)	observe D	.213	(45/211)
send C	.513	(121/236)	C observed	.631	(94/149)
send D	.153	(19/124)	D observed	.223	(47/211)

Table 4: Conditional Relative Frequencies of Cooperation—Stag Hunt (All rounds)

Cheap Talk Cell			Observation Cell		
receive C	.978	(226/231)	observe C	.827	(167/202)
receive D	.128	(5/39)	observe D	.559	(38/68)
send C	.918	(212/231)	C observed	.871	(176/202)
send D	.282	(11/39)	D observed	.412	(28/68)

and they choose D about 54% of the time after receiving a D signal. According to either the sign test or the Wilcoxon signed–ranks test, applied to the session–level results, players are significantly more likely to cooperate after receiving a C signal than a D signal (p=0.125, the smallest p–value possible when there are three sessions for each treatment) in both treatments.

More surprising (from a game–theoretic standpoint) is the extent to which subjects in the PD cells consider their additional information (see Table 3). Consistent with similar results reported

Table 5: Conditional Relative Frequencies of Cooperation—Chicken (All rounds)

Cheap Talk Cell			Observation Cell		
receive C	.623	(132/212)	observe C	.599	(136/227)
receive D	.304	(45/148)	observe D	.579	(77/133)
send C	.623	(132/212)	C observed	.780	(177/227)
send D	.568	(84/148)	D observed	.526	(70/133)

by other researchers, subjects in the observation cell chose to cooperate about 60% of the time when they observed that their opponent cooperated in the previous round. When they observed that their opponent defected in the previous round, subjects defected almost 80% of the time. (The overall frequency of defections in all PD cells was around 67%.) It certainly appears that subjects are willing to forgo some (expected) monetary payoff by cooperating with people who have previously cooperated. Subjects in the cheap talk cell also take their additional information into consideration. They are much more likely to cooperate with someone who sends a cooperate signal than with someone who sends a defect signal (50.4% vs. 16.1%). Again, players are significantly more likely to cooperate after receiving a C signal than a D signal (p=0.0625, the smallest p–value possible when there are four sessions of each treatment) in both treatments.

In the CH cells, subjects also appear to take their additional information into account (see Table 5). In the observation treatment, subjects cooperate slightly more after receiving a C signal than after receiving a D signal (59.9% vs. 57.9%), but this difference is not significant.[20] Subjects in the communication treatment are substantially more likely to cooperate after receiving a C signal than after receiving a D signal (62.3% vs. 30.4%), but again, this difference is not significant.

The fact that subjects in both CH treatments are at least as likely to cooperate after receiving a C signal as when they receive a D signal runs counter to the Nash equilibrium prediction: both pure–strategy equilibria of this game require the two players to choose different actions, so if players believe that messages and previous–round actions are positively correlated with current–round actions, they should do the opposite—choose D when receiving a C signal and choose C when receiving a D signal. This result suggests that subjects' choices of actions may be influenced

---

[20] When the CH–observation data are broken down by sessions, an interesting pattern emerges. In three of the four sessions, players whose previous–round D signals were observed usually defect in the current round (more than 75% of the time in each of the three sessions) and observers of D signals usually cooperate (at least two thirds of the time in each of the three sessions). In the fourth session, however, the opposite happens. Players whose previous–round D signals were observed *cooperate* over 70% of the time in that session, and observers *always* defect. The degree of cooperation by senders of D signals in this session is even higher than that by senders of C signals in that session (the only session in which this is true), and is enough to make the difference in the level of cooperation between senders of C and D signals in the CH–observation cell insignificant. These profiles of sender/receiver behavior are examples of the two possible types of the "symmetry–breaking" phenomenon that one–way observation allows, and which has been observed in BoS experiments.

by nonpecuniary aspects of outcomes (such as fairness or reciprocity).

Having examined the effect of signals on receiver behavior, we now consider the effect of signals on *sender* behavior.

Result 6 *Subjects tend to send truthful signals even when they have no incentive to do so.*

In particular, in every cell of the communication and observation treatments, senders of C signals cooperate more frequently than senders of D signals. As noted in Section 3, Stag Hunt is the game in which we expect cheap talk messages to be the most truthful. Indeed they usually are; subjects who send C messages actually choose C over 90% of the time, and those who send D messages actually choose D over 70% of the time. According to either the sign test or the Wilcoxon signed–ranks test (applied to the session–level results), players in the cheap talk cell are significantly more likely to cooperate after sending a C signal than a D signal (p=0.125, the smallest p–value possible when there are three sessions of each treatment).

In Chicken, cheap talk is less likely to be credible than in Stag Hunt, but players still have an interest in coordination, so there remains some incentive to send truthful messages. In fact, messages in the CH cheap talk cell do tend to be truthful, but not nearly as frequently as in the SH cheap talk cell; in CH, 62.3% of C messages are followed by C actions, but only 43.2% of D messages are followed by D actions, so the overall fraction of truthful messages is just over one–half (54.4%). The difference in cooperation between senders of C messages and senders of D messages in CH is not significant.

In Prisoner's Dilemma, messages are not self–signaling, and C messages are not even self–committing. Indeed, given the way receivers of messages react (Result 5), there is a strong incentive to always signal Cooperate and then choose the Defect action. Nevertheless, in the PD cheap talk cell, Cooperate messages are truthful about half the time, and Defect messages are truthful almost 85% of the time. Once again, according to either the sign test or the Wilcoxon signed–ranks test, players in the communication cell are significantly more likely to cooperate after sending a C signal than a D signal (p=0.0625, the smallest p–value possible when there are four sessions of each treatment).

In the SH–observation cell, subjects who sent C signals subsequently chose to cooperate around 87% of the time, while those who sent D signals subsequently chose to cooperate only about 40% of the time. This difference is significant at the 12.5% level (p=0.125, three sessions of each treatment). In the PD–observation cell, subjects who sent C signals subsequently chose to cooperate 63.1% of the time; those who sent D signals subsequently chose to cooperate only 22.3% of the time. This differ ence is significant (p=0.0625, four sessions of each treatment). In the CH–observation cell, subjects who sent C signals subsequently chose to cooperate 78.0% of the time while those who sent D signals subsequently chose to cooperate 52.6% of the time. Again, this difference is significant (p=0.0625, the smallest p–value possible when there are four sessions of each treatment). We note, however, that the correlation between observed previous–round actions and current–round actions does not differ from the correlation between *unobserved* previous–round actions and current–round actions in any of our cells. In other words, players' actions tend to be positively autocorrelated, and this autocorrelation does not depend on whether individuals were observed or not. Of course, this does not reduce the utility of observed actions in forecasting current–round actions.

Given Result 6, that subjects' signals tend to be truthful and hence informative, Result 5, that receivers condition their actions on the signals they receive, seems reasonable. Using the actual relative frequencies of cooperation, conditional on players' sent signals (the bottom two rows of Tables 3, 4, and 5), we can calculate the expected payoff to cooperating or defecting for the *receivers* of signals. If players' choices depend primarily on strategic considerations, we would expect to see that the probability of cooperation depends on the difference in expected payoff between C and D—the expected payoff to C minus the expected payoff to D. The higher this difference, the more likely cooperation ought to be. If players' choices depend solely on strategic considerations, we would expect them to choose C if and only if the difference in expected payoff from playing C rather than D is positive.

These expected–payoff differences are shown in Table 6. For each cell and for both signals received, the difference in expected payoff between cooperating and defecting is shown, given the

Table 6: Payoff Differences and Strategy Choices (All rounds)

Cell	Signal Received	Expected Payoff (C)	Expected Payoff (D)	Difference	Frequency of C
PD–communication	C	40.76	60.51	–19.75	0.504
PD–communication	D	19.19	46.12	–26.94	0.161
PD–observation	C	47.85	65.23	–17.38	0.597
PD–observation	D	23.36	48.91	–25.55	0.213
SH–communication	C	65.06	55.00	+10.06	0.978
SH–communication	D	26.92	55.00	–28.08	0.128
SH–observation	C	62.28	55.00	+7.28	0.827
SH–observation	D	34.71	55.00	–20.29	0.559
CH–communication	C	62.45	64.91	–2.45	0.623
CH–communication	D	61.35	62.70	–1.35	0.304
CH–observation	C	65.59	71.19	–5.60	0.599
CH–observation	D	60.53	61.05	–0.53	0.579

actual observed conditional relative frequencies of C and D by all signal senders over all rounds.[21] The actual relative frequency of C by receivers is also shown. It is obvious from the table that it is not solely strategic considerations that matter to signal receivers. In the PD and CH cells, the payoff–maximizing action is always D (since the payoff difference is negative in every situation), but cooperation often occurs, in some cases more than half the time. However, expected payoffs do matter. In the four cases of the PD game, there is a perfect (ordinal) relationship between the difference in expected payoff and the observed frequency of C. In the SH game, this relationship is also perfect. This perfect relationship in both games between difference in expected payoff and the observed frequency of cooperation corresponds to a Spearman rank–order correlation coefficient of 1.000, which is significantly different from zero at the 10% level. In the CH game, there is

---

[21] Of course, the exact values of the expected payoffs to C and D reflect a lot of information that subjects do not know, such as play in future rounds and in other sessions. Therefore the results we report here, though suggestive, must be interpreted with some caution. In Section 6 we develop a model of individual learning, in which agents choose actions based only on the information they actually receive.

not a perfect relationship between the difference in expected payoff and the observed frequency of cooperation. Rather, the relationship is nearly a perfect inverse relationship; in the two situations in which the difference in expected payoff is highest, the actual frequency of C is lowest. This negative relationship corresponds to a Spearman coefficient of -0.600, which is not significantly different from zero.[22]

5.1 Do Actions Speak Louder Than Words?

While both messages in the cheap talk sessions and previous–round actions in the observation sessions are helpful in predicting current–round actions in all three games, the question asked in the title of this paper remains: *do actions speak louder than words?* That is, which type of signal is a better indicator of the action the sender actually plays? One obvious measure of the value of a signal is its "truthfulness"—the likelihood that the signal is the same as the action that gets played. According to this criterion, in the SH cells, words speak louder than actions; cheap talk messages are the same as current–round actions 88.9% of the time, while observed previous–round actions are the same as current–round actions only 80.0% of the time. On the other hand, actions speak louder than words in the other two games. In the PD cells, observed previous–round actions are the same as current–round actions 71.7% of the time, while messages are the same as current–round actions only 62.8% of the time. In the CH cells, observed previous–round actions are the same as current–round actions 66.7% of the time, while messages are the same as current–round actions only 54.4% of the time. The relatively small difference in "truthfulness" across observation cells is consistent with our hypothesis in Section 3 that the correlation between previous–round actions and current–round actions would not vary with the game. In contrast, the comparatively large

---

[22] One possible explanation for these findings is reciprocity. Notice that in all three games, signal receivers are more likely to play C when receiving a signal of C (either a message or an observed action) than when receiving a signal of D. Since we know from Tables 3, 4, and 5 that senders' signals tend to be truthful, we can hypothesize that receivers have some taste for cooperating with senders whom they believe are likely to cooperate, and defecting with senders whom they believe are likely to defect. Conditional on this assumption about preferences, the relationship between expected payoff differences and the frequency of cooperation now becomes perfect; for a given game and signal, the cell (communication or observation) with the higher expected payoff difference is always the cell with the higher frequency of cooperation.

differences in truth–telling across cheap talk cells are consistent with our hypothesis in the same section that the correlation between cheap talk messages and actions would vary with the game; indeed, in Stag Hunt, where messages are both self–signaling and self–committing, truth–telling was much more frequent than in the other two games.

The criterion of "truthfulness" is not the only measure of the value of a signal. An arguably more important criterion is the usefulness of signals in predicting senders' current–round actions, which we term the *functionality* of that signal. The functionality of a signal may be related to the truthfulness of the signal, but it need not be. For example, if Subject A cooperates 50% of the time following either signal, while Subject B always defects after sending a C signal and always cooperates after sending a D signal, Subject A's signals are more truthful (they are truthful 50% of the time, whereas B's never are), but Subject B's signals are more functional (his current–round actions can be perfectly predicted from his signal, while Subject A's signal provides no information at all). An advantage of functionality as a criterion is that it can be used in a more general setting than the one we consider—in situations where signals may not have literal meanings. Recall that in our experiment, both actions and messages in the communication cells are either "1" or "2" (R1 and C1 corresponding to C, R2 and C2 corresponding to D), so that a 1 message would reasonably be expected to correspond to a 1 action, and a 2 message to a 2 action. It would be easy to transform the communication treatment of our experiment by changing the 1 and 2 messages to, say, "@" and "#", while keeping the actions unchanged.[23] In this case, there would be no natural correspondence between messages and actions. It may happen that some convention (a one–to–one correspondence between $\{@, \#\}$ and $\{1, 2\}$) arose over time, but until this occurred, it would be problematic to speak of the "truthfulness" of signals. However, the "functionality" of signals would still be meaningful.

We now look at a way to measure the "functionality" of signals in either the cheap talk or observation treatment. We first note that sets of signals, for example the set $\{C,\ D\}$, can be thought of as candidate *forecast rules* for the sender's current–round action. We therefore examine some

---

[23] It is less clear how one might go about transforming the observation treatment in such a way that context was absent, while still allowing for some kind of observation to occur.

concepts that can be used to judge the quality of these forecast rules—calibration and resolution. *Calibration* is a measure of the accuracy of a forecast rule; for a well–calibrated set of signals, Defect signals should usually correspond to Defect actions, and Cooperate signals to Cooperate actions.[24] *Resolution* is a measure of the ability of a rule to partition the sample of prediction–event pairs into subsets in which the events are mostly of one type; for a given signal from a well–resolved set of signals, the actual frequency of Cooperate actions should be either close to zero or close to one. The specific notions of calibration and resolution we use are known as Sanders calibration (C_s) and Sanders resolution (R_s), defined as follows:

$$\begin{aligned} C_s &= Pr(\mathrm{C})(1 - Pr(\mathrm{C}|\mathrm{C}))^2 + Pr(\mathrm{D})(Pr(\mathrm{C}|\mathrm{D}))^2 \\ R_s &= Pr(\mathrm{C})Pr(\mathrm{C}|\mathrm{C})(1 - Pr(\mathrm{C}|\mathrm{C})) + Pr(\mathrm{D})Pr(\mathrm{C}|\mathrm{D})(1 - Pr(\mathrm{C}|\mathrm{D})) \end{aligned}$$

where, for a given set of signals, $Pr(\mathrm{C})$ ($Pr(\mathrm{D})$) is the probability of a C (D) signal being sent, and $Pr(\mathrm{C}|\mathrm{C})$ ($Pr(\mathrm{C}|\mathrm{D})$) is the probability of the current–round action being C after a C (D) signal is sent. Lower values of C_s and R_s reflect better calibration and resolution, respectively.

Calibration generalizes the notion of "truthfulness." When all signals can be interpreted as predictions of either C or D, calibration quantifies the accuracy of these predictions. One can also measure the calibration of probabilistic predictions, such as the forecast rule "always guess that C and D are equally likely." Calibration has the same disadvantage as truthfulness; in order to determine the calibration of a set of signals, we must be able to determine which signal corresponds to which action. In contrast, resolution quantifies our notion of "functionality." Resolution does not depend at all on the meaning of signals; if all C signals were changed to D signals and all D's changed to C's, resolution would be unaltered, while calibration would generally change. Also, if the possible signals were, for example, "@" and "#", the notion of calibration would be meaningless (unless some convention were present), while resolution could still be determined. In our hypothetical

---

[24] Yates (1982) provides a general description of the notions of calibration and resolution, as well as some of their mathematical properties. Murphy and Winkler (1977) use calibration and resolution to evaluate the quality of meteorologists' forecasts of the probability of precipitation, and Lichtenstein, Fischhoff, and Phillips (1982) do the same for the quality of students' judgements concerning general–knowledge questions (see also Camerer (1995, pp. 590–593)). Feltovich (2000) uses measures of calibration and resolution to test the forecasting ability of several models of individual learning in relation to experimental data.

example, Subject A, who cooperates 50% of the time following either signal, has better–calibrated signals than Subject B, who always defects after sending a C signal and always cooperates after sending a D signal. However, Subject B's signals are better–resolved.

Table 7: Effectiveness of Forecast Rules

Cell	Forecast Rule	Calibration (C_s)	Resolution (R_s)
PD–communication	messages	.164	.208
	assume C	.373	.238
	assume D	.151	.238
PD–observation	observed actions	.085	.198
	assume C	.370	.238
	assume D	.153	.238
SH–communication	messages	.017	.094
	assume C	.030	.144
	assume D	.682	.144
SH–observation	observed actions	.055	.145
	assume C	.060	.185
	assume D	.571	.185
CH–communication	messages	.216	.239
	assume C	.160	.240
	assume D	.360	.240
CH–observation	observed actions	.133	.200
	assume C	.099	.215
	assume D	.471	.215

Table 7 shows the calibration and resolution of each set of signals in each game. Also shown for each game, for comparison purposes, are the calibration and resolution of what we call "non–signal" forecast rules which might be used when no signals are sent (as in the control sessions); these are the forecast rules "assume C" and "assume D." Because observation and cheap talk never both took place in the same session, and because there were differences in the overall level of cooperation between observation and cheap talk cells, care must be taken in comparing the predictive value of observed actions vs. messages. However, we can compare the two types of signals indirectly by seeing how much improvement (or lack thereof) each type of signal yields over the "non–signals."

In Table 7 we see that in all three games, and in both cheap talk and observation treatments, the calibration of either type of signal is always better than that of at least one of the non–signals, though not always better than both. In the SH cells, both messages and observed actions are better–calibrated than the respective non–signals, but messages represent more of an improvement. In the PD cells, observed actions improve calibration more than messages, and messages are actually worse–calibrated than a forecast rule of "always D." In the CH cells, both messages and observed actions are less well–calibrated than a forecast rule of "always C," but messages worsen calibration more than observed actions. Also, the resolution of either type of signal is better than that of either type of non–signal, though the difference is sometimes small (as it is in the CH cells). In the SH cells, messages improve resolution more than observed actions do, relative to the corresponding non–signals. In the PD and CH cells, observed actions improve resolution more than messages.[25]

We thus see that the answer to the question, *do actions speak louder than words?* is, "it depends on the game." According to any of the criteria discussed in the previous section, in Stag Hunt, words are relatively more useful than actions, while in Chicken and Prisoner's Dilemma, actions are relatively more useful than words. It is worth noting that there is a correspondence between the usefulness of a set of signals (according to any of these criteria) and its ability to effect high–payoff outcomes; in all three games, the more useful set of signals is also the one that leads to higher payoffs (recall the last column of Table 2), though the difference in payoffs may not be significant. Indeed, it should not be surprising that both sets of signals lead to payoffs that are higher than when no signals are available, since both actions and words are at least somewhat informative in

---

[25] A remark about other types of non–signals is in order here. There is actually a continuum of such forecast rules, of the form "guess that the probability of C is p," which includes as special cases $p = 1$ and $p = 0$, the "assume C" and "assume D" rules used in Table 7. These rules all have the same resolution, but their calibration depends on p. If p is chosen to be the actual observed relative frequency of cooperation, the resulting forecast rule is *perfectly* calibrated, but it is no better for actually predicting the sender's current–round action than any of its 'siblings'. This is one reason to question the wisdom of using calibration as the sole criterion of the usefulness of signals.

all three games (in the sense that improvements in resolution represent extra information).

5.2 Discussion

The results of this experiment shed some light on the inadequacy of standard game theory and provide us with some insight as to how the theory might be extended. Two of the major deficiencies of game theory are (1) its indeterminacy of prediction in games with multiple equilibria, and (2) its inability to account for the degree of cooperation observed in social dilemmas. Experimental evidence and casual empiricism suggest that individuals often find ways in which to solve coordination problems and that cooperation is frequently observed. Our contention is that the presence in many environments of additional information, such as costless cheap talk or observation of prior behavior, plays a role in the solution of these coordination problems and in the fostering of cooperation. Game theory has little to say about how additional information may help to promote these "good outcomes." On the one hand, the theory tells us that additional information expands the set of equilibria to include some in which cooperation occurs, so that it is possible to theoretically rationalize cooperative outcomes. On the other hand, an expansion in the set of candidate equilibria only serves to exacerbate the problem of multiplicity of equilibria, leaving us to ponder how it is that individuals achieve coordination on a particular equilibrium.

6 Learning Model Analysis

The experimental findings suggest that static Nash equilibrium cannot explain the actions of subjects in the short span of time allowed by our experiment. Since certain aspects of subject behavior in several of our cells change from round to round (as seen in, for example, Figure 2), we conjecture that subjects are learning over time how to play these games. In this section, we develop a model of learning that allows subjects to adjust their behavior in response to some or all of the information they receive. When behavior in our learning model is initialized to match initial behavior in the experiment, the learning model tracks subject behavior reasonably well over the ten rounds of the experiment. We therefore feel justified in using this learning model to consider what happens over much longer periods of time as may be necessary for coordination on a Nash equilibrium to occur. Another purpose for developing this model is to determine whether, and how, agents use the additional information we provide to adjust their behavior in the communication and observation treatments.

6.1 The Basic Model

Our basic model is a special case of the reinforcement learning model of Roth and Erev (1995), which has been shown to be useful for characterizing aspects of subject behavior in previous experiments (See, for example, Roth and Erev (1995), Erev and Roth (1998), Slonim and Roth (1998), and Feltovich (1999)).[26] The play of a given Player 1 is as follows. (For ease of exposition, we develop the model from the point of view of Player 1; the model used for Player 2 play is identical.) The i-th Player 1 has *propensities* $q_i^t(C)$ and $q_i^t(D)$ for playing actions C and D in round t. The *strength of propensities* in round t is the sum of these propensities: $Q_i^t = q_i^t(C) + q_i^t(D)$. Player 1's probabilities of playing C ($p_i^t(C)$) and D ($p_i^t(D)$) are proportional to her propensities:

$$p_i^t(C) = \frac{q_i^t(C)}{Q_i^t} \text{ and } p_i^t(D) = \frac{q_i^t(D)}{Q_i^t}.$$

Propensities are updated as follows. After Player 1 plays round t of the game and receives payoff U_i^t, she adds that payoff to the corresponding propensity after multiplying both propensities by $1 - \delta$.[27] If Player 1 played C in round t,

$$q_i^{t+1}(C) = (1-\delta)q_i^t(C) + U_i^t \qquad \text{and} \qquad q_i^{t+1}(D) = (1-\delta)q_i^t(D).$$

---

[26] This is only one of several reasonable learning models. Other models that have been proposed, and tested against experimental data, include the "cautious fictitious play" model of Fudenberg and Levine (1995, 1998), and the "experience weighted attraction" model of Camerer and Ho (1995, 1999). We chose the Roth–Erev model because it has the fewest free parameters, and because it turns out to characterize subject behavior rather well.

[27] As usual in reinforcement learning models, U_i^t represents agent i's monetary payoff, or more precisely, her increase in expected monetary payoff. An alternative approach, suggested by the results of Section 5, would be to use a utility function that incorporated agents' tastes for "equitable" outcomes, such as the model of Fehr and Schmidt (1999). We chose not to do so for two reasons. First, generalizing utility functions would have added more free parameters to our model, and would have distracted attention from the roles played by cheap talk and observation in our experiment. Second, though behavior in the experiment was sometimes far from Nash equilibrium (such as in the PD cells), it was usually the case that behavior was at least moving in the direction of a Nash equilibrium, so that a model of learning that eventually (or asymptotically) reached Nash equilibrium would seem to be called for.

If Player 1 played D in round t,

$$q_i^{t+1}(C) = (1-\delta)q_i^t(C) \qquad \text{and} \qquad q_i^{t+1}(D) = (1-\delta)q_i^t(D) + U_i^t.$$

The parameter δ reflects the weight assigned to payoffs in recent rounds relative to earlier ones. If $\delta = 1$, then only payoffs from the most recent round matter; if $\delta = 0$, all rounds are given equal weight; if $\delta \in (0,1)$, more recent rounds are given more weight than earlier ones; and if $\delta < 0$, the reverse is true.

6.2 The Augmented Model

According to the basic model, agents adjust their strategies over time in response to their own current–round action and the resulting payoff. This *basic* learning model is well suited for simulating play in our control cells, where players are only informed of their current–round action and resulting payoff. In the other two cells, however, players have additional information: either cheap talk messages or opponent's previous–round actions. It is possible that our basic learning model, which does not take account of such information, nonetheless provides a good fit to the experimental data from the cheap talk and observation cells. However, this possibility rests on the assumption that the main difference across treatments is in the initial propensities; in the basic learning model, the manner in which agents adjust their play after the first round will be the same in all three information conditions.

An alternative possibility is that the effect of the information treatment is seen not only in agents' initial behavior, but also in changes in their behavior from round to round. For example, it might be the case that an agent who receives a C signal, responds with a C action, and earns a high payoff will become more likely to respond to a C signal with a C action in the future, but her likelihood of choosing a C action in response to receiving a D signal, or *sending* a C signal, will not change. In this case, the basic model is too simple; we need to enrich this model to account for agents' ability (should they desire) to condition their actions on the additional information provided in the communication and observation cells.[28] Our augmented model allows for these possibilities.

---

[28] Slonim and Roth (1998) examine whether changes in the scale of payoffs in a simple bargaining game affect just initial play, or both initial play and adjustment.

Allowing players to condition their choice of action on observed previous–round actions involves only a slight modification to the basic model. Players now have propensities to play C and D conditional on the information that is available. In our notation, the i-th Player 1 has propensities $q_i^1(C)$ and $q_i^1(D)$ concerning her play in round 1, propensities $q_i^t(C|C\ sent)$, $q_i^t(D|C\ sent)$, $q_i^t(C|D\ sent)$, and $q_i^t(D|D\ sent)$ concerning her play in round $t > 1$ after being observed to have chosen C or D in the previous round, and propensities $q_i^t(C|C\ received)$, $q_i^t(D|C\ received)$, $q_i^t(C|D\ received)$, and $q_i^t(D|D\ received)$ concerning her play in round $t > 1$ after observing her opponent to have chosen C or D in the previous round. Strengths of propensities, probabilities, and next–round propensities are computed as in the basic model, based on the relevant propensities. (As usual in reinforcement learning models, we assume that players are not forward–looking; they do not choose their round–t strategies taking into account that their round–$(t+1)$ opponent may observe them.)

Allowing players to condition their choice of action on cheap talk messages requires a more extensive modification to the model. For each player, in round 1 and in rounds in which that player *receives* a message, propensities and probabilities are analogous to those in the observation treatment. In rounds in which the player *sends* a message, however, we must determine not only the player's action, but also her message. As in the observation treatment, a player sending a message has propensities $q_i^t(C|C\ sent)$, $q_i^t(D|C\ sent)$, $q_i^t(C|D\ sent)$, and $q_i^t(D|D\ sent)$ concerning her own play in round $t > 1$ after sending her opponent a message of either C or D. She also has propensities for sending the messages themselves: $q_i^t(send\ C)$ and $q_i^t(send\ D)$. Propensities for sending C and D messages are updated according to the payoff for the round, just like propensities for actions, and probabilities of sending C and D messages are calculated analogously to those of actions.

6.3 Simulation Design

We conducted a simulated experiment in which the subjects were constrained to behave according to either the basic or the augmented learning model. We had two aims. First, we wanted to examine whether either learning model provided a better characterization of the experimental data than the static Nash equilibrium predictions. Second, we wanted to more closely examine the mechanism by which observation or cheap talk may affect subject behavior. If changes in the information treatment affect only the way in which subjects approach the game initially, and do not alter their behavior over time, then the basic model described in Section 6.1 should describe play in the three treatments at least as well as the augmented model described in Section 6.2. If, on the other hand, changes in the information treatment affect not only the way in which subjects approach the game initially, but also in the way they adjust their behavior over time, the augmented model should describe play in the cheap talk and observation treatments better than the basic model.

There are a total of five simulation cells: control/basic, communication/basic, communication/augmented, observation/basic, and observation/augmented. For initial propensities, we assume that round–1 probabilities are given by the average relative frequencies we found over the first three rounds of our experiment. This is akin to assuming that behavior in this environment is initially guided by some type of social norm, and over time affected by the player's experience. To obtain initial propensities, we then multiply each player's initial probabilities by a strength of initial propensities parameter Q^0. Following Roth and Erev (1995), who set the strength of initial propensities to a level roughly three times expected payoff, we set $Q^0 = 200$. To determine the appropriate value for δ, we performed a grid search for the value that generated simulated data that matched several features of the experimental data on a round–by–round basis. The criterion we used was the square root of the sum of the mean squared deviations from the experimental relative frequency of cooperation, the experimental relative frequency of coordination on a pure–strategy Nash equilibrium (in the SH and CH cells only), and the experimental payoff efficiency, in each of the ten rounds, and in each game and information condition. We restricted δ to be the same for all players, games, and information conditions, but we did not restrict δ to be the same in our basic model as in our augmented model.

6.4 Simulation Results

A total of 5000 simulations of each cell were performed (one simulation corresponds to ten rounds of each of the three games and one information condition, played by twenty subjects). Our grid search found that in the basic model, the best value of δ is 0.25, while it is 0.30 in the augmented model.[29] In order to assess the fit of our learning models to the experimental data, we compare their predictions with those of two alternative models in Table 8. The first, which we use as a baseline, is the "50–50 model": i.i.d. play in which Cooperate and Defect are played with equal probability. Since the predicted probability of cooperation is one–half, the predicted probability of coordination is one–half in the SH and CH cells, and the predicted payoff efficiency is one–third in the PD cell, two–fifths in the SH cell, and two–thirds in the CH cell. The second model, which we call the "no–learning" model, also assumes stationary play, but in this model the probability of cooperation is the same as that used to initialize the two learning models.

Table 8: Fit of Behavioral Models to Experimental Data (MSD)

Cell	Basic Learning Model	Augmented Learning Model	No–Learning Model	50–50 Model
PD–control	.052	.052	.065	.150
PD–communication	.043	.048	.053	.063
PD–observation	.072	.044	.082	.070
SH–control	.032	.032	.030	.094
SH–communication	.103	.048	.099	.213
SH–observation	.063	.074	.062	.134
CH–control	.032	.032	.031	.032
CH–communication	.042	.036	.042	.052
CH–observation	.044	.057	.044	.070
All PD cells	.099	.083	.117	.177
All SH cells	.125	.094	.120	.269
All CH cells	.069	.075	.068	.093
All control cells	.069	.069	.078	.180
All communication cells	.119	.077	.120	.228
All observation cells	.105	.103	.112	.167
All cells	.173	.146	.182	.335

---

[29] The level of precision we use is 0.05. Sensitivity analyses indicate that using a finer grid would not substantially improve closeness of fit to the data.

We report in Table 8 a measure of the goodness–of–fit of each of our four models to the experimental data. Our fitness criterion is the same one that was used to select values of δ for the learning models—the square root of the sum of the mean squared deviations (MSD) from the experimental relative frequencies of cooperation, coordination (when appropriate) and payoff efficiency. We also report in Table 8 the overall goodness of fit for each model for each game, each information treatment, and over the entire nine–cell experiment. Since lower values for our goodness–of–fit measure imply a better fit to the data, we see that the 50–50 model performs much worse than any of the other models overall, and for almost all of the individual cells. The no–learning model and the basic learning model give very similar results in the SH and CH cells (where play doesn't change drastically over time), but the basic model does better with the PD data (where cooperation and payoff efficiency decrease over time). The augmented learning model performs about the same or slightly worse than the basic model and the no–learning model in a few cells, but substantially better in others (particularly the cheap talk cells), and provides the best overall fit.

Since our augmented learning model provides a good quantitative fit to the experimental data, we now consider its *qualitative* performance. Figure 5 shows the mean frequency of cooperation in each of the nine cells, in each of the ten rounds as well as over all rounds. As before, the Nash equilibria are also shown, as horizontal lines marked by "x". Though the augmented learning model was the closest model to the experimental data, its success in qualitatively characterizing the data is mixed. The levels of cooperation seen in the simulation results have some features in common with the experimental results shown in Figure 2, but in some respects they differ. Consider first the PD cells. In both the experimental and simulated data, cooperation decreases over time in all three information conditions, there is little difference between the observation and cheap talk treatments, and cooperation is much less frequent in the control treatment. However, the decline in cooperation over time is much less pronounced in the simulated data than in the experimental data. In the simulated SH cells, cooperation occurs most frequently in the cheap talk cell, as was true in the experiment. However, the frequency of cooperation is greater in the control treatment than in the observation treatment, whereas the opposite was true in the experiment; also, cooperation does not increase over time in any of the SH simulation cells, whereas it increased in the experimental SH–

Figure 5: Average Relative Frequency of Cooperation (Augmented Model Simulations)

Rounds 1–10, 10–Round Averages, and Equilibrium Predictions

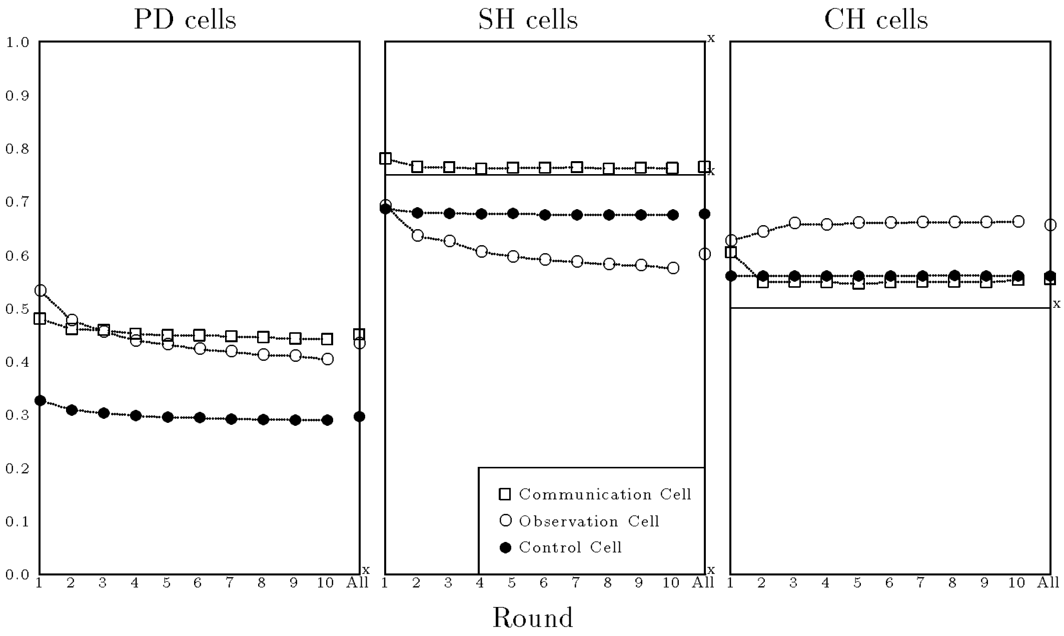

communication and SH–observation cells. In the CH simulations, there is little difference between the control and cheap talk cells, and the frequency of cooperation is highest in the observation cell, both of which were true of the experimental data.[30]

Since the augmented model appears to provide a reasonable quantitative and qualitative fit to the 10–round experimental data, we thought it might be useful to run the augmented model simulations out beyond the 10 rounds of the experiment. We can therefore examine how the augmented model, initialized using the experimental data, performs over spans of time that are not feasible in laboratory experiments with paid human subjects. The reason for conducting such an exercise is that the amount of time that may be necessary for the learning model to converge may be quite long, indeed, beyond the small amount of time available in an experimental session.

---

[30] The main *qualitative* differences in frequency of cooperation between the basic model and the augmented model are that according to the basic model, cooperation is more likely in the PD–communication cell than in the PD–observation cell, equally likely in the SH–control and SH–observation cells, and equally likely in the CH–control and CH–communication cells. Thus, the basic model does a little bit worse than the augmented model qualitatively.

Figure 6 shows, for each cell, the average relative frequency of cooperation from the continuation of 5,000 simulations of the augmented model at 10 rounds (as also shown in Figure 5), 100 rounds, 1000 rounds and 5000 rounds. Again, the Nash equilibria are also shown as horizontal lines marked by an “x”. We see that in all three PD cells, there is further movement toward the unique Nash equilibrium as the number of rounds in the simulated experiment are increased. By contrast there is not much change in the average frequency of cooperation in the SH cells and the CH cells as the number of rounds played is increased beyond the 10 rounds of the experiment. An exception is the SH–observation cell, where there is some further movement in the direction of the (D,D) Nash equilibrium.

Figure 6: Average Relative Frequency of Cooperation (Augmented Model Simulations)

Rounds 10, 100, 1000, and 5000, and Equilibrium Predictions

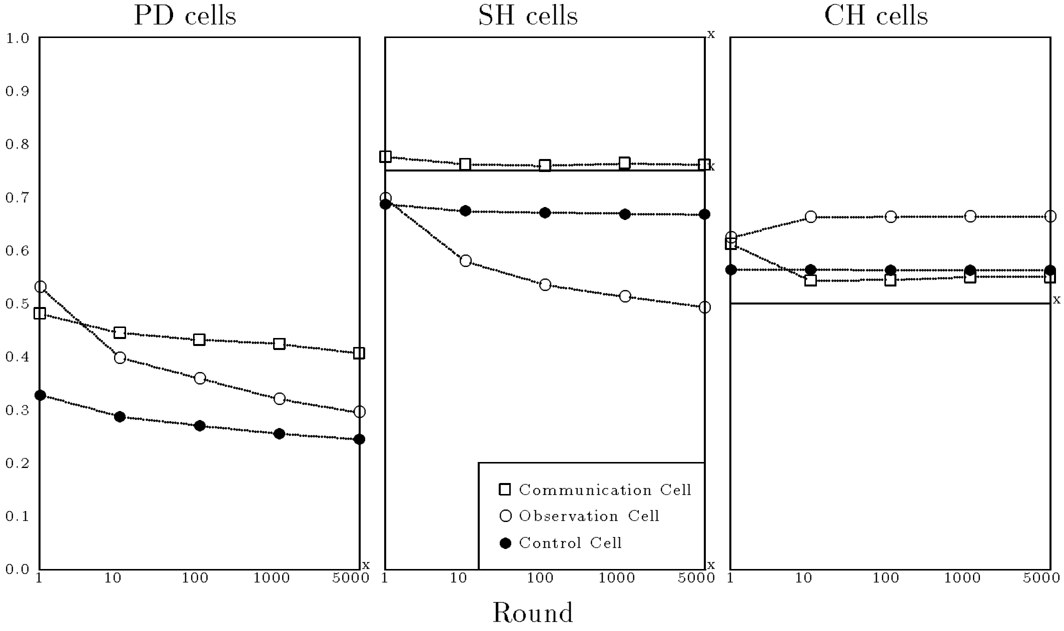

In summary, the augmented learning model, which assumes that agents not only learn over time, but also condition their choice of action on the information that is available (messages or previous–round actions), provides a better quantitative fit to three important features of the experimental data than either a model that assumes no learning over time or a model that assumes learning, but no conditioning. The model also performs reasonably well (but not perfectly) qualitatively, tracking the frequency of cooperation over time in a manner that is similar in several ways to the experimental data. Finally, there appears to be relatively little change in the frequency of cooperation as the number of rounds in the simulated experiment are increased beyond the 10 rounds allowed in the experiment, leading us to believe that additional rounds of experimental play would not have led to substantial changes in our main results.

7 Conclusion

Understanding how individuals achieve good outcomes in strategic situations is crucial to addressing a number of important questions in economics including issues of contract and mechanism design, the origin of standards and conventions, even the possibility of self–fulfilling macroeconomic fluctuations. Previous experimental studies have primarily focused on the role of cheap talk. In this paper we propose an alternative (and overlooked) mechanism fo r achieving good outcomes: observation of an opponent's past actions.

We design an experiment in which subjects play simple 2 × 2 games under one of three information treatments: no information about opponents, one–way cheap talk, or one–way observation of an opponent's previous–round action (when matched with a different player). The games we use differ substantially in the extent to which cheap talk ought to be credible. We find that both cheap talk and observation result in better outcomes—more cooperation, more coordination on pure–strategy Nash equilibria (in games with multiple equilibria) and higher payoffs—than when no additional information about an opponent is available, but that the effectiveness of cheap talk relative to observation in obtaining such good outcomes depends on the game played. In answer to the question posed by our title, *do actions speak louder than words*, we find that in games such as the Stag Hunt, where messages are credible, cheap talk is the relatively better device—words speak louder than actions. On the other hand, in games such as Chicken and Prisoner's Dilemma, where there is less (or no) reason to believe senders' messages, observation is relatively more effective than cheap talk—actions speak louder than words.

In an effort to better understand what is going on in our experiment, we developed a learning model that augments the reinforcement learning approach advocated by Roth and Erev (1995), to allow players to recognize when extra information is present and, if they wish, condition their choice of action on the available information. We find that this augmentation *does* improve the model; that is, the effect of changing the information condition is felt not only in changes in initial behavior, but also in the way subjects adjust their behavior over time. Simulations of our model yield patterns of play that are close to actual experimental subject behavior both in a quantitative sense and in a qualitative sense.

Economists typically assume that individuals' past histories of play are private information. However, in reality, players' recent past histories may be known to some or even all of the participants. The results of this study suggest that such knowledge of others' past history of play can be an important device in resolving certain types of strategic problems (achieving good outcomes). In particular, observation of past actions rather than cheap talk may be the relatively more useful coordination mechanism in situations where there are imperfect incentives for cheap talkers to send truthful messages. Furthermore, since incentive compatibility is a desideratum of mechanism design, allowing observation of past actions may be preferred to allowing preplay communication in certain kinds of strategic environments. We hope that the results of our experiment will encourage other researchers to consider the role played by observation of others' past actions when thinking about how economic agents go about solving coordination problems.

Appendix

The following instructions were used in all experimental sessions involving “one–way observation”. The instructions used in the “control” and “cheap talk” sessions were similar.

General Instructions

You are about to participate in an experiment in the economics of decision–making. Funding for this experiment has been provided by the National Science Foundation. If you follow these instructions carefully and make good decisions you might earn a considerable amount of money that will be paid to you in cash at the end of the session. If you have a question at any time, please feel free to ask the experimenter. We ask that you not talk with one another during the experiment.

This experimental session consists of three different games. Each game consists of a number of rounds. At the start of each game, you will be randomly assigned a player type, either “row player” or “column player.” Your type will not change during the course of game. In each round of this experiment you will be randomly matched to a player of the opposite type. You will be matched with a *different* player in every round of a game. We will refer to the person you are paired with in a round as your “partner.” Your score in each round will depend on your choice and the choice of your partner in that round. You will not know the identity of your partner in any round, even after the end of the session.

Sequence of Play in a Round

- At the beginning of each round, the computer program randomly matches each player to a partner.
- Following the first round of each game, and every round thereafter, one member of each newly matched pair of players will be able to observe the action their partner chose in the last round of play, when this partner was matched with another player. The other member of the pair will not be able to observe the action chosen by his partner in the last round of play. In each round, there is a 50% probability that you get to observe the action your partner chose in

the last round of play and there is a 50% probability that your partner gets to observe the action that you chose in the last round of play.

- You and your partner play the game. If you are a row player, you choose which row of the payoff table to play, R1 or R2. If you are a column player, you choose which column of the payoff table to play, C1 or C2.
- After all players have chosen actions, each player's payoff or *score* is revealed. Your score is determined by your action and the action of your partner according to the given payoff table.
- A random number is drawn. You earn $1.00 if your score is greater than or equal to this random number, and you earn nothing if your score is less than this random number.
- Provided that the last round of the game has not been reached a new round of the same game will then begin. You will be matched with a different partner in the new round.

The Games

The payoff table for each game you play will be shown on your computer screen and will also be drawn on the chalkboard. We will begin by playing the first game for 10 rounds. We will then play the second game for 10 rounds followed by 10 rounds of the third and final game. In every round of a game, both you and your partner have a choice between two possible actions. If you are designated as the row player, you must choose between actions R1 and R2. If you are designated as the column player, you must choose between actions C1 and C2. Your action together with the action chosen by your partner determines one of the four boxes in the payoff table. In each box, the first number represents your score and the second number represents your partner's score.

Observation

Beginning with the second round of each game, some players will be shown the action that their current partner chose in the previous round when matched with a different player. Either you will see your partner's action in the previous round, or your partner will see your previous round action.

These previous–round actions will be observed *before* players are called on to make decisions in the current round. We ask that you record the action you observed or the action your partner observed at the beginning of each round on your record sheets under the heading "Action Observed."

Earnings in Each Round

Your score (payoff) in a round is a number between 0 and 100. This is your percent chance of earning $1.00 in that round. Once your score is determined, we ask that you record your score for that round on your record sheet under the heading "Score." After all players have recorded their scores, a random number between 1 and 100 will be chosen and announced. Record the random number that is announced on your record sheet under the heading "Lottery Number." If your score is greater than or equal to the announced number (the lottery number), you earn one dollar ($1) or that round. If your score is less than the randomly chosen number, you earn zero for that round. Record your earnings for each round (either $1 or 0) in the last column of your record sheet under the heading "Earnings." Notice that the more points you earn in a round the greater is your probability of winning the $1 prize.

The Computer Screen

The top of your computer screen shows your player type, either row player or column player, your player ID number, and the round number. Please write your ID number on your record sheet. The middle of your computer screen contains the payoff table for the current game. Following this information is a prompt asking you to choose an action. If you are a row player you will choose between R1 and R2 and if you are a column player you will choose between C1 and C2. To indicate your choice, you type either 1 or 2 at the prompt and then you press the Enter key. After making your decision you must confirm your decision by typing Y for yes and then pressing Enter. If you want to change your decision, type N for no at the confirm prompt and then press Enter. If you choose not to confirm your decision you will have the opportunity to change the row or column you want to play. Following the first round, the bottom of the screen will show information about the results of your earlier rounds of play.

Payments

If you complete this experiment, you are guaranteed to receive a $5 participation payment. You will also be paid the sum of your earnings from each round. At the end of the session we will ask you to total up your earnings from all 30 rounds and record the sum at the bottom of your record sheet. All earnings will be paid in cash at the end of the session.

ARE THERE ANY QUESTIONS BEFORE WE BEGIN?

References

Arvan L., L. Cabral, and V. Santos. (1999), “Meaningful Cheap Talk Must Improve Equilibrium Payoffs,” *Mathematical Social Sciences* 37, pp. 97–106.

Aumann, R. (1990), “Nash Equilibria Are Not Self–Enforcing,” *Economic Decision–Making: Games, Econometrics, and Optimization: Contributions in Honor of Jacques H. Dreze*, J.J. Gabszewicz, J.-F. Richard, and L.A. Wolsey eds, North–Holland, pp. 201–206.

Camerer, C. (1995), “Individual Decision Making,” *Handbook of Experimental Economics*, J. Kagel and A. E. Roth. eds, Princeton University Press, pp. 587–703.

Camerer, C. and T.-H. Ho (1999), “Experience–Weighted Attraction Learning in Normal Form Games,” *Econometrica*, 67, pp. 167–188.

Cho, I.-K., and D. Kreps (1987), “Signaling Games and Stable Equilibria,” *Quarterly Journal of Economics* 102, pp. 179–221.

Cooper, R. (1999), *Coordination Games: Complementarities and Macroeconomics* New York: Cambridge University Press.

Cooper, R., D.V. Dejong, R. Forsythe, and T.W. Ross (1989), “Communication in the Battle of the Sexes Game: Some Experimental Results,” *Rand Journal of Economics* 20, pp. 568–587.

Cooper, R., D.V. Dejong, R. Forsythe, and T.W. Ross (1992), “Communication in Coordination Games,” *Quarterly Journal of Economics* 107, pp. 739–71.

Cooper, R., D.V. DeJong, R. Forsythe and T.W. Ross (1993), “Forward Induction in the Battle of the Sexes Game,” *American Economic Review* 83, pp. 1303–16.

Crawford, V.P. (1998), “A Survey of Experiments on Communication via Cheap Talk,” *Journal of Economic Theory* 78, pp. 286–98.

Crawford, V.P. and B. Broseta (1998), “What Price Coordination? The Efficiency–Enhancing Effect of Auctioning the Right to Play,” *American Economic Review* 88, pp. 198–225.

Crawford, V.P. and J. Sobel (1982), "Strategic Information Transmission," *Econometrica* 50, pp. 431–1451.

Dawes, Robyn M., 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, pp. 1–11.

Duffy, J. and N. Feltovich (1999a), "Does Observation of Others Affect Learning in Strategic Environments? An Experimental Study," *International Journal of Game Theory* 28, pp. 131–152.

Duffy, J. and N. Feltovich (1999c) "Words, Deeds and Lies," working paper, University of Pittsburgh and University of Houston.

Eckel, C. and P. Grossman (1996), "The Relative Price of Fairness: Gender Differences in a Punishment Game," *Journal of Economic Behavior and Organization* 30, pp. 143–58.

Erev, I. and A.E. Roth (1998), "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *American Economic Review*, 88, pp. 848–881.

Farrell, J. (1987), "Cheap Talk, Coordination and Entry," *Rand Journal of Economics* 18, pp. 34–39.

Farrell, J. and M. Rabin (1996), "Cheap Talk," *Journal of Economic Perspectives* 10, pp. 103–118.

Fehr, E., S. Gächter, and G. Kirchsteiger (1997), "Reciprocity as a Contract Enforcement Device: Experimental Evidence," *Econometrica* 65, pp. 833–860.

Fehr, E., and K.M. Schmidt (1999), "A Theory of Fairness, Competition and Cooperation," forthcoming, *Quarterly Journal of Economics.*

Feltovich, N. (1999) "Equilibrium and Reinforcement Learning in Private Information Games," *Journal of Economic Dynamics and Control*, 23, pp. 1605-1632.

Feltovich, N. (2000) “Reinforcement–Based vs. Beliefs–Based Learning Models in Experimental Asymmetric–Information Games,” forthcoming, *Econometrica.*

Fudenberg, D. and D. Levine (1995), “Consistency and Cautious Fictitious Play,” *Journal of Economic Dynamics and Control* 19, pp. 1065–1089.

Fudenberg, D. and D. Levine (1998), *The Theory of Learning in Games (Economics Learning and Social Evolution)*, MIT Press, Cambridge.

Haiman, J. (1998), *Talk is Cheap: Sarcasm, Alienation and the Evolution of Language*, Oxford University Press.

Holt, C. (1995), “Industrial Organization: A Survey of Laboratory Research,” *Handbook of Experimental Economics*, J. Kagel and A. E. Roth. eds, Princeton University Press, pp. 349–443.

Kahneman, D., J. Knetsch, and R. Thaler (1986), “Fairness and the Assumptions of Economics,” *American Economic Review* 80, pp. 218–233.

Kamecke, U. (1997), “Rotations: Matching Schemes that Efficiently Preserve the Best Response Structure of a One Shot Game,” *International Journal of Game Theory* 26, pp. 409–417.

Ledyard, J. (1995), “Public Goods: A Survey of Experimental Research,” *Handbook of Experimental Economics*, J. Kagel and A. E. Roth. eds, Princeton University Press, pp. 111–194.

Lichtenstein, S., B. Fischhoff, and L.D. Phillips (1982), “Calibration of Probabilities: The State of the Art to 1980,” *Judgement Under Uncertainty: Heuristics and Biases*, D. Kahneman, P. Slovic, and A. Tversky, eds, Cambridge University Press, pp. 306–334.

Mehlmann, A. (1997), *Wer Gewinnt das Spiel? Spieltheorie in Fabeln und Paradoxa*, Vieweg.

Murphy, A.H. and R.L. Winkler (1977), “Scoring Rules in Probability Assessment and Evaluation,” *Acta Psychologica* 34, pp. 273–286.

Orbell, J.M., A. van de Kragt, and R.M. Dawes (1988), “Explaining Discussion–Induced Cooperation,” *Journal of Personality and Social Psychology* 54, pp. 811–819.

Ostrom, E., R. Gardner and J. Walker (1994) *Rules, Games and Common–Pool Resources* Ann Arbor: University of Michigan Press.

Rabin, M. (1994), "A Model of Pre–Game Communication," *Journal of Economic Theory*, 63, pp. 370–391.

Roth, A.E. and I. Erev (1995), "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8, pp. 164–212.

Roth, A.E. and M.W.K. Malouf (1979), "Game–Theoretic Models and the Role of Bargaining," *Psychological Review* 86, 574–94.

Siegel, S. and N.J. Castellan, Jr. (1988), *Nonparametric Statistics for the Behavioral Sciences*, McGraw–Hill, New York.

Slonim, R. and A.E. Roth (1998), "Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic," *Econometrica*, 66, pp. 569–596.

Swensson, R.G. (1967) "Cooperation in the Prisoner's Dilemma Game I: The Effects of Asymmetric Payoff Information and Explicit Communication," *Behavioral Science* 12, 314–322.

Van Huyck, J., R.C. Battalio, and R.O. Beil (1993), "Asset Markets as an Equilibrium Selection Mechanism: Coordination Failure, Game Form Auctions and Tacit Communication," *Games and Economic Behavior* 5, 485–504.

Wilson, R. and C. Rhodes (1997), "Leadership and Credibility in *N*–Person Coordination Games," *Journal of Conflict Resolution*, 41, 767–791.

Wilson, R. and J. Sell (1997), " 'Liar, Liar...' Cheap Talk and Reputation in Repeated Public Goods Settings," *Journal of Conflict Resolution*, 41, 695–717.

Yates, J.F. (1982), "External Correspondence: Decompositions of the Mean Probability Score," *Organizational Behavior and Human Performance*, 30, 132-56.