

Sopher, Barry; Zapater, Inigo

Working Paper

Communication and Coordination in Signalling Games: An Experimental Study

Working Paper, No. 1998-03

Provided in Cooperation with:

Department of Economics, Rutgers University

Suggested Citation: Sopher, Barry; Zapater, Inigo (1998) : Communication and Coordination in Signalling Games: An Experimental Study, Working Paper, No. 1998-03, Rutgers University, Department of Economics, New Brunswick, NJ

This Version is available at:

<https://hdl.handle.net/10419/94292>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Communication and Coordination in Signaling Games: An Experimental Study

by

Barry Sopher* and Iñigo Zapater**

October 2000

*Rutgers University, New Brunswick, NJ and **Deloitte and Touche, Chicago, IL
The research reported on here was completed while the authors were visiting New York University Department of Economics. The support of the C.V. Starr Center for Applied Economics at New York University and the Rutgers Research Council is gratefully acknowledged.

Proposed running title: "Communication in Signaling Games"

Send correspondence to:

Barry Sopher
Department of Economics
Rutgers University
New Jersey Hall
New Brunswick, NJ 08901

Abstract

This paper reports on an experiment testing *Credible Message Rationalizability* (CMR) in one-shot and repeated signaling games with costless communication. CMR is a non-equilibrium theory combining rationalizability with the idea that players will, using the natural meaning of the language, transmit enough information to coordinate on a mutually beneficial outcome. Two conventions of behavior arise in the games, *Partial* and *Full Truth-telling*. In one-shot games, subjects manage to communicate roughly as predicted by CMR, which predicts that only the maximum amount of true information consistent with rationality will be transmitted. In repeated games, Full truth-telling is more prevalent. The repeated game gives the receiver power to impose a preferred convention. Full truth-telling is especially prevalent when there is a conflict between conventions.

1. Introduction

In recent years the study of costless communication, or "cheap talk", in games has attracted a lot of interest. Game theorists are interested in knowing the effect, if any, of cheap talk on the outcome of a game. That is, does allowing non-binding, non-verifiable statements lead to different outcomes than not allowing such statements. Indeed, one of the justifications sometimes given for the use of Nash equilibrium as a solution concept is that the players get together and engage in some form of non-binding communication that results in the play of strategies that form a Nash equilibrium. Thus, it is natural that the study of cheap talk has been conducted largely in an equilibrium framework.

The multiplicity of Nash equilibria in many games has spawned a large literature on refinements, but the implausibility of the computational requirements of Nash refinements leads to questions about the practical usefulness of such an approach. Banks, Camerer and Porter (1994) found some empirical support for refinements, but they concluded that those refinements that "worked best" in experiments were those that led to outcomes consistent with other possible stories: e.g., simple rules-of-thumb.

A natural theoretical response to such findings is to formally embed behavioral rules in the description of the game, and see what results. This is the basic idea of Rabin's (1990) *Credible Message Rationalizability* (CMR), which we test in the experiments reported on in this paper. Rabin argues that if we use Nash equilibrium as our solution concept, we are assuming a lot of coordination among the players, when what we want to analyze is precisely the role of communication in coordinating actions. He propose an alternative way of analyzing these games. The key features of CMR are rationalizability, augmented by the assumptions that players will (i) use the natural meaning of the language to (ii) transmit enough information to coordinate on a mutually beneficial outcome¹. How much information will be transmitted depends on the rationality of the players. In general, it is the maximum amount of truth-telling compatible with rationality.

This paper reports on an experiment that tests CMR, and a generalization of this theory by Zapater (1993), *Generalized CMR* (GCMR). We restricted our analysis to simple *signaling* games, in which there is a *sender* with a privately known random *type* and a *receiver* with a prior probability distribution for the sender's type. A sender sends a message containing a statement about his type, and the receiver then takes one of several available *actions*, with payoffs determined only by the sender's type and the receiver's action.

In the experiment, subjects were given a menu consisting of one message for each possible subset of the set of types. We considered one-shot and repeated versions of 4 different games. For some games, CMR predicts that each sender type will separate, truthfully revealing his private information, and that the receiver will believe this and respond appropriately. For other games, CMR predicts that at least some types will pool, not revealing their types. GCMR provides sharper predictions for some situations than CMR.

Although CMR is not an equilibrium theory, it does have common knowledge requirements, noted above: rationality and assumptions (i) and (ii). These assumptions combined generate a focal policy that says the sender will reveal the maximum amount of information compatible with rationality. There are other possible focal policies that players could follow. For example, "full truth-telling" is a very simple focal policy. If CMR is to work, then, players have to learn (jointly) that CMR's focal policy is, in some sense, the best policy to follow. Further, even if a player believes in the spirit of CMR, it may take some time to figure out all the implications of CMR, especially for a complicated game. Therefore, two types of learning may be required for CMR to work: the coordination problem of settling on a focal policy, and the computational problem of deducing the implications of a policy.

We allowed subjects to play (and learn) in two environments. Some subjects played a finitely repeated game against a series of randomly selected opponents, while other subjects played a finitely repeated game against the same opponent. The results of the experiment, briefly, indicate

generally strong support for CMR, but the degree of support depends on the environment. The subjects seem to understand that full truth-telling needs to be refined, and in the aggregate they manage to establish some, but not all, of the implications of CMR.

We organize our results around three hypotheses that differ in the amount of truthful communication transmitted:

1. *No communication*: the receiver disregards any message sent by the sender, and always responds with the action that maximizes his expected utility given the initial beliefs.
2. *Full Truth-Telling*: the sender always sends a message that fully reveals his type.
3. *Partial Truth-Telling*: the sender transmits some truthful information, although he does not always reveal his type. CMR and GCMR fall within this category.

Our basic result is that, in the games studied in this experiment, the players understand that it is possible to transmit some truthful information, thus costless communication can help to coordinate the actions of players. They also understand that Full Truth-Telling can not be a consistent policy to follow because it leads to opportunistic behavior. Although our experimental results give support to hypothesis 3, they do not fully support the predictions of CMR and GCMR. The subjects act as if they were unable to discover the full implications of rationality. For example, in a game with two types, even if the players know that type 1 will always send the message "I am type 1," they act as if they failed to understand that any other message must come from type 2.

These conclusions must be qualified somewhat for some of the repeated games, where behavior sometimes deviated markedly from the one-shot games. Specifically, we find that, even with its inconsistencies, the Full truth-telling convention is much more prevalent in the repeated games. Although one of the reasons for looking at repeated games in our experiment is to facilitate learning, the situation represented by the repeated games is quite different from the one represented by the one-shot games. First, in the one shot games, the receiver has very little power in establishing a convention, while in the repeated game the receiver can use his actions to punish deviations from

truth-telling, which is his preferred convention. Second, the repeated version brings Folk Theorem type of results. This means that cooperative behavior, which is not possible in the one-shot games, can develop. For instance, full truth-telling can arise as a convention if the receiver rewards sufficiently the truth-telling behavior of the sender. Also, because of the multiplicity of equilibria in the repeated games, there exists a conflict of conventions that does not exist in the one-shot games. In this regard, the truth-telling convention could prevail because of its simplicity.

The rest of the paper is organized as follows. In Section 2, equilibrium and non-equilibrium theories of communication in signaling games are outlined and illustrated using as examples the games we study in the experiments. Section 3 contains details of the experimental design and the procedures followed in the experiments. Section 4 contains a more detailed account of the experimental results. Section 5 contains a discussion of related theoretical and empirical work on communication and coordination in games. An appendix containing the instructions given to subjects in the experiments is available upon request.

2. Credible Message Rationalizability and Other Theories of Communication

In this paper we present the results of an experiment in which subjects played Games 0, 1, 2 and 3, illustrated in Figure 1. Our main purpose is to test the predictions of CMR and GCMR. In this section we discuss both equilibrium and non-equilibrium theories of communication, using the games we analyze in the experiments to illustrate concepts. Detailed predictions for each game are presented in Section 4. This section is only for illustration of some of the main concepts, and is not intended as a thorough statement of predictions.

Figure 1 about here

A simple communication game consists of two players: An informed player called the

sender, S , who is an element of a finite set of types $T = \{t_1, \dots, t_N\}$, and an uninformed player called the receiver, R . The players have prior beliefs represented by the probability distribution p over T . S sends a message m , which is an element of the finite set of messages M , and R , upon hearing the message responds with some action a that belongs to the finite set of actions A . In contrast with costly signaling games, the utility of the players do not depend directly on the message used by S .

So for example, in game 0, illustrated in Figure 1, the set of types is $T = \{1, 2\}$, the prior belief for type 1 is $p(1) = p(2) = 1/2$, and the set of actions is $A = \{A, B, C\}$. We have restricted the set of possible messages in Game 0 (and all other games) to a minimal set of messages that allow the sender to communicate any partition of the set of possible types, or to remain silent. Thus, in Game 0, the messages are:

Message 1: "I am type 1"

Message 2: "I am type 2"

Message 3: "I am type 1 or type 2"

Message 4: "No message"

Notice that the receiver always wants as much information as possible. On the other hand, it is not always in the sender's interest to reveal his identity. The analysis of these games focuses on what and how much truthful information will be transmitted.

We can divide the solutions proposed for these kind of games into two groups: equilibrium and non-equilibrium solutions. The former look at the set of Nash equilibria of the communication game and use some intuitive arguments to refine away some implausible equilibria. The latter use rationalizability combined with some focal policy to determine the outcome of the game.

The simplest equilibrium selection criterion is Pareto dominance. This has been mentioned in Crawford and Sobel (1982). See also Van Huyck, Battalio and Beil (1990, 1993). If one equilibrium is better than any other equilibria for all types of sender (and the receiver), then it is likely that that equilibrium will result. For instance, in Game 0 there exist pooling equilibria where

R responds to any message with action A, and separating equilibria, where each type of sender and the receiver get their highest possible payoffs. In this case the equilibria can be Pareto-ranked, and all types of sender and the receiver prefer separation.

Neologism-proofness (Farrell (1985)) and a generalization of that concept, Announcement-proofness (A-P) (Matthews, Okuno-Fujiwara, and Postelwaite (1991)) considers each equilibrium of the game and asks what happens if the sender tries to introduce a new word into the vocabulary of the game, in an attempt to convince the receiver to play in a different way from the one prescribed by the equilibrium. That is, how should the receiver interpret such a message? The theory of neologism-proofness proposes that the receiver should reason in the following way: "The unexpected message must come from a type, or set of types, that expect to get a better payoff than in equilibrium." The idea, then, is that the receiver considers what each subset of the possible types, K , have to gain from sending such a message. If the receiver updates his beliefs, conditioning on K , and his resulting best response generates payoffs for all types in the set K that are greater than or equal to their payoffs in the original equilibrium, then the original equilibrium is not neologism-proof. For example, in Game 0, the pooling equilibrium is not neologism-proof. In the pooling equilibrium, type 1 gets a payoff of 2, but he can announce that he is in fact type 1. If the receiver believes him, she will respond with action A, and that will give type 1 a payoff of 8. Notice that type 2 will not send that message because he will get a lower payoff than in equilibrium, hence, the receiver has the correct beliefs. Announcement-proofness allows for more than one message to be sent at the same time. For example, in Game 0, S could announce that if he is type 1 he will send message m , and that if he is type 2 he will send message m' . This is an announcement and it will also break the pooling equilibrium.

Rabin (1990) has criticized the above approach as follows. The reason for analyzing cheap-talk in games is to understand the power of costless communication in coordinating the actions of different players. If we begin by looking at the Nash equilibria of a game, we are already

assuming a lot of coordination. For this reason, he uses rationalizability as his basic solution concept. Rationalizability is a weaker solution concept than Nash equilibrium, and on its own it does not have any predictive power in these kind of a games. It only assumes that the rationality of the players is common knowledge, but it does not say anything about what a message should mean. This implies that any message could come from any type and, in particular, that the set of beliefs that R could have after hearing any message should be $\mathcal{J}(T)$, the set of all probability distributions over the set T . To enrich the environment he assumes that all players share a common trait that he calls "honesty". This behavioral rule says that the players want to use their language to transmit some information because that could be mutually beneficial. The uninformed player, R , always prefers more information to less. On the other hand, S can use the language to induce a profitable response. Honesty should thus be understood as a refinement of full truth-telling, where S transmits the maximum amount of information that is compatible with the rationality of the players. Next, we present an alternative, but equivalent, interpretation of Rabin's (1990) formalization of this notion of honesty.

We assume that the set M is large enough so that the sender can communicate any relevant information. In particular, consider the set 2^T of all the subsets of T , and assign to each element T_i a message m_i in the set M . To understand how much truth-telling is compatible with rationality, we start by assuming that R will consider every message m_i to be true. Then for every nonempty T_i , we can calculate the set of best responses to the message m_i , $BR(m_i)$, by conditioning the initial beliefs on T_i . Corresponding to the element $\emptyset$ of 2^T there exists a noninformative message $m_\emptyset$. The literal meaning of the message $m_\emptyset$ is "I will not reveal my type". In this case, because the message does not contain any information, we are back in the world of rationalizability. Therefore, the set of best responses to $m_\emptyset$ coincides with the set of undominated actions in the set A . Assuming that a message is true allows us to associate to each message in the set M the subset of actions $BR(m)$. Given the sets $BR(m)$ we can define preferences of type over messages. We say that type t prefers m_i to m_j if

by sending message m_i he always gets at least the maximum payoff that he could get with m_j , and sometimes more.

Although Rabin (1990) never talks explicitly about preferences we can reinterpret his credibility conditions by saying that for a message m_i to be credible it is necessary that m_i is preferred to any other message in the set M , for all $t \in T_i$. If some type t prefers m_i to $m_\emptyset$, then it must be the case that m_i gives type t his highest possible payoff. Also, if m_i is credible, every type in T_i will use the message m_i with probability one.

In Game 2, although it appears that type 3 should be able to separate from types 1 and 2, Rabin's theory does not make that prediction because types 1 and 2 cannot send any truthful message that with certainty will give them their highest possible payoff. Thus, the message "I am type 3" is as good as any other message for type 1 or type 2. In particular, a type 1 or type 2 might hope to avoid the receiver's safe action, D, by saying "I am type 3." To formalize this intuition behind this kind of game, Zapater (1993) extends Rabin's approach by introducing, in addition to the kind of messages just considered, a weaker kind of message. Given a subset T_i of T , the weak message wm_i associated with t_i generates the set of beliefs $J(T_i)$. Zapater (1993) relaxes Rabin's credibility condition by just requiring that for a weak message to be credible, it must be true that for every type t in the associated set T_i , the message wm_i is not dominated by any other message. Notice that under these circumstances a type $t \in T_i$ need not send the message wm_i with probability one. This justifies that the meaning of the weak message wm_i should be $J(T_i)$. If we consider Game 2 again we see that if t_1 or t_2 use a weak message that means $J(\{t_1, t_2\})$, then they will never send the message "I am type t_3 ," and type t_3 could reveal his identity.

3. Experimental Design and Procedures

Design

There are four different parameter sets, corresponding to the four games considered in

Section 2. We were primarily interested in the performance of CMR and GCMR in each game. We were also interested in the relative performance between the one-shot (with players matched against different opponents in each period) and repeated games (with players always matched with the same opponent) for a given parameter set. The main reason for studying these two types of games, each with "repeated" play, is to allow players to learn more sophisticated behavior than full-truth-telling or silence. CMR is really a theory of play in the one-shot game and when the same game is repeated with the same players, some features unique to a repeated game may arise. The one-shot predictions are, nonetheless, one possible outcome for the repeated game, so it is of interest to see whether the one-shot predictions are robust to this extension of the game.

In five of six experimental groups, subjects were senders in some periods and receivers in other periods, for sets of 5 to 10 periods at a time. In one experimental group (#6), players were randomly assigned one role (sender or receiver) for the entire experiment, and played from 5 to 25 periods at a time.

Each experiment consisted of 2 sessions. In Session 1 of each experiment, Game 0 was played. Session 1 consisted of 10 periods of play, with each subject playing as a sender for 5 consecutive periods and a receiver for the other 5 periods. The details of pairings, etc., are contained in the appendix. Group 6 played this game for only 5 periods. We considered this game to be mainly a warm-up exercise for subjects.

Three parameter sets were used in Session 2, corresponding to Games 1, 2 and 3 above. Session 2 consisted of 20 periods of play, with each subject playing as a sender for 10 consecutive periods and as a receiver for the other 10 periods. Group 6 played a modified version of Game 1 (Game 1') in which the entry (10,4) in row 2, column C, was changed to (15,4). The game played by Group 6 also differs from the games played by other groups in that Session 2 consisted of 25 consecutive periods of play, instead of 10.

There were one-shot and repeated versions of Games 1 and 3, plus a longer repeated version

using the modified Game 1'. There was a one-shot version of Game 2, but not a repeated version. The predictions for each parameter set are summarized in Table 1. The schedule of experiments, including date conducted, parameter set, whether one-shot or repeated, etc. is contained in Table 2.

Tables 1 & 2 about here

Procedures

All of the experiments were conducted at New York University in April of 1992. The subjects were undergraduate students from economics classes. Each experimental group was recruited in classes within one week of the scheduled time of the experiment. At the time they were recruited, subjects were told that they would be participating in an experiment concerned with "the economics of decision-making," that the experiment was expected to last about 2 hours, and that average earnings would be about \$25. Only one experiment lasted longer than 2 hours (by 15 minutes), and average earnings turned out to be slightly more than \$26 per subject.

The experiments were conducted in a conference room equipped with cardboard dividers to give privacy to each subject. In each experiment there were 8, 10 or 12 subjects, with half seated on each side of the table. Instructions were passed out, read aloud to subjects, and then questions answered. The instruction period typically took about 20 minutes. After instructions were read and questions answered, Session 1 began. Occasionally some additional questions arose in the early periods of Session 1, that required some additional answering of individual questions. We do not believe there was any significant confusion as to what the procedures were and how earnings were determined after the first period or two of Session 1. Full details of the step-by-step conduct of the experiment are contained in an appendix, available on request from the authors.

4. Experimental Results

In this section, discussion of the predictions of CMR and GCMR for the games, and the results of our experiment, are discussed. We have not conducted formal statistical tests on the data, since our sample sizes are relatively small, and the structure of the games do not lend themselves to straightforward or enlightening hypothesis testing. Instead, we present aggregate cross-tabulations of two sorts for each game--sender type against message sent, and message sent against action taken--and base most of the discussion on these cross-tabulations. We also discuss some of the individual pair data.

4.1 Game 0 Results (Session 1)

We first consider the results of Session 1 (Game 0), shown in Figure 1. The purpose of Game 0 was mainly to train participants in the mechanics of the experiment (sending messages, taking actions, recording information). Notice that in this game the CMR prediction coincides with full truth-telling.

In Game 0 there are pooling equilibria where no information is transmitted, and separating equilibria where the sender always reveals his type.

There are several equilibria of each type, many of which do not use the "natural meaning" of the language. For example, one separating equilibrium is for type 1 to send message 2 ("I am type 2"), for the receiver to respond with action A, and for type 2 to send message 1 ("I am type 1"), and for the receiver to respond with Action B. Restricting messages to their natural meanings, there is one separating equilibrium, and two pooling equilibria (only message 3 or 4 sent).

Action C is a safe strategy for the receiver, so if risk or regret looms large in the receiver's mind, one might expect to see action C chosen all the time-- a pooling equilibrium in which all messages have the same meaning--i.e., nothing. CMR and GCMR both predict separation in this game.

Sender Behavior

Table 3 shows a cross-tabulation of sender type against the messages chosen by the senders. The table shows this cross-tabulation separately for those groups that played one-shot games (Groups 1-3) and those that played repeated games (Groups 4-6). In Game 0 there were 2 possible sender types and (thus) 4 possible messages (shown at the bottom of the table). In Game 0, each sender type has a credible message, which is to state its true type. For both one-shot and repeated games, this is what most of the senders do: about 75% in the one-shot, 79% in the repeated games state their true types. The remaining senders either lie (and state that they are the "other" type, or they do not identify themselves (say they are "type 1 or type 2", or say nothing). There appears to be no difference between the cross-tabulations in the one-shot vs. the repeated games.

Table 3 about here

Receiver Behavior

The bottom of Table 3 shows a cross-tabulation of the message received and action taken by the receiver, again separately for the one-shot and repeated games. Of the set of messages that purport to identify a sender (i.e., Message 1 and Message 2) that were sent, about 80% (in one-shot games) and 77% (in repeated games) were responded to as if they were believed (assuming subjects were best-responding). Again, as with the senders' message-sending behavior, there appears to be no big differences between the one-shot and repeated games with respect to action-taking behavior by receivers. The one possible exception to this is that receivers show a slightly stronger tendency to take action 'C' in response to all messages in the repeated games. In the game against different opponents a message that is not full-truth telling could be due to the shortsightedness of one player and would not necessarily carry over to the next match, while in the repeated game with the same opponent, the same kind of message might make the receiver doubt the rationality of the sender.

Action C could thus be either a punishment for not being fully truthful, or it could be a safe response to an irrational sender.

Discussion

Although the aggregated results are consistent with separation and truth-telling, there are plenty of exceptions. For example, some senders say they are type 2 when in fact they are type 1². This could be due to learning (i.e., just mistakes). On the other hand, these exceptions could be significant in that a subject may, besides learning the rules of the experiment, gather information about the rationality of the other subjects in the experiment. Further, one might learn that what one considers focal in the game is not necessarily focal for the other subjects. This is important because CMR assumes that, from the structure of the payoff matrix, players can infer everything they need in order to play the game. That is, by examining the payoff matrix, people can infer that certain types should play one way, and the behavior of other players in the game can then be deduced from rationality. Even if some people understand the reasoning of CMR, they may not play as predicted by CMR if they believe the players they will face are not rational with probability 1. A comparison to behavior in ultimatum bargaining games seems appropriate: Even if one understands that a rational person should prefer "something to nothing" and accept any offer greater than zero in an ultimatum bargaining game, it does not follow that making the smallest possible offer is the best policy, when you do not believe the other player is rational with probability 1.

4.2 Game 1 Results (Session 2)

In Game 1, the set of equilibria range from separating, where the sender always reveals his type, to pooling equilibria where no information is transmitted. There are also some semipooling equilibria. In this game, truth-telling does not coincide with the CMR prediction. Thus, Game 1 requires a higher degree of sophistication in interpreting messages. The basic CMR prediction for

this game is as follows: A type 1 sender sends message 1, while type 2 sends any message except message 1. The receiver responds to message 1 with action A, and responds to any other message with action B. Notice that this implies that the receiver must understand that although a sender may "hide" his identity, message 1 will only come from type 1. That is, any message other than message 1 comes from type 2. Truth-telling is like CMR, except that type 2 always sends message 2, rather than "anything but message 1". Since truth-telling requires type 2 to directly reveal his identity, a receiver who believes that truth-telling is focal and receives message 3 or 4 will need a new interpretation of the situation. This may include questioning whether type 1 is really behind message 1.

As in Game 0, a risk-averse receiver might want to respond to every message with action C. Game 1 represents the first challenge to CMR because, in contrast to Game 0, the preferences of all players are not perfectly aligned. While the receiver always prefers more information, the sender, *ex ante* (before he knows his type), prefers the pooling outcome. *Ex post*, type 1 prefers to separate, while type 2 prefers to pool.

Sender Behavior

Table 4a shows cross-tabulations of sender type by message sent in the one-shot and repeated versions of Game 1. Note that there are two versions of the repeated game: Group 4 played the standard version, in which a given pair plays for 10 periods, while Group 6 played a longer version, in which a given pair plays for 25 consecutive periods. There appears to be strong support for CMR in both the one-shot and repeated versions of the game, for sender behavior. For example, between 86% and 92% of type 1 senders send message 1, "I am type 1", and manage to separate themselves from type 2 senders, as predicted by the theory. CMR predicts nothing specific for type 2 senders, except that they should not send message 1. In both the one-shot and repeated versions of the game, very few type 2 senders send message 1, consistent with CMR. But there is a slight difference

between the one-shot and repeated games, as follows: type 2 senders appear slightly more inclined to send message 2, identifying themselves as type 2, in the repeated game than in the one-shot game.

In the repeated version of Game 1, a sender could use "no message" (message 4), regardless of his type, in an attempt to induce action C, the one that maximizes his *ex ante* expected utility. Notice, though, that this strategy does not satisfy backwards induction, by a familiar argument.³ Alternatively, the sender could stick to sending message 1 early in the game, even when he is not of type 1, in an attempt to establish that he does not intend to communicate any information about his type, and thus try to induce action C in this way. Notice that this will tend to ruin the meaning of message 1, so perhaps this would be more credible than sending message 4. On the other hand, the sender would have to be "lucky," and draw type 2 early in the game, if this method were to have any hope of working.

The purpose of the long repeated game (Group 6) was to challenge the theory a bit more. Specifically, it seems that a sophisticated sender could "pool" his two types in the repeated game, always sending the same message, and maybe get the receiver to take action C, giving the sender a higher expected payoff. The first thing we did was to change the (10,4) entry to (15,4), as an additional allure to the sender. We also extended the game, from 10 to 25 periods, to give time for the pooling convention, if it is to arise, to be established. In fact, we were not able to upset the predicted CMR outcome: The frequencies of sender behavior in the long repeated game are virtually identical to the regular repeated game. Thus, comments on the repeated version of Game 1 below refer to both the regular and the extended version, unless otherwise noted.

Tables 4a & 4b about here

Receiver Behavior

Table 4b shows cross-tabulations of message received by action taken for the one-shot and

repeated versions of Game 1. In all versions of the game, message 1 is responded to as if it were true a majority of the time: 75% for the one-shot version and 87-88% of the time in the repeated versions. According to CMR, any other message than 1 ought to be taken to mean the sender is type 2. In all versions of the game, message 2 is interpreted in this way (76% in one-shot, 82-90% in repeated).

One-Shot vs. Repeated Games

The most notable difference between the one-shot and repeated versions of Game 1 is that the receiver in the repeated game is much less likely to choose action C in response to a message other than 1 or 2. But responses to messages 3 and 4, which do not explicitly identify the sender, are very different across games. Only 36% of responses to message 3 or 4 in the one-shot game are action B, while from 76 to 84% in the repeated versions are action B. So, repeating the game seems to help the CMR predictions, in Game 1. One possible explanation is that it is easier for the receiver to tell that message 3 or 4 is sent by a type 2, presumably because he has noticed that message 1 is sent in those periods when (as he soon learns) the sender is type 1. Another possibility is that the receiver simply has more power in the repeated game, in the sense that he can impose the separating outcome. That is, the receiver may simply not choose action C, not allowing the sender to benefit from pooling his types. We do not observe any systematic evidence that receivers are "punishing" in this way, but this, of course, does not preclude the possibility that senders did not try to pool for fear of such punishment.

Discussion

Overall, then, only one part of the CMR prediction really works. Type 1 sends message 1, and the receiver responds with action A, but the receiver fails to follow the CMR logic all the way through. In fact, analysis of the individual data reveals that almost all of the subjects, when playing

the role of a receiver, choose both actions B and C in response to messages 3 and 4. This suggests that almost all of the receivers are not entirely sure of how to interpret such messages, mixing between the optimal CMR response of action B and the safe response of action C. Things do work a bit better in the repeated game, however. Both frameworks involve repetition, and hence the possibility of learning. The differences may be that with different opponents, it is not clear how much the experience with one subject should affect behavior with the next opponent. When one is playing against the same opponent this difficulty disappears. At the same time, new features appear, such as experimentation and possible attempts at enforcing different conventions.

Truth-telling does not fare well in the one-shot game, as 59% of the time type 2 senders hide their identity (i.e., send message 3 or 4). In the repeated game the comparable figure is 44%, which suggests that truth-telling is considered a better strategy in that context. This may just be a reflection of learning over time in the repeated game: if it is clear that any message other than message 1 is coming from type 2, then the sender might as well send message 2, and avoid any confusion. Truth-telling, in general, goes from 39% to 58% in the second half of the game, which is consistent with this dynamic learning story.

Pure risk considerations predict that receivers ought to choose action C. We note that C is used most frequently in the one-shot game as a response to message 3 or 4, suggesting that it is a response to uncertainty about the sender's type. In the repeated game, this use of action C is much less frequent. Our interpretation is that safety is a fall-back policy when a player discovers that his view of what is focal is not shared by other players. This is similar to the findings of Van Huyck, Battalio and Beil (1990,1993) that a safe strategy, in the absence of an external focal coordinating device to aid in selecting a payoff-dominant equilibrium, appears to be a fall-back strategy.

4.3 Game 2 Results (Session 2)

There is no separating equilibrium in Game 2, therefore there is no equilibrium that

corresponds to truth-telling. Nevertheless, there exist partially-separating outcomes, where types 1 and 2 send the same message, or mix over the same set of messages, and the receivers respond with some action (or mixture) in the set $\{A, B\}$. Type 2 sends a different message from types 1 and 2, and the receiver responds with action C. These outcomes correspond to partial truth-telling. A safe response, on the other hand, is for the receiver to choose action D all of the time.

CMR, in fact, has no prediction for this game, though the generalized CMR does make predictions. Specifically, sender types 1 and 2 could send messages 1, 2 or 4, and the receiver can respond with any message in $\{A, B\}$. Sender type 3 sends message 3, and the receiver responds with action C. Finally, note that all players prefer the partially separating equilibrium to the pooling equilibrium, but sender types 1 and 2 have conflicting interests (i.e., type 1 would like message 2 to be sent and believed by the receiver, while type 2 would like message 1 to be sent and believed by the receiver).

Sender and Receiver Behavior

Table 5 shows cross-tabulations for both sender type against message, and for message received against action taken, for the one-shot version of Game 2. We have not conducted a repeated version of Game 2, partly because it seemed to work so well (vis-a-vis CMR), and we did not think that the repeated version would change the aggregate results.

Although type 1 and 2 senders do not use the explicit pooling message, #4, "I am type 1 or type 2", very often (16%), the set of messages they do send appear to be treated as having this meaning. Specifically, messages 1, 2 and 4 seem to identify the sender as type 1 or type 2, since 100% of the time, the response to these messages is either action A or action B. Also, 100% of the time the response to message 3 is action C.

Table 5 about here

Discussion

Overall, the subjects in Game 2 behave very closely to the GCMR predictions. Risk (or safety) seems to play only a minor role, since it is only when message 5 ("I am type 2 or 3") is sent that action D is taken, and this only happens twice. The striking thing here is the systematic nature of the results (which are not predicted by CMR). Consistent with GCMR, people do things to get a relatively high payoff, even though there is no consistent way for them to get their highest payoff. For example, although both type 1 and type 2 senders lie, they do so in a consistent way: messages 1, 2 and 4 have the same meaning, as evidenced by the systematic pattern of responses to these messages by the receivers. This is in contrast to Game 3, considered next, in which only message 4 will allow types 1 and 2 to pool.

4.4 Game 3 Results (Session 2)

Game 3 is a greater challenge to CMR than the other games. First, note that there are two Nash equilibria: pooling, and partially-separating, in which only types 1 and 2 pool. Unlike the other games, there is no safest action here, though action C is a best response with no communication. CMR predicts that types 1 and 2 send message 4, and that any message other than 4 should be interpreted as coming from type 3. Truth-telling is not optimal for any type of sender. By Pareto efficiency, the sender *ex ante* prefers the partially-separating equilibrium, but *ex post*, type 3 does not get his highest payoff with this equilibrium (since he can be identified).

Game 3 is a bigger challenge to CMR because now types 1 and 2 must pool by sending message 4. That is, they must hide some information. At the same time, type 3 need not directly reveal his identity--thus, the CMR prediction is very different from truth-telling in this game. If type 1 decides to be an honest guy, he can expect trouble, since type 3 would like to be mistaken for type 1, and thus it will be hard for type 1 to establish that as a convention.

Sender Behavior

Table 6a shows cross-tabulations of sender type by message sent for a one-shot and a repeated version of Game 3. A striking result is that type 1 and 2 senders seem to have trouble using message 4. Analysis of individual data provides some insight. Of the 8 subjects in the one-shot version of Game 3, four of them always used message 4 when they were type 1 or 2, two of them never used message 4, and the other two used message 4 early in the game, but abandoned it later, possibly because they did not get action A in response.

Note that type 3 will do best if he can get the receiver to believe that he is a type 1, so that the receiver chooses action B. On the other hand, type 3 can identify himself as type 3, and, if believed, the receiver will choose action C, and the sender will receive his second best payoff. In the one-shot game, many type 3 senders (56%) choose message 1, presumably hoping to fool the receiver. No type 3 senders send message 3, "I am type 3", in the one-shot game.

The data indicate that although some subjects understand and use message 4, a considerable proportion do not. If the subjects cannot use message 4 because its literal meaning is not shared by all the subjects, then it becomes extremely hard for the subjects to establish any kind of communication.⁴ In contrast to Game 2, the receiver cannot treat messages 1, 2 and 4 as being equivalent, because type 3 would take advantage of the situation.

Tables 6a & 6b about here

Receiver Behavior

Table 6b shows cross-tabulations of message received by action taken in the one-shot and repeated version of Game 3. The main things to note in this table are as follows. First, the one-shot version: Message 4, apparently, is taken to be truthful 60% of the time (17/28 times the action A is chosen). Message 1, apparently, is taken to mean that the sender is type 3 in 64% of case (18/27

times the action C is chosen). Message 3 is sent only 2 times. In the repeated version, message 4 receives a response of action A 77% of the time (10/13 cases), while message 3 receives a response of action C 90% of the time (18/20 cases).

Receivers seem to understand that message 4 comes from type 1 or type 2, but they are not certain, since 40% of the responses are different from action A. In fact, type 3 never sends message 4, so it must be either that the receiver is playing safe (choosing action C and ignoring messages), or the receiver is playing hunches (choosing actions B and C in the hopes of getting a higher payoff).

One-Shot vs. Repeated Games

Something interesting happens in the repeated version. Specifically, type 1 and type 2 senders are much more likely to reveal their true type in a message in the repeated game. Overall, 58% of type 1 or 2 senders in the one-shot version, versus 16% in the repeated version, send the pooling message, "I am type 1 or 2". On the other hand, 25% of type 1 or 2 senders in the one-shot version, versus 54% in the repeated version, state their true type explicitly.

This "Full-truth-telling" behavior is hard to understand on its face: If a sender believes the receiver will choose a (myopic) best-response, then he can never expect to do better by telling the truth in Game 3. For example, if a sender says he is type 1 and is believed, then the receiver would respond with action B. Similarly, the best response to the message "I am type 2", if it is believed, is to choose action C. In both cases, the sender does worse than if he sent message 4: "I am type 1 or type 2", to which the receiver's best response is action A.

One possible explanation is that receivers in the repeated game punish "Partial-truth-telling" (saying "I am type 1 or type 2"), and reward Full-truth-telling, by allocating actions between the one that gives the sender a good payoff and the one that gives the receiver a good payoff. In the repeated game, actions A, B and C all account for a significant proportion of responses. 50% of the responses to message 1 are action A, which suggests an attempt to establish trust, or to reward the sender for

Full-truth-telling, since action B would, in fact, be the best response if message 1 were believed. As in Game 1, we do not see any explicit evidence of punishment, though this is not necessary: it would only be necessary for a sender to believe that such punishment is possible. More generally, we suspect that the repeated version of the game gives the sender less commitment power (or less of a first-mover advantage) than the one-shot version does.

The behavior of type 3 senders is different in the two versions of Game 3 as well. In the repeated version of Game 3, only 32% of type 3 senders send message 1, but 36% choose message 3.

Discussion

Perhaps the best way to summarize the results of Game 3 is to compare it again with Games 1 and 2. In Games 1 and 2, the GCMR predictions can be implemented by truth-telling, while in Game 3 this is not true. In fact, truth-telling can turn out to be a complex, confusing matter in Game 3. To the extent that truth-telling conventions arise, there is an incentive to cheat on the convention, so the convention turns out to be unstable. In general, the failure of GCMR in Game 3 seems to be related to the subtlety of behavior required: one must, as a type 1 or 2 sender, use message 4, and nothing else. If types 1 and 2 do not all understand this, then everyone can get confused.

5. Related Work on Communication and Coordination

Theoretical work on communication and coordination can be divided into two main categories: *equilibrium* and *non-equilibrium* approaches, with the former constituting most of the research. Equilibrium approaches include the seminal papers of Crawford and Sobel (1982) and Green and Stokey (1980), who first showed that in a signaling game where the sender's and receiver's interests do not coincide, there exist equilibria in which some truthful information can be

transmitted; that is, cheap talk can matter.

Farrell (1985), continuing with the equilibrium approach, developed the idea of a *Neologism-Proof* equilibrium. Multiplicity of equilibria is a problem with cheap-talk equilibria, since for any equilibrium in a simple cheap talk game, where some signal is sent with probability zero, we can construct an equivalent equilibrium where all the signals are sent with positive probability. Imagine that, perhaps through evolution, an equilibrium is reached where certain words of a language common to the players are used with their true meaning. Farrell's idea is that an equilibrium should prevail if no subset of sender types prefer to start using a "new word" than to use the equilibrium messages. This reduces the set of equilibria. Matthews, Okuno-Fujiwara and Postelwaite (1991) refined the idea of Neologism-Proofness by allowing deviations (called "announcements") that contain more than one message. Their concept of *announcement-proofness* is thus more restrictive than Farrell's.

Blume and Sobel (1995) suggest that if communication in a signaling game is costless, the sender should be able to continue to send messages once some information has been transmitted. For example if there is a message that means that the sender belongs to some subset of possible types, the sender could propose a new equilibrium, given that the sender must belong now to a smaller set. Roughly, an equilibrium would not prevail if once a message has been sent the types sending that message in equilibrium strictly prefer to keep on talking. An equilibrium where this kind of renegotiation is not possible is called *Communication-Proof*.

The approach of Rabin (1990), which we have focussed on in the experiments, is completely different from the above mentioned equilibrium approaches. He criticizes the equilibrium approach, as we have already noted in Section 2, on the grounds that truthful communication is basically a coordination problem, and that use of an equilibrium concept assumes coordination. Thus, this view suggests that the messages themselves are the primal objects. When sent and interpreted appropriately, the communication process leads one to an outcome, which need not be a full

equilibrium. The equilibrium approach, on the other hand, suggests that the equilibria are the primal objects, and that, somehow, coordination is achieved by players managing to choose one of the possible equilibria.

There has been some recent work experimentally testing the role of communication in games. For example, Cooper, Dejong, Forsythe and Ross (1989, 1992) have examined the efficacy of one- and two-way communication in achieving coordination in simultaneous-move coordination games. Our experiment differs from the coordination games studied by Cooper, et al., most importantly, in that in the sender/receiver games only one player can make a binding decision (the receiver) that affects payoffs. Moreover, this decision is made after the sender has made his (non-binding) decision of which message to send. This makes it easier to sort out how one player is responding to the other than in simultaneous-move games. The result that one-way communication seemed to work better in the Cooper, et al., experiments would seem to be due to the commitment power implicit in the one-way structure of communication. The breakdown of CMR in the repeated version of one of our games might be viewed as a related phenomenon, with the repeated nature of the game diluting the commitment power of the sender's message, compared to the one-shot game.

Palfrey and Rosenthal (1991) have tested the effects of cheap talk in public goods games. Introduction of cheap talk in such games was found to affect play. Specifically, although the messages sent by subjects were not systematic (i.e., no single theory explained the message behavior), subjects responded to messages in a systematic way. We found similarly mixed results in the repeated version of Game 2, where the conflict between the conventions of Partial- and Full-truth-telling led to a different kinds of message-sending behavior.

6. Conclusions

We have characterized the results in terms of full and partial truth-telling. We think that the majority of subjects realized that some kind of refinement of full truth-telling would be the best

policy. CMR assumes that a particular form of honesty is focal, and that players, with this in mind, can figure out the rest. Repetition in our games was intended to give subjects enough time to figure out what is focal (through interaction with other players), and what the full implications of the focal policy are.

Our results show that full understanding and implementation of CMR-type play is a difficult and demanding task for subjects. It should be emphasized that the use of the full truth-telling convention does not appear to be driven by ethical considerations, since a majority of subjects did something other than "telling the whole truth" at some point. Further, it matters what type of repetition subjects experience. For example, in the repeated games with the same opponent, new features not present in the one-shot games, such as the power of the receiver to attempt to impose a preferred convention, or concern on the part of the sender about the effect of not being fully revealing on future interactions with the same individual, arise. When there is a conflict of preferred conventions, then the game may look more like bargaining or negotiation, and it is not clear if more experience would lead to play consistent with CMR. Establishment (not to mention discernment) of a focal policy is tricky, so full truth-telling tends to be used because it is easy.

We think that CMR and GCMR are still worthy candidates for further testing, but other ways of establishing a focal policy should be investigated. The games studied in this paper were relatively simple-- simpler, for example, than the costly signaling games analyzed in Banks, et al. (1994)-- yet there were still substantial departures from the theoretical predictions. We suspect that the proximate cause of this is a failure of some of the common knowledge requirements of the theory. Specifically, how does one decide that the other players in the game have the same understanding of the game as oneself? Players could be allowed to discuss in detail, prior to play, what the "right" way to play would be, or a menu of possible ways of playing the game could be suggested. Play by experienced subjects in repeated sessions could shed light on this as well, since it would allow one to assess the predictive value of CMR for relatively sophisticated subjects, who presumably have

more sophisticated notions of the implications of various possible policies. Even in this case, however, players would need to know that the other players in the game are experienced players as well.

References

- Banks, Jeffrey, Colin Camerer, David Porter, "An Experimental Analysis of Nash Refinements in Signaling Games", *Games and Economics Behavior*, 6(1), 1-31, 1994.
- Blume, Andreas, Joel Sobel, "Communication-Proof Equilibria in Cheap-Talk Games", *Journal of Economic Theory*, 65(2), 359-382, 1995.
- Cooper, Russell, Douglas V. Dejong, Robert Forsythe, Thomas W. Ross, "Communication in the Battle of the Sexes Game: Some Experimental Results", *Rand Journal of Economics*, 20(4), 568-587, 1989.
- Cooper, Russell, Douglas V. Dejong, Robert Forsythe, Thomas W. Ross, "Communication in Coordination Games", *Quarterly Journal of Economics*, 107(2), 739-771, 1992.
- Crawford, Vincent and Joel Sobel, "Strategic Information Transmission," *Econometrica*, **50**, 1431-1452, 1982.
- Farrell, Joseph, "Credible Neologisms in Games of Communication," MIT working paper 386, 1985.
- Green, Jerry and Nancy Stokey, "Two-Person Games of Information Transmission," Harvard and Northwestern Universities, 1980.
- Matthews, S., M. Okuno-Fujiwara and A. Postelwaite, "Refining cheap talk equilibria," *Journal of Economic Theory*, 1991.
- Palfrey, Thomas R., Howard Rosenthal, "Testing for Effects of Cheap Talk in a Public Goods Game with Private Information", *Games and Economic Behavior*, Vol. 3, No. 2, 1991, pp. 183-220.
- Rabin, Matthew, "Communication Between Rational Agents", *Journal of Economic Theory*, Vol. 51, No. 1, 1990, pp. 144-170.
- Rabin, Matthew, "Corrigendum", *Journal of Economic Theory*, Vol. 53, No. 1, 1992, pp. 110-111.
- Van Huyck, John, Ray Battalio and Richard Beil, "Tacit Coordination Games, Strategic Uncertainty and Coordination Failure," *American Economic Review*, 80:234-248, 1990.
- Van Huyck, John, Ray Battalio and Richard Beil, "Asset Markets as an Equilibrium Selection Mechanism," *Games and Economic Behavior*, 5(3), 485-504, 1993.
- Zapater, Iñigo, "Generalized Communication between Rational Agents", Brown University, Department of Economics, Working Paper 93-4R, 1993.

Figure 1

		Session 1, Game 0						Session 2, Game 1			
		Action						Action			
Type		A	B	C		Type		A	B	C	
	t1	8,7	4,2	2,5			t1	8,7	1,2	4,5	
	t2	1,3	6,6	3,4			t2	1,1	4,6	10,4	
		Session 2, Game 2						Session 2, Game 3			
		Action						Action			
Type		A	B	C	D	Type		A	B	C	
	t1	8,6	5,7	1,1	0,5		t1	7,8	4,10	3,2	
	t2	6,7	7,6	2,1	0,5		t2	7,6	3,2	5,8	
	t3	1,0	1,1	6,7	2,5		t3	1,1	6,2	4,7	

Table 1: Predictions

Game	Credible Message Rationalizability	Generalized Credible Message Rationalizability	No Communication Baseline
0	Separation	Separation	Pooling
1	Type 1 separates Type 2 ?	Type 1 separates Type 2 may separate	Pooling
2	Types 1 and 2 ? Type 3 separates	Types 1 and 2 may pool Type 2 separates	Pooling
3	Types 1 and 2 pool Type 3 ?	Types 1 and 2 pool Type 3 may separate	Pooling

Table 2: Experiments Conducted

Group	Date	Session 1	Session 2	Description
1	April 6, 1992	Game 0	Game 3	one-shot games
2	April 6, 1992	Game 0	Game 1	one-shot games
3	April 7, 1992	Game 0	Game 2	one-shot games
4	April 8, 1992	Game 0	Game 1	repeated games
5	April 13, 1992	Game 0	Game 3	repeated games
6	April 14, 1992	Game 0	Game 1'	long repeated games

Table 3: Session 1, Game 0
(One-shot vs. Repeated Games)
Types vs. Messages

Groups 1, 2 and 3 (One-Shot Games)

Message Type	1	2	3	4	Total
1	44	8	4	1	57
2	7	46	5	5	63
Total	51	54	9	6	120*

Groups 4, 5 and 6 (Repeated Games)

Message Type	1	2	3	4	Total
1	57	5	7	2	71
2	5	50	6	3	64
Total	62	55	13	5	135

Messages vs. Actions

Groups 1, 2 and 3 (One-Shot Games)

Action Message	A	B	C	Total
1: "I am type 1"	42	4	5	51
2: "I am type 2"	5	42	7	54
3: "I am type 1 or 2"	5	1	3	9
4: "No message"	6	0	0	6
Total	58	47	15	120*

Groups 4, 5 and 6 (Repeated Games)

Action Message	A	B	C	Total
1	52	0	10	62
2	5	38	12	55
3	1	1	11	13
4	2	0	3	5
Total	60	39	36	135

* 10 observations dropped in periods 4 and 5 in Group 2, due to messages being sent to wrong subjects.

**Table 4a: Session 2, Game 1
(One-shot vs. Repeated Games)
Types vs. Messages***

Group 2 (one-shot)

Message Type	1	2	3	4	Total
1	38	0	2	4	44
2	2	21	23	10	56
Total	40	21	25	14	100

Group 4 (repeated)

Message Type	1	2	3	4	Total
1	38	0	2	1	41
2	4	29	13	13	59
Total	42	29	15	14	100

Group 6 (long repeated**)

Message Type	1	2	3	4	Total
1	55	1	4	1	61
2	5	33	13	13	64
Total	60	34	17	14	125

**Table 4b: Session 2, Game 1
(One-shot vs. Repeated Games)
Messages* vs. Actions**

Group 2 (One-Shot Games)

Action Message	A	B	C	Total
1: "I am type 1"	30	3	7	40
2: "I am type 2"	3	16	2	21
3: "I am type 1 or 2"	7	9	9	25
4: "No message"	0	5	9	14
Total	40	33	27	100

Group 4 (Repeated Games)

Action Message	A	B	C	Total
1: "I am type 1"	37	2	3	42
2: "I am type 2"	1	26	2	29
3: "I am type 1 or 2"	1	11	3	15
4: "No message"	2	11	1	14
Total	41	50	9	100

Group 6 (Long Repeated Games)

Action Message	A	B	C	Total
1	52	2	6	60
2	2	28	4	34
3	1	14	2	17
4	0	13	1	14
Total	55	57	13	125

** Entry for Row 2, Column C of Game 1 was changed from (10,4) to (15,4). Also, games were 25 periods long, instead of 10 periods.

**Table 5: Session 2, Game 2
(One-shot Games Only)
Types vs. Messages***

Group 3 (One-Shot Games)

Message Type	1	2	3	4	5	6	7	8	Total
1	6	17	0	5	1	0	0	0	29
2	13	8	1	4	0	0	0	0	26
3	0	1	23	0	1	0	0	0	25
Total	19	26	24	9	2	0	0	0	80

Messages* vs. Actions

Group 3 (One-Shot Games)

Action Message	A	B	C	D	Total
1: "I am type 1"	5	14	0	0	19
2: "I am type 2"	13	13	0	0	26
3: "I am type 3"	0	0	24	0	24
4: "I am type 1 or 2"	2	7	0	0	9
5: "I am type 2 or 3"	0	0	0	2	2
6: "I am type 1 or 3"	0	0	0	0	0
7: "I am type 1, 2 or 3"	0	0	0	0	0
8: "No message"	0	0	0	0	0
Total	20	34	24	2	80

**Table 6a: Session 2, Game 3
(One-shot vs. Repeated Games)
Types vs. Messages***

Group 1 (One-Shot Games)

Message Type	1	2	3	4	5	6	7	8	Total
1	8	2	0	10	0	1	0	0	21
2	1	4	2	18	2	0	0	0	27
3	18	1	0	0	3	5	3	2	32
Total	27	7	2	28	5	6	3	2	80

Group 5 (Repeated Games)

Message Type	1	2	3	4	5	6	7	8	Total
1	23	3	2	4	3	1	2	0	38
2	9	18	2	8	1	0	0	0	38
3	14	2	16	1	3	6	0	2	44
Total	46	23	20	13	7	7	2	2	120

**Table 6b: Session 2, Game 3
(One-shot vs. Repeated Games)
Messages* vs. Actions**

Group 1 (One-Shot Games)

Action Message	A	B	C	Total
1: "I am type 1"	3	6	18	27
2: "I am type 2"	0	0	7	7
3: "I am type 3"	0	0	2	2
4: "I am type 1 or 2"	17	5	6	28
5: "I am type 2 or 3"	0	0	5	5
6: "I am type 1 or 3"	0	2	4	6
7: "I am type 1, 2 or 3"	0	0	3	3
8: "No message"	0	0	2	2
Total	20	13	47	80

Group 5 (Repeated Games)

Action Message	A	B	C	Total
1: "I am type 1"	23	12	11	46
2: "I am type 2"	10	2	11	23
3: "I am type 3"	0	2	18	20
4: "I am type 1 or 2"	10	1	2	13
5: "I am type 2 or 3"	2	1	4	7
6: "I am type 1 or 3"	1	1	5	7
7: "I am type 1, 2 or 3"	1	0	1	2
8: "No message"	1	1	0	2
Total	48	20	52	120

Endnotes

1. Rabin (1990) uses the assumption of "honesty" as a focal policy. We prefer not to use the term "honesty" since its technical meaning in CMR is not the everyday sense of the word, "telling the whole truth."
2. In fact, individual data shows that less than 50% of players used a truth-telling strategy throughout the one-shot games, so it is not just that there are a few outliers in the data.
3. This is the familiar chain-store paradox. If a sender is type 1 in the last period, he will want to send message 1 then, and thus a reputation for always "fighting entry" by sending "no message" will not be credible.
4. Under CMR, the message "I am type 1 or type 2" should generate the beliefs $p(1) = p(2) = 1/2$, but some might interpret the message as generating the beliefs $(\{1,2\})$. If the typical subject has trouble with this subtlety this could account for message 4 not being used very often.