

Words, Deeds and Lies: Strategic Behavior in Games with Multiple Signals

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 77204–5019, USA
nfelt@mail.uh.edu

September 6, 2004

Abstract

We report the results of an experiment in which subjects play games repeatedly against changing opponents. In one treatment, “senders” send messages to “receivers” indicating intended actions in that round, and receivers observe senders’ previous-round actions (when matched with another receiver). In another treatment, the receiver additionally observes the sender’s previous-round message to the previous opponent, enabling him to determine whether the sender lied in the previous round. We find that allowing more than one signal leads to better outcomes when signals are *aligned* (all pointing to the same action), but worse outcomes when signals are *crossed*. We also find that senders’ signals tend to be truthful, though the degree of truthfulness depends on the game and treatment, and receivers’ behavior combines elements of payoff maximization and reciprocity.

Journal of Economic Literature classifications: D83, C72, C73.

Keywords: observation, cheap talk, truthfulness, deception.

*Corresponding author. Funding for this project was provided by the National Science Foundation under grants SBR–9809622 and SBR–9813622. We thank Chris Anderson, Juergen Bracht, Nadège Marchand, participants at several conferences and seminars, and two anonymous referees for helpful comments and suggestions. Any remaining errors could have been observed in our previous paper.

1 Introduction

Many situations can be modeled as one-shot games in which players' interests are at least partially aligned. People in these situations often manage to coordinate successfully on one equilibrium when several exist. They may even obtain higher payoffs than in any equilibrium (for example, by overcontributing relative to equilibrium in public good situations). How do they achieve such good outcomes? One promising theory is that they use available information to determine the actions others are likely to choose.¹ This information can take many forms. Players may be able to communicate their intentions to others with costless nonbinding "cheap talk". Alternatively, they might use their opponents' past behavior to infer their likely future behavior.

This paper is an examination of these two types of information—cheap talk and observation. We design and run an experiment in which subjects play games against changing opponents under two information treatments. In our first, "words and deeds" (WD), treatment, each "sender" sends a cheap-talk message to her opponent, indicating the action the sender intends to play. The "receiver" also observes the sender's previous-round action choice (when matched with someone else). The receiver can use the sender's previous behavior to make inferences about the likely truthfulness of the sender's cheap-talk message, and thence the action the sender will likely choose in the current round. In our second, "words, deeds, and lies" (WDL) treatment, receivers see the cheap-talk message and previous-round action, plus a third piece of information: the sender's *previous-round message* to her then-opponent. Observation of the previous round in this treatment therefore consists of two different pieces of information: what action the sender actually chose, and whether the sender lied about her intentions to the previous receiver. Receivers are able, if they wish, to use the truthfulness (or lack thereof) of the previous-round message to evaluate the sender's current-round message.

Our examination consists of two distinct, though related, issues. In both cases, we build upon an earlier experiment (Duffy and Feltovich (2002)), which used the same games, but with three different information treatments: cheap talk only, observation of previous-round actions only, and a control treatment with neither cheap talk nor observation. One issue we examine is how the availability of multiple pieces of information, or "signals", affects the outcomes subjects reach.² There are two (not mutually exclusive) possibilities. Giving subjects additional information might enable them to make better decisions, so that the likelihood of good outcomes improves, relative to when they have less information. (This is part of the motivation behind our WDL treatment.) On the other hand, allowing multiple signals opens up the possibility that they might be "crossed"—for example, a current-round message different from the previous-round action. When signals are crossed, the receiver may have no better idea of the sender's likely action than if he had been given no information at all, so the resulting outcome may actually be worse than if only one piece of information had been available.³ The vast majority of theoretical and experimental treatments of signaling have examined situations with only one signal available. Some theorists have

¹Another possibility is that some outcome serves as a "focal point" for players; each player perceives that the other player will choose that outcome and so they choose it as well. See Sugden (1995) for a first effort in this direction. We consider this theory to be complementary, since the question then becomes one of how players recognize the existence of a focal point; such recognition may be easier if additional information is available for players to use.

²For ease of exposition, we use the term "signal" to encompass both messages and observed actions, and more generally, any piece of information sent to another player, whether intentionally or not, as long as both sender and receiver know that it is sent and received. We note here that signals in our experiment are costless and nonbinding.

³As an old expression goes, the man with one watch always knows what time it is, while the man with two watches never does.

looked at multiple signals, but there has been almost no experimental work designed to see how people actually behave in these situations. (Exceptions are discussed in Section 3.) We address the issue of multiple signals by comparing the results of our current experiment with those of the previous experiment.

The other issue we examine deals with the question: “do actions speak louder than words?” Our earlier paper attempted to answer this question indirectly, by comparing the behavior of senders and receivers in the cheap-talk-only treatment with those in the observation-only treatment. We found that the answer varied with the game, and specifically on whether its strategic structure provided incentives for senders to be truthful with their cheap-talk messages. When the game’s structure implied that cheap talk ought to be credible, cheap talk was more effective than observation—words spoke louder than actions. When its structure implied that cheap talk ought not to be credible, observation was more effective than cheap talk—actions spoke louder than words.

However, this earlier paper had the limitation that no *direct* comparisons between the two signaling devices were possible. Our current paper remedies this deficiency, by giving receivers both pieces of information. Our new treatments allow us to see whether receivers consider both types of information with equal weight, or alternatively, whether they tend to rely more on one type, and how, if at all, this depends on which game they are playing. We can also see whether receivers’ behavior is warranted, by examining the way senders’ actual choices are related to the signals they have sent—in particular, whether words or actions are better indicators of senders’ current-round behavior—and again, how this depends on the game. Additionally, the WDL treatment allows us to ascertain whether the behavior of senders and receivers depends on the ability of receivers to verify senders’ past truthfulness.

Our main findings are as follows. The new, multiple-signal, treatments lead to more cooperation, more coordination, and higher payoffs than when no signals are available (the control treatment). However, they do not improve upon the levels that were obtained when exactly one signal was available (the cheap-talk and observation treatments); in fact, they frequently make matters worse! This aggregate-level result masks sharp differences in outcomes following “aligned” signal combinations (signals all pointing to the same action) and those following crossed signal combinations. Outcomes following aligned signals are generally as good as, or better than, those in the one-signal treatments, while outcomes following crossed signals are worse, and comparable to those in the no-signal treatment. Even after accounting for changes in levels of cooperation, successful coordination is more likely, and average payoffs higher, after aligned signals than after crossed signals. In other words, it is not the *amount* of information people receive that is important, but rather its *content*.

At a more disaggregated level, we find that senders’ actions are positively correlated with their current-round messages and previous-round actions. In particular, messages tend to be truthful—even when the structure of the game provides incentives to lie—though truthfulness varies with the game and increases in the WDL treatment, where lies are detectable. Receivers, for their part, take into account the pieces of information they receive. The way they do so depends on the game and information treatment, but in all cases reflects a combination of payoff maximization and reciprocity: receivers cooperate more with senders who are likely to also be cooperative.

2 The games

We use three games: Prisoners’ Dilemma (PD), Stag Hunt (SH), and Chicken (CH) (see Figure 1). In each game, players choose between two strategies, which we label Cooperate (C) and Defect (D). These games were chosen because they are well-known, symmetric, 2×2 games in which choosing Cooperate always weakly increases the

other player’s payoff. The games differ in strategic aspects such as best–reply correspondence, in the number and nature of Nash equilibria, and in the predicted effect of cheap talk on play (as we will discuss in Section 3). Desirable outcomes for these games (from the players’ standpoint) include cooperation, coordination on a pure–

		Player 2		Player 2		Player 2					
		C	D	C	D	C	D				
Player	C	70,70	10,80	Player	C	70,70	10,55	Player	C	70,70	50,80
1	D	80,10	40,40	1	D	55,10	55,55	1	D	80,50	40,40
		<i>Prisoners’ Dilemma</i>				<i>Stag Hunt</i>				<i>Chicken</i>	

Figure 1: The games

strategy Nash equilibrium, and high payoffs. By “good outcomes” we refer to outcomes in which as many as possible of these features are present. Table 1 reports the extent to which they are present in the Nash equilibria

Table 1: Characteristics of Nash equilibria of the games

Game	Equilibrium	P(Cooperate)	P(Coordinate)	Expected payoffs	Efficiency
PD	(0,0)	0	—	(40,40)	.000
	(1,1)	1	1	(70,70)	1.000
SH	(0,0)	0	1	(55,55)	.600
	(.75,.75)	.75	.625	(55,55)	.600
CH	(1,0)	.5	1	(50,80)	.833
	(0,1)	.5	1	(80,50)	.833
	(.5,.5)	.5	.50	(60,60)	.667

Nash equilibria are presented in the form (P(row player cooperates), P(column player cooperates)).

of these three games. The overall frequency of C choices is found under the heading P(Cooperate). The probability of coordination, P(Coordinate), refers to the likelihood that players play a pure-strategy Nash equilibrium in the two games with multiple equilibria: (C,C) or (D,D) in Stag Hunt, and (C,D) or (D,C) in Chicken. Payoff efficiency is defined as the sum of row and column players’ payoffs, normalized so that the maximum possible joint payoff in a given game has an efficiency of one and the minimum has an efficiency of zero.

3 The information treatments: theory and hypotheses

From a theoretical standpoint, allowing signals need not affect outcomes in the games we consider. Because the games are finitely repeated, and this fact is public knowledge, the set of subgame perfect equilibria corresponds closely to the set of sequences of stage–game Nash equilibria. Allowing current–round messages does not increase the number of action sequences consistent with equilibrium, though the number of equilibria may increase due to the possibility of players’ conditioning their actions on the message. Previous–round actions (by themselves,

or in conjunction with previous-round messages), on the other hand, do open the possibility of equilibrium play containing action profiles that are not stage-game Nash equilibria: for example, (C,C) in Chicken.⁴

Aumann (1990) and Farrell and Rabin (1996) propose conditions for cheap talk to be truthful in situations where messages have literal meanings (some convention exists for translating each message into a unique intended action), as they do in our design. We will adopt Farrell and Rabin’s nomenclature here. Their conditions make use of the reasonable assumption that, if the receiver believes the sender’s message to be truthful (the same as the sender’s subsequent action), he will choose an action that is a best response to that message. Their first condition, *self-commitment*, is satisfied when the sender’s message is, in turn, a best response to the receiver’s action (so they form a Nash equilibrium). Their second condition, *self-signaling*, is satisfied when (a) a sender intending to be truthful prefers the receiver to play his best response, and (b) a sender who intended her message to be deceptive would not prefer the receiver to play his best response to the message. In Stag Hunt, both C and D messages are self-committing and self-signaling; in Chicken, both are self-committing but not self-signaling; and in Prisoners’ Dilemma, C messages are neither self-committing nor self-signaling, while D messages are self-committing but not self-signaling. Following Farrell and Rabin, who note that “a message that is both self-signaling and self-committing seems highly credible” (p. 112), we therefore expect that messages in Stag Hunt should most often be truthful and believed, messages in Chicken and D messages in Prisoners’ Dilemma should less often be truthful and believed, and C messages in Prisoners’ Dilemma should least often be truthful and believed.

By contrast with cheap talk, observed previous-round actions are credible by their very nature. However, observed actions also differ from messages in the extent to which they can be considered signals of the sender’s likely action. There is no question that a cheap-talk message is a signal; that is its only function. In contrast, while a previous-round action may be taken to be a signal by the receiver, that was not necessarily its intent. Even if observed, a previous-round action plays a dual role of signal for the next round and action choice for the current round; the latter role directly affects the sender’s payoff. Thus, the receiver of such a signal must bear in mind that though credible, it need not be a perfect forecast of the sender’s current-round intended action. Since this is equally the case in the three games of our experiment, we expect that the extent to which previous-round actions correlate with current-round actions will not vary systematically with the game. Therefore, we hypothesize that the relative efficacy of observation versus cheap talk in facilitating good outcomes should depend simply on how credible cheap talk is 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.

In addition to our earlier paper (Duffy and Feltovich (2002)), which gave evidence broadly supporting this hypothesis, we know of two previous papers that have attempted to compare observation and cheap talk.⁵ Wilson

⁴One way this can happen is as described by Okuno-Fujiwara and Postlewaite (1995), who look at infinitely-repeated games with discounting and random matching. In their model, each player carries a “status” with her, which is observable to opponents and can be updated in response to her actions. Her status functions as a proxy for her history of play, so that “punishment” strategies are possible even when players play each other only once. The resulting “norm equilibria” correspond to subgame perfect equilibria in standard infinitely-repeated games (with fixed opponents); in particular, mutual cooperation can be enforced. In our setup, players play only finitely many times, so mutual cooperation is not possible in an equilibrium in Prisoners’ Dilemma; however, it is possible in Chicken.

⁵There are many experimental papers that examine the role of either cheap talk or observation of the past actions of others in simultaneous-move games, but these papers do not examine them together. Papers that experimentally examine the role of cheap talk include Cooper et al. (1989, 1992), Charness (2000), Burton et al. (1999), and Blume and Ortmann (1999). See also Crawford (1998) for a survey of experiments involving cheap talk. Papers that consider how players react to the observed actions of other players

and Sell (1997) examined cheap talk and observation in a public-good game, in which the same group of subjects repeatedly decided how much to contribute to the 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 neither cheap talk nor observation were present. However, they also found that either cheap talk alone or observation alone actually *decreased* the amount of cooperation, so that there were substantial positive social returns to adding the second type of information, given that the first was already present. Their experimental setup was substantially different from ours, however (for example, their control treatment gave subjects no feedback at all, while our subjects learned their opponents’ current-round actions after they took place), so their results don’t carry much implication for our experiment. Çelen, Kariv, and Schotter (2003) examine observation and cheap talk in an information cascade experiment, where incentives are such that message senders’ interests are closely aligned with receivers’. (Senders received a payment if the receiver guessed correctly.) They found, among other results, that messages tended to be truthful and treated by receivers as such, that allowing cheap talk improves payoffs much more than allowing observation, and that allowing observation on top of cheap talk makes little further improvement (though adding cheap talk on top of observation does improve payoffs). Because this was a situation in which messages are expected to be extremely credible, their results are encouraging. In this paper, we go even further, examining the role of signals in strategic environments where cheap talk signals need not be so credible.

4 Experimental procedures

We used a 3×2 design in which we vary the order in which the games were played (PD-SH-CH, SH-CH-PD, or CH-SH-PD) and the information condition (WD or WDL).⁶ Each experimental session involved 20 subjects playing ten rounds of each game under a single information condition. Subjects were primarily University of Pittsburgh undergraduates. No subject participated in more than one session, nor did any participate in this experiment who participated in the experiment of our previous paper. In each game, ten of the subjects were row players, and the other ten were column players. Subjects were randomly assigned one of these roles and remained in the same role throughout a game. Each row player faced each column player exactly once in each game.⁷

Sessions were conducted in the Pittsburgh Experimental Economics Laboratory (PEEL), 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. The computer screen displayed the current game’s payoff matrix, the results of the player’s own previous rounds of play of that game,

include Kahneman et al. (1986), Eckel and Grossman (1996), Fehr et al. (1997), Duffy and Feltovich (1999), Huck et al. (1999, 2000), Bosch-Doménech and Vriend (2003), and Simonsohn et al. (2004).

⁶We actually split our WD treatment into “random” (WDr) and “nonrandom” (WDnr) subtreatments. In WDr, the roles of sender and receiver were determined randomly at the beginning of each round, as in our previous experiment. In WDnr, roles were determined randomly prior to the first round of a game and remained the same for all ten rounds played. This was done so that we could make direct comparisons with data from other treatments. The WDr treatment was chosen to facilitate comparison with the treatments from the previous experiment, while the WDnr treatment was chosen to facilitate comparison with the WDL treatment, where the presence of previous-round messages made it necessary to fix the roles of sender and receiver in all rounds of a game. As it turned out, we were unable to find any differences between the WDr and WDnr data; therefore, we simply pooled these data.

⁷Kamecke (1997) shows that the matching technique we used—“rotation”—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 the game.

and signals sent or received. The current payoff matrix was also drawn on a blackboard for all to see. Subjects input their actions by choosing which row or column of the payoff matrix they wanted to play. Row players' actions were labeled R1 and R2, and Column players' were labeled C1 and C2, in both cases corresponding to C and D respectively. In describing the actions to subjects we avoided reference to the labels "cooperate" or "defect," and we referred to a player's opponent as his or her "partner." Also, subjects were not given advice about how they should make use of the information they were given; for example, subjects in the WDL treatment were not told that previous-round actions might be used to assess the truthfulness of previous-round messages.

In the first round of each game, each row player was randomly paired with a column player. No signals were sent or received in the first round.⁸ In rounds 2–10 of each game, cheap talk and observation of previous-round actions took place before subjects chose their current-round actions. Observation of previous-round messages in the WDL sessions also took place at this time, from the third round on. After all subjects had chosen their actions for the current round, each subject was informed of her payoff and her opponent's action in that round.

Each point in the payoff matrix represented a 1% chance of winning \$1.00. At the end of each round, an integer between 1 and 100 inclusive was randomly drawn. Subjects whose payoff in that round was greater than or equal to the chosen number earned \$1.00 for the round; those whose payoff was lower earned nothing for the round.⁹ At the end of the session, subjects received in cash their total earnings from all rounds, as well as a \$5 show-up fee. Subjects earned an average of about \$25; sessions typically lasted between 60 and 75 minutes.

5 Results

The experiment consisted of six sessions each of the WD and WDL treatments; within either treatment, there were two sessions using each of the three orderings.¹⁰ We address the issue of differences in play due to changes in the ordering of the games in Sections 5.3 and 5.4; for now, we pool the data from sessions with different orderings, so that our cells are PD–WD, PD–WDL, SH–WD, SH–WDL, CH–WD, and CH–WDL.

5.1 Population aggregates

The aggregate frequencies of cooperation and coordination, as well as payoff efficiencies, are presented in Table 2. For comparison, we include corresponding results from the cheap-talk-only, observation-only, and control (neither cheap talk nor observation) treatments of Duffy and Feltovich (2002). Superscripts in the table refer to significance of differences between two information treatments; for a given game and statistic, entries sharing a superscript are not significantly different, while entries with letters earlier in the alphabet correspond to significantly lower values.¹¹ (For example, a statistic with a *b* superscript is significantly higher than one with an *a* superscript, but

⁸In the first round, observation of past actions is not possible. To maintain the symmetry of treatment of the two types of signal, we chose to suspend cheap talk in the first round as well. An implication of this feature of the design is that, in the WDL cell, previous-round messages could not be observed until the third round.

⁹This binary lottery procedure is intended to induce risk neutral behavior among hypothetical expected-utility maximizing agents. See, e.g., Roth and Malouf (1979) for a discussion.

¹⁰The instructions used in the experiment, and the raw data, are available from the authors upon request.

¹¹We 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

neither is significantly different from one with an *ab* superscript.)

Table 2: Relative frequencies of cooperation, coordination and efficiency (all rounds)

Game	Treatment	Cooperation	Coordination	Efficiency
PD	WD	.354 ^{ab}	—	.230 ^{ab}
	WDL	.391 ^b	—	.266 ^b
	<i>Cheap talk</i>	.400 ^b	—	.260 ^b
	<i>Observation</i>	.404 ^b	—	.266 ^b
	<i>Control</i>	.222 ^a	—	.113 ^a
SH	WD	.782 ^{ab}	.798 ^{bc}	.752 ^{abc}
	WDL	.823 ^b	.873 ^c	.828 ^c
	<i>Cheap talk</i>	.835 ^{ab}	.840 ^c	.803 ^{bc}
	<i>Observation</i>	.757 ^{ab}	.667 ^{ab}	.636 ^{ab}
	<i>Control</i>	.607 ^a	.513 ^a	.453 ^a
CH	WD	.577 ^{ab}	.507 ^{ab}	.746 ^{ab}
	WDL	.604 ^{ab}	.498 ^{ab}	.770 ^{ab}
	<i>Cheap talk</i>	.564 ^{ab}	.532 ^b	.741 ^{ab}
	<i>Observation</i>	.634 ^b	.438 ^a	.780 ^b
	<i>Control</i>	.537 ^a	.475 ^a	.696 ^a

Italicized treatments from Duffy and Feltovich (2002). Within each game and statistic, entries with no superscripts in common are significantly different at the 10% level (two-sided robust rank-order test, session-level data); superscripts earlier in the alphabet correspond to significantly lower values.

One of the main findings of our previous paper was that allowing either cheap talk or observation improved outcomes relative to the control: cooperation and coordination became more likely and payoff efficiency increased. We see from Table 2 that allowing both cheap talk and observation of previous-round actions (the WD treatment) almost never significantly improves outcomes over the control, never significantly improves upon either the observation-only or the cheap-talk-only treatments, and is often significantly worse than one or the other of these “one-signal” treatments. Adding observation of previous-round messages (the WDL treatment) improves matters a bit. Outcomes in the WDL treatment are never significantly worse than in any other treatment, they are often significantly better than those in the control, and sometimes significantly better than in the observation-only treatment. However, they are never significantly better than those in either the cheap-talk-only treatment or the WD treatment. Thus, while the value of allowing signaling is high if no signaling is currently available, incremental social returns to additional signals beyond the first are small or even negative.

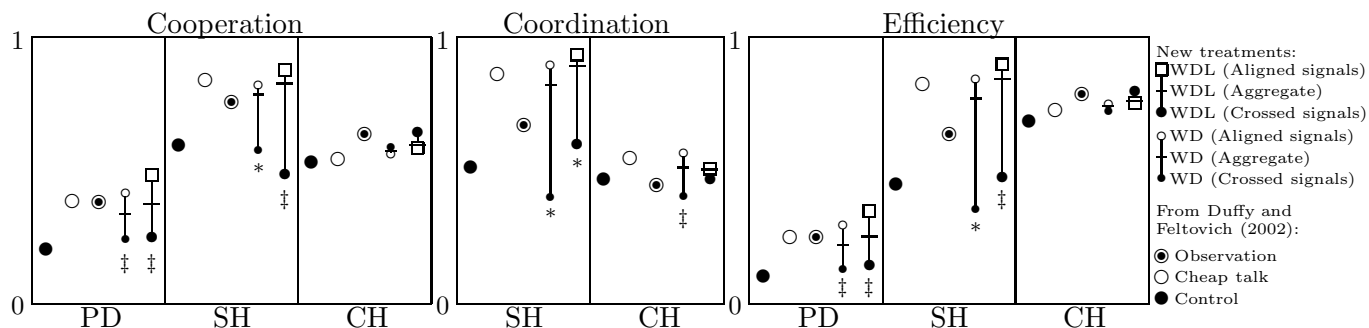
One explanation for this finding is that it is not the *amount* of information that matters, but rather its *content*.¹² When players have access to only one signal, the interpretation of the signal is unambiguous (though it need not

nonparametric statistical tests used in this paper. See Feltovich (2003) for a simulation-based comparison of the robust rank-order and Wilcoxon-Mann-Whitney tests, under a variety of distributional assumptions. All of our nonparametric statistical tests are performed on data at the *session* level, rather than at the level of the individual subject. While this has the disadvantage of substantially reducing the power of the tests, it avoids the problem of the likely interdependence of play among subjects in the same session.

¹²Another possible explanation is the phenomenon of “information overload” (Earl (1990)): as the *amount* of information provided increases, decision-makers tend to devote less careful attention to deciding whether they need to make changes to their strategies.

be truthful). With more than one signal, however, the potential exists for signal combinations with no clear implication, such as a ‘C’ message and ‘D’ observed action. Hence, we must distinguish between “aligned” (all ‘C’ or all ‘D’) and “crossed” (at least one ‘C’ and one ‘D’) signal combinations. Aligned signal combinations can be interpreted as a single signal (though again, not necessarily truthful), or even as a signal that’s been reinforced. On the other hand, crossed signal combinations may contain little or no information value.¹³

It is therefore reasonable to expect that when signals are aligned, resulting outcomes should be comparable to, or even better than, those in the earlier cheap talk and observation treatments (whose single signals are, by definition, aligned), while when signals are crossed, outcomes should be similar to those in the control treatment. If so, the overall lack of improvement from the earlier cheap talk and observation treatments to the WD and WDL treatments—seen in Table 2—may mask better outcomes following aligned signal combinations, and worse outcomes following crossed signal combinations. We address this issue in Figure 2, which shows the levels of



*: Sig. diff. between aligned and crossed signal combinations (one-tailed Wilcoxon summed-ranks test, session-level data, $p < 0.10$).
 ‡: Sig. diff. between aligned and crossed signal combinations (one-tailed Wilcoxon summed-ranks test, session-level data, $p < 0.05$).

Figure 2: Relative frequencies of cooperation, coordination and efficiency (rounds 2–10)

cooperation, coordination, and efficiency for rounds 2–10 of the WD and WDL treatments of each game, broken down according to whether the signals received were aligned or crossed. Also shown are the overall levels of cooperation, coordination, and efficiency for rounds 2–10 of the WD and WDL treatments, as well as the control, cheap talk, and observation treatments from Duffy and Feltovich (2002). In the PD and SH games, levels of cooperation, coordination, and efficiency are always significantly higher when signals are aligned than when signals are crossed. (Indeed, if WD and WDL session-level data are pooled, the difference is always significant at the 1% level.) When signals are aligned, outcomes are comparable to those from the earlier observation-only and cheap-talk-only treatments, but when signals are crossed, levels of cooperation, coordination, and efficiency are much lower and comparable to those from the control treatment. The CH results are more complicated. There are no significant differences ($p > 0.10$ for WD alone, WDL alone, and both together) in either cooperation or efficiency between aligned and crossed signal combinations. In fact, they are sometimes actually slightly (though not significantly) lower after aligned signals than after crossed signals. Coordination in the CH-WD cell after

¹³In the WDL treatment, some crossed signal combinations may carry information. For example, a receiver viewing the combination of D previous-round message, D previous-round action, and C current-round message may reason that, because the sender was truthful in the previous round, she will be truthful in the current round also, and therefore her current-round action should be C. In Section 5.4, we look at behavior following such signal combinations in the WDL treatment. For now, we simply point out that, even if not completely uninformative, they are likely to be less informative than aligned signal combinations.

aligned signal combinations is significantly higher than after crossed signal combinations, and is comparable to coordination in the earlier observation and cheap talk treatments. However, there is no significant difference in coordination between aligned and crossed signal combinations in the CH–WDL cell. (If CH–WD and CH–WDL data are pooled, coordination after aligned signal combinations is higher than after crossed signal combinations at the 5% level.) For the most part, though, it does appear that aligned signal combinations lead to especially high levels of cooperation, coordination, and efficiency, while crossed signal combinations lead to little or no improvement over no signals at all.

5.2 A closer look at coordination and efficiency

We saw in Table 2 and Figure 2 that coordination and efficiency were higher in the WD and WDL treatments than in the control, and were particularly high following aligned signal combinations. We will see now that these levels are high not only in an absolute sense, but also relative to what would have been expected *given the observed frequency of cooperation*.

First, consider coordination. In Stag Hunt, coordination means play by the two players of either a (C,C) or a (D,D) strategy pair. *If players are not able to make use of the extra information available in our treatments*, then sender and receiver actions should be uncorrelated. In that case, the observed frequency of (C,C) pairs should be approximately equal to the observed frequency of senders’ C choices multiplied by the observed frequency of receivers’ C choices, and the observed frequency of (D,D) pairs should be roughly the observed frequency of senders’ D choices multiplied by the observed frequency of receivers’ D choices. Hence, if q_1 and q_2 denote the observed frequencies of cooperation in a SH cell by senders and receivers, respectively, then the observed frequency of coordination ought to be $q_1q_2 + (1 - q_1)(1 - q_2)$. We will call this the *predicted* frequency of coordination.¹⁴ In Chicken, coordination means play of either a (C,D) or a (D,C) pair, so that if r_1 and r_2 are the observed frequencies of cooperation in a CH cell by senders and receivers, the *predicted* frequency of coordination will be $r_1(1 - r_2) + r_2(1 - r_1)$. In both games, if players are able to use their information well, then the observed frequencies of coordination will be higher than these predicted frequencies.

Next, consider efficiency. In Prisoners’ Dilemma, the efficiency of a (C,C) outcome is 1, that of a (D,D) outcome is 0, and that of a (C,D) or (D,C) outcomes is $\frac{1}{6}$. So, if s_1 and s_2 are the observed frequencies of cooperation by senders and receivers, and their actions are not correlated with each other, then the *predicted* level of efficiency is $s_1s_2 + \frac{1}{6}[s_1(1 - s_2) + (1 - s_1)s_2]$. Similarly, predicted levels of efficiency are $q_1q_2 + \frac{3}{5}(1 - q_1)(1 - q_2)$ in Stag Hunt and $r_1r_2 + \frac{5}{6}[r_1(1 - r_2) + (1 - r_1)r_2]$ in Chicken. Comparison of the observed levels of efficiency to these predicted levels will tell us whether the information subjects are given enables them to earn especially high payoffs, beyond what would be expected based on the observed frequency of cooperation.

The left panel of Figure 3 shows predicted and observed levels of coordination (for the SH–WD, SH–WDL, CH–WD, and CH–WDL cells) and efficiency (for all six cells) for rounds 2–10; the center and right panels show these levels following aligned and crossed signals. Also shown is the 45° line, where predicted and observed levels are equal. If there is no systematic relationship between the actions of Row and Column players, symbols should

¹⁴There is some room for confusion here, since “predicted” can also be used to refer to the theoretical (Nash equilibrium) levels of coordination and efficiency, shown in Table 1. We emphasize that the predicted levels of coordination and efficiency discussed in this section are derived from the observed levels of cooperation by sender and receiver, the assumption of zero correlation between sender and receiver actions, and (in the case of efficiency) the game’s payoff matrix.

be centered on the 45° line. The left panel shows that overall, this is not the case; observed levels of coordination and efficiency are higher than predicted levels in every cell, though sometimes the difference is small. As before,

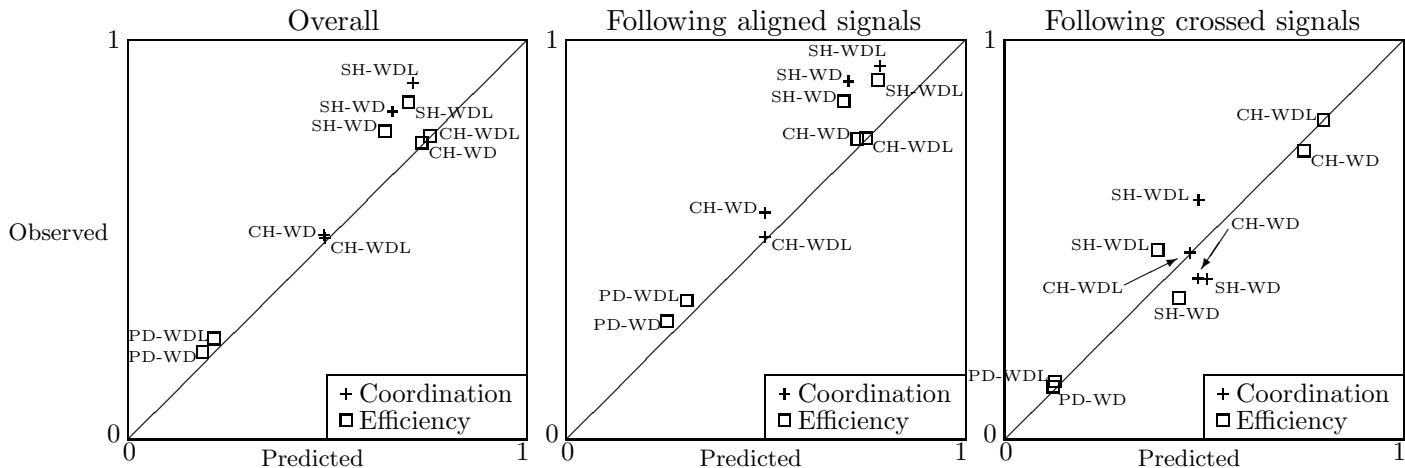


Figure 3: Predicted and observed coordination and efficiency, rounds 2–10

the content of information is what determines whether subjects are able to coordinate and earn high payoffs. Aligned signal combinations lead to an improvement of observed over predicted coordination and efficiency in every game and information treatment. Following crossed signal combinations, in contrast, observed coordination and efficiency are higher than predicted only about half the time (6 cases out of 10).

To further illustrate the relationship between predicted and observed frequencies of coordination and efficiency, and to test significance, we disaggregate the data further. Figure 4 shows the predicted and observed frequencies of coordination for Stag Hunt and Chicken (left and left-center panels) and efficiency for all three games (right-center and right panels), broken down by session and by whether signals were aligned or crossed. As in Figure 3,




Figure 4: Predicted and observed coordination and efficiency, session-level data, rounds 2–10

we see here that following aligned signals, observed coordination and efficiency are typically higher than would be predicted based on the observed levels of cooperation; this is true both overall (one-sided Wilcoxon summed-ranks test, session-level data, $p < 0.001$) and for each game individually ($p < 0.01$). On the other hand, there are no significant differences between predicted and observed efficiency following crossed signals either overall or for either Stag Hunt or Chicken individually ($p > 0.10$), though observed efficiency in Prisoners' Dilemma following crossed

signals *is* significantly higher than predicted ($p < 0.05$).

To summarize, aggregate and session-level frequencies of coordination and payoff levels are greater than what would be expected based on observed levels of cooperation *and the assumption of no correlation between row and column player actions*. However, these overall increases are comprised of large increases following aligned signal combinations, and no increase following crossed signal combinations. When signals are aligned, subjects are relatively successful in coordinating on pure-strategy equilibria and achieving high payoffs. When signals are crossed, subjects fail to coordinate any more often than they would at random, and payoffs are appropriately low.

5.3 Individual behavior I: do actions speak louder than words?

These results—(1) players are able to coordinate more successfully and earn higher payoffs than would be expected based on the observed frequency of cooperation, and (2) the extent to which they can do so depends partly on whether signals are aligned or crossed—imply that senders and receivers condition their actions on the signals sent and received. We now examine how they do this. Table 3 shows senders’ and receivers’ relative frequencies of cooperation, conditional on the current-round message and on the observed previous-round action.

Table 3: Relative frequencies of cooperation, conditional on signals sent or received (rounds 2–10)

Senders	PD-WD	PD-WDL	SH-WD	SH-WDL	CH-WD	CH-WDL
C message	.449** (176/392)	.552** (212/384)	.920** (402/437)	.920* (427/464)	.749** (218/291)	.810** (234/289)
D message	.135 (20/148)	.154 (24/156)	.223 (23/103)	.250 (19/76)	.301 (75/249)	.187 (47/251)
C obs. action	.619** (125/202)	.773** (197/255)	.938** (391/417)	.964** (429/317)	.691** (219/317)	.809** (237/293)
D obs. action	.210 (71/338)	.137 (39/285)	.276 (34/123)	.179 (17/95)	.332 (74/223)	.178 (44/247)
Receivers	PD-WD	PD-WDL	SH-WD	SH-WDL	CH-WD	CH-WDL
C message	.380** (149/392)	.372** (143/384)	.913** (399/437)	.920* (427/464)	.574 (167/291)	.654 (189/289)
D message	.115 (17/148)	.160 (25/156)	.233 (24/103)	.276 (21/76)	.635 (158/249)	.677 (170/251)
C obs. action	.485** (98/202)	.424** (108/255)	.904** (377/417)	.930** (414/445)	.571 (181/317)	.659 (193/293)
D obs. action	.201 (68/338)	.211 (60/285)	.374 (46/123)	.358 (34/95)	.646 (144/223)	.672 (166/247)

*: *C* is signif. more likely after a *C* than after a *D* signal (one-tailed Wilcoxon summed-ranks test, session-level data, $p < 0.10$).

** : *C* is signif. more likely after a *C* than after a *D* signal (one-tailed Wilcoxon summed-ranks test, session-level data, $p < 0.05$).

Several patterns stand out. First, the top half of the table shows that senders’ signals—both messages and observed actions—are quite useful in forecasting their current-round actions. In fact, both types of signal tend to be truthful.¹⁵ For all three games and both information treatments, and for both messages and observed actions, senders are substantially—and significantly—more likely to choose *C* after sending a *C* signal than after sending a *D* signal. It is perhaps not surprising that senders’ previous-round actions are truthful, as this may simply mean that players’ action choices are positively autocorrelated (as indeed they are, both senders’ and receivers’). What is surprising is that cheap-talk messages are so truthful, even in the PD cells, where messages are not expected to be credible. Table 4 shows the overall frequencies of truthful messages, observed actions, and aligned pairs (current-round message and observed previous-round action that are either both *C* or both *D*) for each cell. The

¹⁵We call a signal “truthful” if it matches the subsequent current-round action; for example, a *C* current-round message is truthful if the current-round action is also *C*. It seems uncontroversial to use this term for current-round messages, but there is a slight abuse of vocabulary in using it to refer to an observed previous-round action that matches the current-round action, since as mentioned in Section 3, observed actions have other purposes besides signaling.

overall frequency of truthful messages varies substantially across games and information treatments, but is well over one-half in all cells, even in the PD cells. Notice also that aligned pairs are even more informative than either messages or observed actions alone.

Table 4: Senders’ frequencies of truthful words, deeds, and aligned pairs

Game	WD treatment			WDL treatment		
	words	deeds	aligned pairs	words	deeds	aligned pairs
PD	.563	.726	.773	.637	.820	.864
SH	.893	.889	.961	.896	.939	.971
CH	.726	.681	.811	.811	.815	.906

If receivers understand how senders’ signals correlate with their subsequent actions, they ought to condition their own actions on these signals. Indeed, they do so. However, while senders’ behavior can be concisely, and reasonably accurately, described as truthfulness, receivers respond to senders’ signals in a more complex manner that combines aspects of reciprocity and best response. By “reciprocity”, we mean choosing a cooperative action when matched with a sender who is expected to choose a cooperative action, and choosing an uncooperative action when the sender is expected to choose an uncooperative action. In our three games, choosing C weakly increases the other player’s payoff relative to choosing D, so we consider C to be the cooperative action (for both players, since the games are symmetric). Reciprocity therefore implies that in all three games, the receiver chooses the same action as the one he expects the sender to choose. Best responses, on the other hand, vary by the game (recall Figure 1). In Prisoners’ Dilemma, the best response to either action is D; in Stag Hunt, the best response to either action is the same action; in Chicken, it is the opposite action.

Because senders’ current-round messages and previous-round actions tend to be truthful, we can combine the predictions of reciprocity and payoff maximization to predict how receivers should react to senders’ signals. In Prisoners’ Dilemma, payoff maximization implies no difference between responses to C signals and those to D signals, so that the only effect is that of reciprocity. Receivers should tend to choose C more often in response to C signals than D signals. In Stag Hunt, reciprocity and payoff maximization point in the same direction, so their effects reinforce each other. Receivers should tend to choose C in response to C signals more often than in response to D signals, possibly even more so than in Prisoners’ Dilemma. In Chicken, payoff maximization and reciprocity point in opposite directions, and will work to cancel each other out. We should thus expect little or no difference between responses to C signals and to D signals. This is exactly what the bottom half of Table 3 shows. In the SH and PD cells, receivers are substantially (and significantly) more likely to cooperate after seeing a C signal than after seeing a D signal. In the CH cells, on the other hand, there is no strong relationship between signal and receiver action; receivers actually tend to choose the opposite action slightly more than the same action.¹⁶

In order to assess the effects of combinations of signals, we estimate probit regressions in which the binary

¹⁶For messages in the CH-WD cell, receivers are actually significantly more likely to choose C following a D message than following a C message (two-tailed Wilcoxon summed-ranks test, session-level data, $p < 0.10$), suggesting that best response might be stronger than reciprocity in this case; however, see the discussion of Table 5 below. For observed actions, and for messages in the CH-WDL cell, there is no significant difference between responses to C and D signals.

dependent variable is whether a player chose to play C. We focus on the role played by *two*-signal combinations of current-round message and previous-round action, in both the WD and WDL treatments. (Later, in Section 5.4, we add *three*-signal combinations to our regressions using data from the WDL treatment only.) Independent variables in our probit regressions include a constant, the round number (to control for any changes over time), a dummy “wdl” equal to 1 if the data came from the WDL treatment, dummies “shfirst” and “chfirst” equal to 1 if the order of the games played was SH-CH-PD or CH-SH-PD, and “paydiff”, a measure of the difference in expected payoff between a C choice and a D choice.¹⁷ Finally, we included dummy variables CC, CD, and DD, that were equal to 1 whenever the (message, observed action) combination was CC, CD or DD.¹⁸ We also estimated this model under the restriction that the three signal-combination dummies were all zero, and we compare the performance of this restricted model relative to the unrestricted model using likelihood-ratio tests.

Table 5: Probit results—frequency of cooperation (WD and WDL cells, rounds 2–10)

	Prisoners' Dilemma		Stag Hunt		Chicken		Prisoners' Dilemma		Stag Hunt		Chicken	
	Senders ($N=1080$)		Senders ($N=1080$)		Senders ($N=1080$)		Receivers ($N=1080$)		Receivers ($N=1080$)		Receivers ($N=1080$)	
const.	1.195 (0.964)	3.670*** (1.042)	-0.179 (0.315)	1.130*** (0.304)	-0.420 (0.261)	0.362 (0.328)	1.051 (1.040)	2.568*** (0.919)	-0.572* (0.328)	0.529*** (0.165)	0.005 (0.225)	0.060 (0.197)
round	-0.073*** (0.026)	-0.064** (0.028)	-0.044 (0.031)	0.039 (0.032)	-0.029 (0.021)	-0.029 (0.024)	-0.093*** (0.027)	-0.073*** (0.024)	0.055** (0.028)	0.038** (0.019)	0.021 (0.018)	0.021 (0.018)
wdl	0.155 (0.181)	0.228 (0.280)	0.005 (0.215)	0.172 (0.251)	0.089 (0.196)	-0.144 (0.382)	-0.101 (0.218)	-0.103 (0.190)	0.106 (0.230)	0.095 (0.119)	0.136 (0.165)	0.154 (0.165)
shfirst	-0.292 (0.229)	-0.304 (0.423)	0.474 (0.365)	1.896*** (0.409)	0.396 (0.259)	-0.181 (1.952)	-0.395 (0.282)	-0.308 (0.241)	0.353 (0.371)	0.234 (0.195)	0.312 (0.215)	0.316 (0.215)
chfirst	0.061 (0.212)	0.048 (0.299)	0.403* (0.235)	1.357*** (0.317)	0.220 (0.226)	-0.590 (0.437)	-0.495** (0.256)	-0.390* (0.221)	0.334 (0.267)	0.146 (0.137)	-0.078 (0.193)	-0.080 (0.193)
paydiff	0.092** (0.048)	0.179*** (0.056)	0.030* (0.016)	0.081*** (0.019)	-0.024 (0.044)	-0.082 (0.120)	0.064 (0.051)	0.122*** (0.046)	0.029* (0.018)	0.054*** (0.038)	-0.078** (0.0381)	-0.065* (0.038)
CC	1.589*** (0.224)	—	2.421*** (0.278)	—	1.124*** (0.170)	—	1.106*** (0.227)	—	2.169*** (0.268)	—	-0.087 (0.137)	—
CD	0.652*** (0.235)	—	0.738** (0.324)	—	1.196*** (0.217)	—	0.287 (0.223)	—	0.707*** (0.273)	—	0.128 (0.173)	—
DD	-0.044 (0.242)	—	-0.887*** (0.326)	—	-0.857*** (0.210)	—	-0.668*** (0.248)	—	-1.347*** (0.292)	—	0.175 (0.142)	—
-ln(L)	495.91	553.34	193.87	257.22	461.73	542.21	499.80	569.33	268.91	466.15	629.10	632.27
p -value‡	< 0.001		< 0.001		< 0.001		< 0.001		< 0.001		0.097	

* (**, ***): Coefficient is significantly different from zero at the 10% (5%, 1%) level.

‡: Likelihood-ratio test of no difference between unrestricted and restricted models.

Note: Signal combinations are (current-round message, previous-round action).

Coefficient estimates are shown in Table 5. In order to account for possible individual effects, standard errors were corrected for clustering (the assumption of independence of observations over time is relaxed for individual

¹⁷Expected-payoff differences were calculated separately for each game, session, and round, as follows. First, we found the relative frequency of C choices and that of D choices in that game and session, over all previous rounds. Next, we treated this pair of relative frequencies as a population mixed strategy, and calculated the expected payoffs to C and D based on an opponent using that mixed strategy. The expected-payoff difference was the expected payoff to C minus the expected payoff to D.

¹⁸To avoid perfect collinearity, we have to leave out the DC signal dummy. Analogously, we do not have a dummy for the third game ordering, PD-SH-CH.

subjects, though not across subjects). Corrected standard errors are shown in parentheses.¹⁹ We also show log-likelihoods (for each regression) and p -values from the likelihood-ratio test of the joint significance of the signal-combination dummies (for each pair of regressions). The results show several regularities. The coefficients for the signal combinations are jointly significant in each case, and are often significant when considered individually. (The main exception is in the regression for Chicken receivers, where no signal coefficient is individually significant, though the three together are jointly significant at the 10% level.) The signs on the signal coefficients reveal that the CC combination tends to have a significant, positive effect on cooperation, while the DD combination has a significant, negative effect. The effect of a CD combination is either significant and positive, or insignificant, depending on the game and whether we are considering senders or receivers. For senders, the coefficient is significant and positive for all three games. This suggests that cooperation by senders is significantly higher following a CD combination than following a DC combination (the baseline case); that is, words speak louder than actions—even in Prisoners’ Dilemma, where cheap talk should not be credible. For receivers, the coefficient for the CD dummy is significant and positive in Stag Hunt but not in the other two games.

Since the signal-combination coefficients are jointly significant in every case, we concentrate on the unrestricted regressions for our discussion of the other variables. (Results are often, though not always, qualitatively similar in the restricted regressions; see Table 5.) The effect of the round number varies. In Prisoners’ Dilemma, it is significant and negative, consistent with many other studies which have shown that cooperation declines over time. In Stag Hunt, it is not significant for senders, but significant and positive for receivers. In Chicken, it is not significant for either senders or receivers. The coefficient on the WDL dummy is never significant; players’ behavior in the WD and WDL treatments is apparently quite similar, once other factors are controlled for. Coefficients for the game-ordering dummies are seldom significant. Coefficients for payoff difference are significant in four cases out of six; when significant, they are positive except for receivers in Chicken. A positive coefficient is consistent with payoff maximization, to the extent that our payoff-difference variable reflects true expected payoffs. A negative coefficient in Chicken is not completely surprising; as mentioned earlier, Chicken is the one game in which reciprocity points in the opposite direction from payoff maximization, so the negative sign suggests simply that reciprocity is relatively powerful here.

5.4 Individual behavior II: lies, damned lies, and statistics

We saw in Section 5.1 that aggregate behavior in the WDL treatment was somewhat different from that in the WD treatment: cooperation and coordination were more likely, and average payoffs higher (though differences were usually not significant) relative to WD. There are two primary ways in which the addition of information about previous-round messages might lead to improvements in outcomes. It could be that receivers in the WDL treatment, who can judge the veracity of senders’ previous-round messages—unlike receivers in the WD treatment, who can’t—are better able to evaluate senders’ current-round messages and choose their own actions accordingly. A second possibility is that senders, knowing that their truthfulness will be observed in the next round, might choose their messages and actions differently. (These possibilities are not mutually exclusive.) In this section, we take a closer look at the behavior of senders and receivers in the WDL treatment.

Table 6, which shows the frequency of senders’ and receivers’ C choices conditional on each three-signal

¹⁹We have also examined probit regressions with random effects and obtained results that are very similar to those reported here.

combination, gives some evidence that information about the sender’s past truthfulness (or lack thereof) is useful in predicting her current actions, above and beyond the information present in the current-round message and previous-round action, and that this information seems to be acted on by receivers. Both senders and receivers are

Table 6: Conditional relative frequencies of cooperation (WDL cells, rounds 3–10)

Combination (prev. message, curr. message, prev. action)	Relative frequency of cooperation (conditional on that signal combination)					
	PD		SH		CH	
	Senders	Receivers	Senders	Receivers	Senders	Receivers
1 (CCC)	.849 (152/179)	.430 (77/179)	.992 (374/377)	.963 (363/377)	.906 (173/191)	.654 (125/191)
2 (CCD)	.075 (7/93)	.172 (16/93)	.056 (1/18)	.444 (8/18)	.500 (13/26)	.692 (18/26)
3 (CDC)	.263 (5/19)	.158 (3/19)	.444 (4/9)	.333 (3/9)	.481 (13/27)	.556 (15/27)
4 (CDD)	.152 (7/46)	.174 (8/46)	.125 (1/8)	.500 (4/8)	.400 (4/10)	.800 (8/10)
5 (DCC)	.643 (9/14)	.500 (7/14)	.857 (6/7)	.857 (6/7)	.538 (7/13)	.615 (8/13)
6 (DCD)	.264 (14/53)	.434 (23/53)	.750 (12/16)	.813 (13/16)	.800 (16/20)	.800 (16/20)
7 (DDC)	.429 (3/7)	.571 (4/7)	.333 (1/3)	.667 (2/3)	.893 (25/28)	.714 (20/28)
8 (DDD)	.058 (4/69)	.072 (5/69)	.048 (2/42)	.095 (4/42)	.012 (2/165)	.661 (109/165)

most likely to cooperate after a CCC combination and least likely to cooperate after a DDD combination; the lone exception is for receivers in Chicken, where best response and reciprocity pull in opposite directions. A few other patterns can be seen in the data, if we combine similar types of signal combination. We classify the 8 possible three-signal combinations into 3 classes: (1) *truth*, where the previous-round message and previous-round action were the same (combinations 1, 3, 6, and 8), (2) *nice lie*, a D previous-round message and a C previous-round action (combinations 5 and 7); and (3) *damned lie*, a C previous-round message and a D previous-round action (combinations 2 and 4). The strongest pattern is that following “truth”, the sender is likely to be truthful again: her current-round action is likely to be the same as her current-round message. This happens 76.6% of the time in the PD sessions, 97.1% of the time in the SH sessions, and 90.8% of the time in the CH sessions. Less strong, but still discernable, are the patterns following lies. Following a “nice lie”, senders tend to choose C; this happens 57.1% of the time in the PD sessions, 70.0% of the time in the SH sessions, and 78.0% of the time in the CH sessions. Following a “damned lie”, senders tend to choose D; they choose C only 10.1% of the time in the PD sessions, 7.7% of the time in the SH sessions, and 47.2% of the time in the CH sessions.

Receivers’ actions also correlate with these classes of signal combinations. In Prisoners’ Dilemma and Stag Hunt, receivers respond to “truth” by choosing an action the same as the sender’s current-round message; this happens 56.2% of the time in the PD sessions and 94.6% of the time in the SH sessions (but only 51.9% of the time in the CH sessions). Receivers in Prisoners’ Dilemma and Stag Hunt are substantially more likely to choose C following a “nice lie” than following a “damned lie”. In the PD sessions, they choose C 52.4% of the time following a “nice lie”, but only 17.3% of the time following a “damned lie”; in the SH sessions, the frequencies are 80.0% and 46.2%. In Chicken, this pattern does not hold; receivers choose C 68.3% of the time following a “nice lie” and 72.2% of the time following a “damned lie”.

These numbers, while suggestive, should be interpreted with caution due to small sample sizes in many cases. In order to draw more solid conclusions, we estimate another set of probit regressions (similar to those reported in

Table 5) using data from the WDL cells only. The binary dependent variable is again whether a player chose C, and we again estimate coefficients separately for senders and receivers and for each of the three games. In addition to the two-signal combination dummies that we considered in our earlier probit regressions, we now add four three-signal dummies, CCC, CCD, CDC, and CDD, which are equal to 1 if the (previous-round message, current-round message, previous-round action) combination is CCC, CCD, CDC, or CDD. Notice that three of these signal combinations correspond to the addition of a C previous-round message to the two-signal combinations CC, CD, and DD; including both the two- and three-signal combinations allows us to evaluate the marginal contribution of the extra piece of information provided in the WDL treatment.²⁰

We estimate coefficients for an unrestricted model, with dummies for both two- and three-signal combinations, and a restricted model with all of the three-signal-combination dummies equal to zero. Coefficient estimates,

Table 7: Probit results—frequency of cooperation (WDL cells, rounds 3–10)

	Prisoners' Dilemma Senders ($N=480$)		Stag Hunt Senders ($N=480$)		Chicken Senders ($N=480$)		Prisoners' Dilemma Receivers ($N=480$)		Stag Hunt Receivers ($N=480$)		Chicken Receivers ($N=480$)	
const.	-4.161** (1.863)	-3.537* (2.052)	-0.564 (0.868)	-0.783 (0.733)	1.062* (0.641)	0.383 (0.456)	2.412 (2.685)	3.011 (2.591)	0.804 (1.064)	0.644 (0.794)	0.369 (0.455)	0.197 (0.402)
round	-0.122*** (0.045)	-0.124*** (0.047)	-0.040 (0.064)	-0.061 (0.064)	-0.060 (0.044)	-0.056 (0.039)	-0.077 (0.052)	-0.068 (0.051)	-0.051 (0.068)	-0.042 (0.059)	0.000 (0.031)	-0.001 (0.031)
shfirst	0.819** (0.387)	0.629 (0.431)	0.872* (0.499)	0.946 (0.603)	-0.028 (0.432)	-0.046 (0.367)	-0.841 (0.743)	-0.802 (0.717)	-1.401 (1.263)	-1.287 (0.991)	0.358 (0.353)	0.338 (0.344)
chfirst	0.662** (0.273)	0.539* (0.307)	0.775* (0.460)	0.782 (0.520)	-0.012 (0.392)	-0.227 (0.343)	-0.794 (0.627)	-0.754 (0.568)	-0.625 (1.443)	-0.585 (0.842)	-0.209 (0.325)	-0.194 (0.325)
paydiff	-0.206** (0.087)	-0.155 (0.099)	0.024 (0.019)	0.038 (0.026)	-0.027 (0.082)	-0.040 (0.070)	0.140 (0.126)	0.153 (0.123)	0.083 (0.057)	0.086** (0.040)	-0.154** (0.066)	-0.136** (0.066)
CC	0.521 (0.604)	1.642*** (0.320)	1.497 (1.006)	3.295*** (0.738)	-1.022 (0.642)	1.190*** (0.328)	1.142* (0.634)	0.843** (0.339)	2.400*** (0.780)	2.881*** (0.634)	0.002 (0.358)	-0.080 (0.238)
CD	-0.460 (0.527)	-0.244 (0.400)	1.365 (0.839)	1.194 (0.850)	0.617 (0.701)	0.819* (0.433)	0.841 (0.604)	0.093 (0.337)	-0.001 (0.848)	0.435 (0.575)	0.154 (0.443)	0.361 (0.306)
DD	-1.501*** (0.549)	-0.606 (0.373)	-1.500* (0.866)	-0.758 (0.781)	-3.272*** (0.586)	-1.775*** (0.381)	-0.632 (0.640)	-0.681* (0.361)	-1.526** (0.711)	-1.377** (0.673)	-0.093 (0.332)	0.033 (0.237)
CCC	0.725* (0.376)	—	1.361* (0.714)	—	1.819*** (0.536)	—	0.171 (0.324)	—	0.636 (0.452)	—	-0.491* (0.255)	—
CCD	-0.772*** (0.273)	—	-2.019*** (0.621)	—	-0.516 (0.497)	—	-0.523* (0.285)	—	0.480 (0.769)	—	-0.004 (0.442)	—
CDC	-0.529 (0.580)	—	0.188 (0.858)	—	-1.218** (0.546)	—	0.596 (0.664)	—	†	—	-0.415 (0.421)	—
CDD	0.698** (0.345)	—	1.172 (0.73)	—	2.810*** (0.653)	—	0.557 (0.426)	—	†	—	-0.190 (0.246)	—
$-\ln(L)$	186.33	194.46	46.41	55.17	135.97	158.05	221.83	224.97	91.94	93.15	259.55	262.47
p -value‡	0.001		< 0.001		< 0.001		0.179		0.299		0.211	

* (**, ***): Coefficient is significantly different from zero at the 10% (5%, 1%) level.

†: This variable had to be omitted due to perfect correlation with $P(C)$.

‡: Likelihood-ratio test of no difference between unrestricted and restricted models.

Note: Signal combinations are ([previous-round message,] current-round message, previous-round action).

standard errors, log-likelihoods, and p -values (from likelihood-ratio tests of the joint significance of the three-

²⁰We use only four of the eight possible three-signal combinations—those with a C previous-round message—to avoid perfect collinearity. (For example, any CD combination in rounds 3–10 of the WDL treatment would have been either CCD or DCD.)

signal–combination dummies) are shown in Table 7. We have again corrected the standard errors for clustering. We see that the coefficient for the round number is almost never significant; the only exception is for senders in Prisoners’ Dilemma, where it is again negative. Coefficients for the game–order dummies (shfirst, chfirst), and expected payoff differences (paydiff) are significant in only a few cases.

The most important finding in Table 7 is that the coefficients for the three–signal dummies are always jointly significant for senders; that is, the previous–round message is useful for predicting senders’ actions. In fact, the signs of these coefficients are consistent with our discussion of Table 6. A CCD combination corresponds to a “bad lie”, so that the sender is likely to choose D, while a DCD combination signals “truth”, so that the sender is likely to be truthful again (choose C). This reasoning implies that the sender is less likely to choose C after a CCD combination than after a DCD combination; that is, the coefficient on the CCD dummy should be negative. Indeed, this is the case (CCD is significantly negative in two of the three games). Similar reasoning predicts that the sender should be less likely to choose C after a CDC combination (“truth” and a D current–round message) than after a DDC combination (“nice lie”), and if we use the fact that the pattern following truth is stronger than that following lies, the sender should be more likely to choose C after a CCC combination (“truth” and a C message) than after a DCC combination (“nice lie”), and more likely to choose C after a CDD combination (“bad lie”) than after a DDD combination (“truth” and a D message). Each of these sign predictions is borne out in the corresponding coefficient estimates, in all cases where the estimate is significantly different from zero.

Surprisingly, we are unable to conclude that receivers base their actions on senders’ previous–round messages; the coefficients for the three–signal dummies in the receiver regressions are not jointly significant at conventional levels of significance. One explanation for this finding may be that senders’ previous–round messages are highly correlated with their current–round messages and, often, consistent with their previous–round actions (they frequently tell the truth). Consequently, while senders’ previous–round messages do provide additional information about their subsequent actions, the value of this information is relatively small, and may be overwhelmed by the cognitive costs of processing it. This suggests that the common knowledge in our experiment that lie detection is possible—via the observation of previous–round messages—is sufficient to constrain the behavior of senders to the point that receivers need not be so careful about checking for lies.²¹

6 Summary

How do individuals achieve good outcomes in strategic situations? A common explanation is that they make use of additional information that is available to them. Cheap talk and observation of past actions represent two types

²¹A corollary of this hypothesis is that when such a constraint is removed—as in the final round of a game—senders will become less truthful. Such behavior by senders can be interpreted as their exploiting reputations for truthfulness built up in earlier rounds. To examine this possibility, we estimated additional probit regressions for senders in Prisoners’ Dilemma (where we would expect this tendency to be strongest), similar to the ones presented in this section and in Section 5.3, with an additional dummy variable for the final round. We found that the coefficient for the final–round dummy was not significant when we used pooled WD and WDL data (as in Section 5.3), but it was significant and negative when we looked at only the WDL data (as in this section). This fits our hypothesis, as building a reputation for truthfulness is possible in the WDL treatment but not the WD treatment. As an illustration, if we concentrate on sender behavior following CCC signal combinations—which is the way we would expect such reputation–building to happen—we find that senders choose C 84.9% of the time in rounds 3–9, and 70.0% of the time in round 10. This is a substantial decrease, though full exploitation would imply senders never cooperate in round 10.

of information that players might use. Whether they use these signals, and the extent to which the signals affect behavior, are clearly empirical questions. Our experiment was designed to address those questions by having subjects send or receive both types of signal. We found that they did indeed make use of them. Signals are typically truthful and believed, even when they ought to be false (or at best uninformative) and disregarded. How truthful the signals are depends on the game’s strategic structure and whether the receiver is given another piece of information: the sender’s previous-round message, which—in conjunction with the previous-round action—allows the receiver to judge the veracity of the sender’s current-round message. As could be expected, current-round messages are more truthful (better correlated with the sender’s subsequent action) when the game provides less incentive to lie and when receivers have access to senders’ previous-round messages. However, observed actions also become better correlated with current-round actions in these cases, so much so that there is no uniform relationship between messages and observed actions as predictors of senders’ actions. Receivers, for their part, respond to signals in systematic ways that seem to combine payoff maximization with reciprocity. When these effects work in the same direction, their common prediction is consistent with receivers’ behavior. When they work in opposite directions, they cancel each other out and receivers’ behavior shows little overall pattern.

Our experimental design also allows us to compare the aggregate outcomes reached in these treatments with those from earlier treatments from our previous paper, in which receivers received either no signals or one. We found that the multiple-signal treatments led to higher levels of cooperation, more frequent coordination on pure-strategy Nash equilibria (when multiple equilibria are present), and higher average payoffs than the earlier control (no-signal) treatment. However, they typically led to lower levels of cooperation, less frequent coordination, and lower average payoffs than the earlier cheap-talk-only and observation-only (one-signal) treatments. That is, once one signaling device is present, adding an additional one lowers the likelihood of good outcomes. The potential for worse outcomes with more signaling opportunities may at first seem paradoxical. However, increased opportunities for signaling lead to the possibility of “crossed” signal combinations, whose elements have conflicting interpretations, so that the receiver might have difficulty inferring anything at all about the sender’s likely choice of action. Examination of our results, broken down into play following crossed signal combinations versus play following “aligned” signal combinations (whose elements have identical interpretations), shows that crossed signal combinations tend to lead to outcomes no better than those in the no-signal treatment, while aligned signal combinations lead to outcomes at least as good as, and often better than, those in the one-signal treatments. Levels of coordination and efficiency tend to be higher after aligned signal combinations than crossed signal combinations, both in an absolute sense and relative to what would be expected from the observed relative frequencies of cooperation in each case. Aligned pairs of current-round message and observed previous-round action are also particularly truthful, even more so than either current-round messages or observed actions alone.

The primary implication of our findings is that it is not merely the *amount* of information that determines the likely outcome of a situation, but rather the *content* of that information. Increases in the dimensions along which individuals can signal may lead to quite complicated behavior with ambiguous welfare consequences. On the one hand, the opportunities provided by additional signaling (for example, lie detection) can aid in the establishment of the trustworthiness of an opponent. Indeed, the fear of being “caught in a lie” may constrain the behavior of signal senders to the point that receivers become less concerned about lie detection.²² On the other hand, increased

²²Seller feedback mechanisms in e-commerce that have the flavor of our WDL treatment in that buyers can see how sellers have performed in past transactions with other buyers, have been credited with increasing buyer participation and satisfaction. Resnick and

signaling opportunities can lead to the possibility of crossed signals, or even possibly of intentional, strategic “signal jamming.” Casual empiricism suggests that individuals involved in strategic encounters are frequently bombarded by multiple signals of various types, so a reasonable direction for future theoretical and experimental research on signaling games is to consider how individuals react to a variety of different signals. This paper represents a first, small step in this direction.

References

- 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.
- Burton, A., G. Loomes, and M. Sefton (1999), “Communication and Efficiency in Coordination Game Experiments” CEDEX Working Paper.
- Blume, A. and A. Ortmann (1999), “The Effects of Costless Pre-Play Communication: Experimental Evidence from a Game with Pareto-Ranked Equilibria, working paper, University of Pittsburgh.
- Bosch-Doménech, A. and N.J. Vriend (2003), “Imitation of Successful Behaviour in Cournot Markets,” *Economic Journal* 113, pp. 495–524.
- Çelen, B., S. Kariv, and A. Schotter (2003), “The Advice Puzzle: An Experimental Study of Social Learning Where Words Speak Louder Than Actions,” working paper, New York University.
- 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. 739–771.
- Charness, G. (2000), “Self-Serving Cheap Talk: A Test of Aumann’s Conjecture,” *Games and Economic Behavior* 33, pp. 177–194.
- Crawford, V.P. (1998), “A Survey of Experiments on Communication via Cheap Talk,” *Journal of Economic Theory* 78, pp. 286–298.
- Duffy, J. and N. Feltovich (1999), “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 (2002), “Do Actions Speak Louder Than Words? Observation vs. Cheap Talk as Coordination Devices,” *Games and Economic Behavior* 39, pp. 1–27.
- Earl, P.E. (1990), “Economics and Psychology: A Survey,” *Economic Journal* 100, pp. 718–755.

Zeckhauser (2001), for instance, report that on eBay, 99.1% of all buyer feedback about sellers is positive.

- 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. 143–158.
- Farrell, J. and M. Rabin (1996), "Cheap Talk," *Journal of Economic Perspectives* 10, pp. 103–118.
- Feltovich, N. (2003), "Nonparametric Tests of Differences in Medians: Comparison of the Wilcoxon–Mann–Whitney and Robust Rank–Order Tests," *Experimental Economics* 6, pp. 273–297.
- Huck, S., H.T. Normann, and J. Oechssler, (1999), "Learning in Cournot Oligopoly—An Experiment," *Economic Journal* 109, pp. 80–95.
- Huck, S., H.T. Normann, and J. Oechssler, (2000), "Does Information About Competitors' Actions Increase or Decrease Competition in Experimental Oligopoly Markets?," *International Journal of Industrial Organization* 18, pp. 39–57.
- 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.
- Okuno-Fujiwara, M. and A. Postlewaite (1995), "Social Norms and Random Matching Games," *Games and Economic Behavior* 9, pp. 79–109.
- Resnick, P. and R. Zeckhauser (2001), Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System, working paper, January 2001.
- Roth, A.E. and M.W.K. Malouf (1979), "Game–Theoretic Models and the Role of Bargaining," *Psychological Review* 86, pp. 574–594.
- Siegel, S. and N.J. Castellan, Jr. (1988), *Nonparametric Statistics for the Behavioral Sciences*, McGraw–Hill, New York.
- Simonsohn, U., N. Karlsson, G. Loewenstein, and D. Ariely (2004), *The Tree of Experience in the Forest of Information: Overweighing Personal Over Vicarious Experience*, working paper.
- Sugden, R. (1995), "A Theory of Focal Points," *Economic Journal* 105, pp. 533–550.
- Wilson, R. and J. Sell (1997), " 'Liar, Liar...' Cheap Talk and Reputation in Repeated Public Goods Settings," *Journal of Conflict Resolution* 41, pp. 695–717.