

LEARNING AND COMMUNICATION
IN SENDER-RECEIVER GAMES: AN ECONOMETRIC INVESTIGATION

Andreas Blume
(andreas-blume@uiowa.edu)
319-335-0931

Douglas V. DeJong
(douglas-dejong@uiowa.edu)
319-335-0919

George R. Neumann
(george-neumann@uiowa.edu)
319-335-0850

N.E. Savin*
(gene-savin@uiowa.edu)
319-335-0855

Tippie College of Business
University of Iowa
Iowa City, IA 52242
USA

Fax: 319-335-1956

January 29, 2000

* Corresponding author

General field: econometrics, game theory and experiments

JEL Classification Numbers: C72, C91, C92

LEARNING AND COMMUNICATION
IN SENDER-RECEIVER GAMES: AN ECONOMETRIC INVESTIGATION

January 29, 2000

By Andreas Blume, Douglas V. DeJong, George R. Neumann and N. E. Savin*

Learning and communication play important roles in coordinating activities. Game theory and experiments have made a significant contribution to our understanding and appreciation for the issues surrounding learning and communication in coordination. However, the results of past experimental studies provide conflicting results about the performance of learning models. Moreover, the interaction between learning and communication has not been systematically investigated. Our long run objective is to overcome the conflicting results and to provide a better understanding of the interaction. To this end, we econometrically investigate a sender-receiver game environment where communication is necessary for coordination and learning is essential for communication. (JEL C72, C91, C92)

LEARNING AND COMMUNICATION
IN SENDER-RECEIVER GAMES: AN ECONOMETRIC INVESTIGATION

1. INTRODUCTION

Learning plays an important role in coordinating activities be it a market or an organization. Communication goes hand in hand with learning when addressing problems of coordination. Game theory has made a significant contribution to our understanding and appreciation for the issues surrounding learning and communication in coordinating activities, e.g., Paul Milgrom and John Roberts (1992). Experiments have picked up this theme with an extensive and growing literature that investigates learning in strategic settings. However, as detailed in the related literature section of our paper, the results of these studies provide conflicting results about the performance of the chosen learning models. Moreover, the interaction between learning and communication has not been systematically investigated. Our long-run objective is to overcome the previous conflicting experimental results on learning and to investigate the interaction of learning and communication in coordinating activities. We chose our sender-receiver environment because communication is necessary for coordination and learning is essential for communication when messages are a priori meaningless.

We address these two problems by focusing on data from sender-receiver games in which documented learning takes place (Andreas Blume, Douglas V. DeJong, Yong-

Gwan Kim and Geoffrey B. Sprinkle, 1998). Sender-receiver games play a prominent role in theory and applied work in accounting, economics, finance and political science.¹

Communication is essential in this game because it is the only way to link the receiver's action to the private information of the sender. The players are members of a population of players, senders and receivers. The population environment mitigates repeated game effects and, therefore, is particularly suitable for evaluating myopic learning rules. The players learn only about actions, not strategies of other players, either privately or at the population level; this separates this paper from bulk of the empirical work on learning in strategic form games. When population information is available, it is only about sender play. In this environment, learning is essential for players to coordinate because messages are costless and *a priori* meaningless.

We compare two learning rules, the stimulus-response (SR) model of Alvin Roth and Ido Irev (1995) and a simple belief-based learning (BBL) model in the spirit of fictitious play (Julia Robinson, 1951) or one of its stochastic variants (Drew Fudenberg and David M. Kreps, 1993). SR learning can be and has been applied in a wide variety of settings. It requires only minimal cognitive abilities on the part of players. This feature of the model is appealing for those who want to show that high-rationality predictions can be derived from a low-rationality model. A closely related feature is that SR learning requires only minimal information. All that players need to know are the payoffs from their own past actions; they need not know that they are playing a game, they need not know their opponents' payoff or their past play. These two characteristics make the SR model a natural benchmark for our investigation.

On the other hand, it seems quite likely that individuals would try to exercise more cognitive ability and hence try to use other available information. In BBL models, players use more information than their own historical payoffs. This information may include their own opponents' play, the play of all possible opponents and the play of all players. Models of this kind embody a higher level of rationality; e.g. fictitious play can be interpreted as optimizing behavior given beliefs that are derived from Bayesian updating.

Our sender-receiver game environment is challenging for the learning models because in formulating the SR model for this extensive form setting we have to decide whether the players are updating actions or strategies. We choose actions on the grounds that this is cognitively simpler than strategies. For BBL it is challenging because senders do not have direct information on receiver behavior at the population level. For all treatments, we ask whether the SR model or the BBL model better describes learning. The initial step in our analysis was to fit the SR and BBL models to the data produced by the various treatments. The most striking fact of our analysis is that for common interest games both models fit the data reasonably well as measured by McFadden's (1974) pseudo- R^2 .

We let the BBL model take a form that is analogous to the SR model. In both cases choice probabilities depend on unobserved propensities. The models differ in how the propensities are updated. In the SR model the propensity for taking an action is solely dependent on a player's own past payoffs from that action, whereas in the BBL model the propensity depends on the expected payoff. Our focus attention for both the SR and BBL model is the data generated by senders. We chose senders because they have the same number of strategies in all the games considered. This makes the econometric analysis and the comparisons much more transparent.

Our empirical findings for sender data can be summarized as follows. In the setting of common interest games when information on history is available, both the SR and BBL models tend to fit the data reasonably well as measured by the coefficient of determination. In the absence of history information, however, SR fits the data substantially better than BBL. We find this reassuring since it is compatible with the view that players use history information in forming their expectations, provided it is available. For divergent interest games, the results are sensitive to the equilibrium selected. Both SR and BBL fit the data reasonably well in the case of separating equilibria, but neither performs well for pooling equilibria. The fit of BBL in the case of pooling equilibria is especially poor.

The SR and BBL models are non-nested in the sense that the SR model is not a special case of the BBL model and vice versa. Nevertheless, the two models can be artificially nested with the consequence that they can be compared using standard testing procedures. This fact has been exploited by Colin Camerer and Teck-Hua Ho (1999), who recognized that SR and BBL models are nested in a general class of models which use propensities to determine choice probabilities. In the context of artificial nesting, the results clearly show that the problem of distinguishing between SR and BBL models is particularly acute in the case of common interest games with history. On the other hand, the test results accept SR and reject BBL in the games with no history and in all but one of the divergent interest games. The results from artificial nesting are consistent with the story suggested by the coefficient of determination.

A point often overlooked in empirical work is that information about learning can only come from observations where learning occurs. Once behavior has converged,

observations have no further information about learning. Including such observations will make the SR and BBL models appear to fit better, while at the same time reducing the contrast between the models, making it difficult to distinguish the models empirically. We call this effect convergence bias. A preliminary examination suggested that our non-nested results were affected to some degree by convergence bias. Accordingly, we eliminated observations where it appeared that learning has ceased and reanalyzed the remaining data. The results of this reexamination were not markedly different from our original finding. In general, however, the convergence effect can strongly influence the test results based on artificial nesting because this kind of nesting tends to induce collinearity between the explanatory variables.

As has been noted, the results clearly show that the problem of distinguishing between SR and BBL models is particularly acute in the case of common interest games with history. The last section of the paper outlines an agenda for how we can learn more about learning and communication through experiments.

2. GAMES AND EXPERIMENTAL DESIGN

Our data are generated from repeated play of sender-receiver games among randomly matched players. Players are drawn from two populations, senders and receivers, and rematched after each round of play. The games played in each period are between an informed sender and an uninformed receiver. The sender is privately informed about his type, θ_1 or θ_2 , and types are equally likely. The sender sends a message, * or #, to the receiver, who responds with an action, a_1 , a_2 or a_3 . Payoffs depend on the sender's private

information, his type, and the receiver's action, and not on the sender's message. The payoffs used in the different treatments are given in Table 1 below, with the first entry in each cell denoting the sender's payoff and the second entry the receiver's payoff. For example, in Game 2, if the sender's type is θ_1 and the receiver takes action a_2 , the payoffs to the sender and receiver are 700,700, respectively.

A strategy for the sender maps types into messages; for the receiver, a strategy maps messages to actions. A strategy pair is a Nash equilibrium if the strategies are mutual best replies. The equilibrium is called separating if each sender type is identified through his message. In a pooling equilibrium, the equilibrium action does not depend on the sender's type; such an equilibrium exists for all sender-receiver games. In Game 2, an example of a separating equilibrium is one where the sender sends * if he is θ_1 and # otherwise and the receiver takes action a_2 after message * and a_1 otherwise. An example of a pooling equilibrium is one in which the sender, regardless of type, sends * and the receiver always takes action a_3 .

A replication of a game is played with a cohort of twelve players, six senders and six receivers. Players are randomly designated as either a sender or receiver at the start of the replication and keep their designation throughout. In each period of a game, senders and receivers are paired using a random matching procedure. Sender types are independently and identically drawn in each period for each sender.

In each period, players then play a two-stage game. Prior to the first stage, senders are informed about their respective types. In the first stage, senders send a message to the receiver they are paired with. In the second stage, receivers take an action after receiving

a message from the sender they are paired with. Each sender and receiver pair then learns the sender type, message sent, action taken and payoff received

To ensure that messages have no a priori meaning, each player is endowed with his own representation of the message space, i.e. both the form that messages take and the order in which they are represented on the screen is individualized. The message space $M = \{*, \#\}$ is observed by all players and for each player either appears in the order $\#,*$ or $*,\#$. Unique to each player, these messages are then mapped into an underlying, unobservable message space, $M = \{A,B\}$. The mappings are designed such that they destroy all conceivable focal points that players might use for a priori coordination, Blume et al. (1998). The representation and ordering are stable over the replication. Thus, the experimental design focuses on the cohort's ability to develop a language as a function of the game being played and the population history provided.

Note that in this setting we learn the players' action choices, not their strategies. Also, the players themselves receive information about actions, not strategies. They do not observe which message (action) would have been sent (taken) by a sender (receiver) had the sender's type (message received) been different. This is important for how we formulate our learning rules; e.g. the hypothetical updating (see Camerer and Ho (1997)) of unused actions that occurs in BBL does not and cannot rely on knowing opponents' strategies but instead uses information about the population distribution of play.

We consider five experimental treatments, each with three replications. Each replication is divided into two sessions, Session I, which is common to all treatments and Session II, which varies across treatments. We concentrate on Session II data. The treatments examined here differ in terms of the players' incentives and the information that

is available after each period of play. For one treatment, the only information available to a player is the history of play in her own past matches. Two questions are examined for these cases. The first is whether learning takes place. If learning does take place, the second question is whether learning is captured by the SR model. In all the other treatments, there is information available to the players in addition to the history of play in their own past matches. For both senders and receivers, this information is the history of play of the population of senders. Three questions are examined for these treatments. The first again is whether learning takes place. If learning does take place, the second question is whether learning is different from that in the previous treatment, and the third is whether the BBL model better describes learning than the SR model.

The data from the experiments in Blume et al. (1998) consists of three replications for each game. Replications for Game 1 and 2 were played for 20 periods and Game 3 and 4 for 40 periods. There were two different treatments conducted with Game 1, one with and one without population history. The treatments are summarized in Table 2.²

In this paper we focus on the analysis of sender behavior. The attraction of concentrating on sender behavior is that senders have the same number of strategies in all of our treatments. A potential drawback of this focus is that since senders do not receive information about the history of receiver play at the population level, they cannot form beliefs based on that information. Instead they have to make inferences from what they learn about sender population.

3. ESTIMATION OF SR MODEL

In this section we report the results of estimation for the SR model of behavior using sender data. The SR and BBL models both use propensities to determine choice probabilities. In our extensive form game setting, we have to make a choice of whether we want to think of players as updating propensities of actions or of strategies. Both choices constrain the way in which the updating at one information set affects the updating at other information sets. If actions are updated, then there are no links across the information sets. If strategies are updated, then choice probabilities change continually at every information set. We chose updating of actions, which amounts to treating each player-information set pair as a separate player. We use the index i to refer to such a pair (ι, θ) , where ι denotes one of the six senders, θ a type realization for that sender, and the pair i is called a player.

By SR we mean that individual play is affected only by rewards obtained from own past play. Specifically, following Alvin E. Roth and Ido Erev (1995) define the propensity, $Q_{im}(t)$, of player i to send message m at time t as:

$$Q_{im}(t) = \phi_0 Q_{im}(t-1) + \phi_1 X_{im}(t-1) \quad (1)$$

where $X_{im}(t-1)$ is the reward player i receives from sending message m at time $t-1$. Time here measures the number of occurrences of a specific type for a fixed sender and ϕ_0 measures the lag effect (i.e. the importance of forgetting). Note that t is not real time. We refer to the parameter of $Q_{im}(t-1)$ as the forgetting parameter and the parameter of $X_{im}(t-1)$ as the reward parameter. Given this specification of propensities, the probability that player i sends message m is a logit function³

$$P_{im}(t) = \Pr(\text{Player } i \text{ sends message } m \text{ at time } t) = \frac{\exp(Q_{im}(t))}{\sum_m \exp(Q_{im'}(t))}. \quad (2)$$

To complete the specification of the SR model we require an initial condition for the propensities- the values of $Q_{im}(1)$. In the spirit of Roth and Erev (1995) we set $Q_{i1}(1) = Q_{i2}(1) = 350$, which is on the scale of rewards received by participants in the experiments.

The senders, who can be of two types, can send message “1” or “2”. Let $y = I\{\text{message} = \text{“2”}\}$, where $I\{A\}$ is the indicator function that takes the value 1 if event A occurs and 0 otherwise. The log likelihood function for the sender data is

$$\ln L(\mathbf{j}_0, \mathbf{j}_1) = \sum_{i=1}^N \sum_{t=2}^T [y_i(t) \ln(P_{i2}(t)) + (1-y_i(t)) \ln(1-P_{i2}(t))] \quad (3)$$

where $P_{i2}(t)$ is the probability of sending message “2”.

To show explicitly how the likelihood function (3) depends on the parameters ϕ_0 , ϕ_1 , it is convenient to rewrite the propensity (2) as a partial sum:

$$Q_{im}(t) = \mathbf{j}_0^t Q_{im}(1) + \mathbf{j}_1 \sum_{j=1}^{t-1} \mathbf{j}_0^{t-1-j} X_{im}(j). \quad (4)$$

Using (4), the probability $P_{i2}(t) = 1/[1 + \exp(\Delta Q(t))]$ where

$$\Delta Q(t) = \mathbf{j}_0^t (Q_{i1}(1) - Q_{i2}(1)) + \mathbf{j}_1 \sum_{j=1}^{t-1} \mathbf{j}_0^{t-1-j} (X_{i1}(j) - X_{i2}(j)) = \mathbf{j}_1 \sum_{j=1}^{t-1} \mathbf{j}_0^{t-1-j} (X_{i1}(j) - X_{i2}(j)) \quad (5)$$

since $Q_{i1}(1) - Q_{i2}(1) = 0$.

From (5), we see that the identification of the parameters ϕ_0 and ϕ_1 depends on the speed of learning. Consider the behavior of the difference in the rewards, $X_{i1}(t) - X_{i2}(t)$, $t = 1, 2, \dots$. If the players converge to equilibrium in the first round, then the

difference in the rewards does not change over time: $X_{i1}(t) - X_{i2}(t) = c$ for $t = 1, 2, \dots$. The consequence is that φ_0 and φ_1 are not identified. Of course, if the reward parameter is zero, $\varphi_1 = 0$, then φ_0 is not identified.

More generally, the speed of learning determines the amount of the available sample information that is relevant for estimating the parameters. Suppose $X_{i1}(t) - X_{i2}(t) = c$ for $t > T^*$ then increasing T beyond T^* , $T > T^*$, will not increase the precision of the estimator. Rapid learning means that T^* is small, and hence there is, relatively speaking, little information available to estimate the parameters. On the other hand, increasing T beyond T^* will appear to improve the fit of the model to the data. We refer to this effect as convergence bias. Convergence bias is discussed in more detail in section 8.

The likelihood function was maximized separately for each of the 15 replications using observations from period 2 onwards. The results of doing so are shown in Table 3. Columns 2 and 3 of the table contain the maximum likelihood (ML) estimates of φ_0 and φ_1 and columns 4 and 5 contain the log likelihood value at the optimum, and McFadden's (1974) psuedo- R^2 statistic. The psuedo- R^2 statistic is defined as $1 - \ln(L^U)/\ln(L^R)$ where L^U is the maximized value of the likelihood and L^R is the value of the likelihood function when $\varphi_1 = 0$ and $\varphi_1 = 0$, $L^R = L(0, 0)$. From (2) the value of the $\ln L(0,0) = -NT\ln(2) = -(6 \times 20)\ln(2) = -83.17$. This is because $\varphi_1 = 0$ and $\varphi_1 = 0$ imply that $Q_{im}(t) = 0$ for all t , and hence that the probability of message "2" is $P_{i2}(t) = 0.5$.

The feature that stands out in Table 3 is the fit of the SR model. When the parameters are chosen optimally, the SR model fits the experimental data well for the common interest games with history, judging by the psuedo- R^2 values reported in column 5. In general, the SR model tends to fit Game 1 with history better than Game 1 without

history. A puzzle is why history should matter for SR. The SR model performs well for the divergent interest games when there is convergence to a separating equilibrium, which occurs in G3R1 and G3R2. For G4R3 the relatively good performance is explained by the fact that one sender type, θ_2 , is almost fully identified even though the pooling equilibrium occurs. It should be noted that the ML estimates maximize the log likelihood, not the psuedo- R^2 .

Figure 1 shows the probability of sending message 2 for each agent type by period for the 15 replications. The smooth line marked with triangles shows the predicted fraction of type 1 agents playing message “2” each period while the smooth line marked with circles shows the model's predicted fraction of type 2 agents playing message “2”. The fraction of type 1 agents actually playing message “2” is shown by the upper jagged line while the fraction of type 2 agents actually playing message 2 is shown by the lower jagged line. Thus, in the game shown in the top left-hand graph in the figure - G1R1 - 50% of type 1 agents play message “2” in round 1, as do 50% of type 2 agents. By period 7, 100% of type 1 agents play message 2, and 100% of type two agents play message 1. A similar pattern appears in replications 2 of Game 1 and in all three replications of Game 2. The players converge to equilibrium less rapidly in replication 3 of Game 1 and in the divergent interest games. A complete discussion of the empirical patterns in these experiments is given in Blume et al. (1998). Figure 1 demonstrates that in many experiments SR fits the data of BDKS reasonably well.

4. ESTIMATION OF THE BBL MODEL

An alternative literature (e.g., John B. Van Huyck, Raymond C. Battalio, and Fredrick W. Rankin (1996), Yin-Wong Cheung and Daniel Friedman (1997)) argues that

variants of fictitious play–BBL–are better characterizations of play in experiments. BBL is expected to dominate SR because BBL uses more information, namely in our setting the experiences of other participants.

In the BBL model we define the propensity, $Q_{ij}(t)$, of player i to play action j at time t as:

$$Q_{im}(t) = \beta_0 Q_{im}(t-1) + \beta_1 \bar{X}_{im}(t-1) \quad (6)$$

where $\bar{X}_{im}(t-1)$ is the expected reward of player i from taking action j based on data available at time $t-1$. The expected reward is calculated using the past frequencies of play. For the senders, $\mathbf{h}^{t-1}(\mathbf{q}|m)$ is the frequency of type \mathbf{q} given message m in period $t-1$, and for receivers, $\mathbf{r}^{t-1}(a|m)$ is the frequency of action a given message m in period $t-1$.

The choice probabilities again are logit as in (2) with (6) replacing (1) as the definition of Q_{ij} , and in the likelihood function (3). The models are similar in that both use propensities to determine choice probabilities.

In our setting, the senders do not observe the history of receiver play and hence calculate the *expected* frequency of action a given message m , $\hat{\mathbf{r}}^{t-1}(a|m)$, from data on past sender play. Thus, the sender formula for the expected reward is

$$\bar{X}_{(??),m}(t-1) = \sum_a u(a, \mathbf{q}) \hat{\mathbf{r}}^{t-1}(a|m) \quad (7)$$

where

$$\hat{\mathbf{r}}^{t-1}(a|m) = \left\{ \begin{array}{l} 1 \text{ if } a = \arg \max_{a'} \sum_{t \leq t-1} \sum_{\mathbf{q}} u(a', \mathbf{q}) \mathbf{h}^t(\mathbf{q}|m) \\ 0 \text{ if } a \notin \arg \max_{a'} \sum_{t \leq t-1} \sum_{\mathbf{q}} u(a', \mathbf{q}) \mathbf{h}^t(\mathbf{q}|m) \\ 0.5 \text{ otherwise} \end{array} \right\}.$$

Table 4 contains the estimates for each of the 15 replications. Columns 2 and 3 of

the table contain the ML estimates of β_0 and β_1 , and columns 4 and 5 contain the log likelihood value at the optimum, and the psuedo- R^2 statistic evaluated at the ML estimates.

Two features stand out in the table. First, in the case of the common interest games, the BBL model performs as expected in the sense that it tends to fit better when history information is available. The comparison of psuedo- R^2 's is suggestive: BBL fits about as well as SR, sometimes better and sometimes worse, in common interest games with population history information while SR wins without that information. SR tends to fit better than BBL in divergent interest games. The fit of this model is illustrated in Figure 2, which shows the relation of predicted response and actual response by period. In Figure 2 what stands out is the poor performance of BBL in Game 1 without history and in Game 4.

Second, β_0 , the forgetting coefficient in the BBL model, is negative in two replications of Game 1 without history, G1NHR1 and G1NHR2, and in two replications of Game 4, G4R1 and G4R2. In every one of these four cases, the fit of the BBL model is very poor, which suggests that the BBL model is not the correct model.

For these four cases, we also note that the null hypothesis that coefficient of the expected reward is equal to zero, $\beta_1 = 0$, is overwhelmingly accepted by the conventional likelihood ratio test. In the conventional test, the asymptotic critical value is based on the chi-square distribution with one degree of freedom. For each of the four likelihood ratio tests, the P-value based on the chi-square (1) distribution is at least 0.30. Note that the null hypothesis $\beta_1 = 0$ is also accepted by the conventional likelihood ratio test for G3R3 as well as G4R3.

The test for $\beta_1 = 0$ (i.e., there is no learning) is interesting because when $\beta_1 = 0$ is true, the forgetting parameter β_0 is not identified. Thus, it not so surprising that a negative estimate of β_0 occurs. Although the above test results for $\beta_1 = 0$ appear to support the conclusion that no learning is taking place, these test results must be treated with utmost caution. Even if the BBL model is correct, which is very doubtful, the conventional test may produce misleading results due to the fact that nuisance parameter β_0 is not identified when the null $\beta_1 = 0$ is true.

As in the case of the SR model, the lack of identification is due to the fact that $Q_{i1}(t) = Q_{i2}(t)$ for all t when $\beta_1 = 0$, which implies that $P_{i2}(t) = 0.5$ for all t . Hansen (1996), Andrews (1997), Stinchcombe and White (1998) and others have investigated the general case of hypothesis testing when the nuisance parameter is not identified under the null. Their results show that the relevant asymptotic distribution theory is nonstandard: the asymptotic distribution of the test statistic, instead of being chi-square, may depend on the value of the nuisance parameter.

5. POOLING

In above analysis, we have treated each replication of a particular game as independent, and obtained separate estimates of the parameters from each. Since the underlying learning rule is assumed to be the same in each replication it may make sense to pool the observations across each replication, and potentially across the different games. If such aggregation were warranted it would admit more precise inference. Whether one can pool the replications of a particular experimental game involving learning is of independent interest. One obvious circumstance where pooling would be inappropriate is one where the participants followed different learning rules. In this case combining the

replications of a game will pool learning processes that are fundamentally different, i.e., have different values of ϕ_i , or β_i , $i = 0, 1$.

Alternatively, even where the learning process is the same, the evolution of play may be quite different because of the incentives that the participants face. For example, in the games that we study, Games 1 and 2 have an equilibrium that both senders and receivers strictly prefer. Note that in Game 2 there is also a pooling equilibrium where action a_3 is taken regardless of the message sent. In Games 3 and 4 by way of contrast the interests of the senders and receivers are not perfectly aligned, with receivers preferring the separating over the pooled equilibrium relatively more in Game 3 than Game 4. The possibility of different equilibrium selection in these games leads to the possibility that the outcomes are situation specific and that the data from the experiments cannot be simply pooled.

To examine whether the replication data can be pooled we fit the SR model to the three replications of each experimental game. Then the standard chi-squared test based on minus 2 times the difference of the log likelihoods is used to compare the constrained and unconstrained values of the likelihood function. There are three replications for each game so there are $2 \times 3 = 6$ parameters estimated in the unconstrained case and 2 in the constrained case, resulting in chi-square tests with 4 degrees of freedom. The results are displayed in Table 5. In both Games 1 and 2 with history the data suggest that a common learning rule is followed. Indeed, the marginal significance level of the test statistics is at least 35% across the tests, suggesting that it is not even a close call. In contrast, the data for Games 3 and 4 are not compatible with a common learning model. In each case the chi-square test rejects the hypothesis of a learning rule with common parameter values

and, again, the rejection is not a close call. Of course, other combinations could be tested for pooling. The message delivered by the above results is that even pooling across the replications of the same experimental game is not always possible, and hence the focus of the subsequent statistical analysis is on individual replications.

6. SHAPE OF LIKELIHOOD FUNCTIONS

The likelihood functions for the SR and BBL models are not globally concave. As a consequence, standard maximization routines are not guaranteed to locate a global optimum. In fact, we frequently found that quasi-Newton methods alone would get stuck at local optima, or wander aimlessly in flat sections of the function. Figure 3 shows the typical shape of the likelihood function for SR models in games of common interest, in this case G1R3, while Figure 4 shows the shape of the likelihood function for the BBL model in the case of G1R3.

Note that although the likelihood functions are shaped quite differently, both exhibit large flat spots and multiple local optima. The SR likelihood function also characteristically has a long ridge, as the contour plots in Figure 3 illustrate.

To obtain the MLE's we initially used a grid search on the parameter space, evaluating the function 2,500 times. This grid search always included the point (0,0) so that the no-learning model was always considered. The maximum maximum of these points was used to pick the starting values for the first derivative method. We used BFGS() with golden section search to obtain an optimum. Graphical inspection of the likelihood surface was used to assess whether a global optimum had been reached.

7. COMPARING SR AND BBL MODELS

The results in Tables 3 and 4 show that SR and BBL tend to fit about equally well for the common interest games with history on the basis of the psuedo- R^2 values, and that BBL fits better in the common interest games with history than without history. Further, SR tends to fit better for the divergent interest games. While this is suggestive, it is difficult to make probabilistic comparisons between the SR and BBL models because the models are *non-nested*. By a *non-nested model* we mean that the model being tested, the null hypothesis, is not a special case of the alternative model to which it is compared. In our setting, the two models can be artificially nested, however, with the consequence that they can be compared using standard testing procedures. This fact has been exploited by Colin Camerer and Teck-Hua Ho (1999), who recognized that SR and BBL models are nested in a general class of models which use propensities to determine choice probabilities

In this section, the approach of Camerer and Ho (1999) is adapted to our experimental data. The updating equation is modified to include both own and expected payoffs:

$$Q_{ij}(t) = \mathbf{g}Q_{ij}(t-1) + \mathbf{g}\bar{X}_{ij}(t-1) + \mathbf{g}X_{ij}(t-1) \quad (8)$$

where $\bar{X}_{ij}(t-1)$ is the expected payoff from making choice j at time $t-1$ and $X_{ij}(t-1)$ is the actual payoff from making choice j at time $t-1$. We refer to the model with the updating equation (8) as the hybrid model since it is a combination of the SR and BBL models.⁴

The likelihood function for the sender data was maximized separately for each of the 15 replications using observations from period 2 onwards. The results of maximization

are shown in Table 6. Columns 2, 3 and 4 of the table contain the estimates of γ_0 , γ_1 and γ_2 and columns 5 and 6 contains the log likelihood value at its optimum and the value of psuedo- R^2 . The fit of the hybrid model is usually better than the fit of either the SR or BBL model. Note that the ML estimation algorithm failed to converge for G2R2 and G2R3 due a singular Hessian.

In Table 7, columns 2 and 3 report the Wald t statistics for testing the hypotheses $H_0: \gamma_1 = 0$ and $H_0: \gamma_2 = 0$. Each null hypothesis is evaluated using a symmetric test at the nominal 5% level and hence by the Bonferroni inequality the probability of simultaneously rejecting both hypotheses is approximately 0.10. The results of this pair of tests are summarized in columns 4, 5, 6 and 7. Rejection of $H_0: \gamma_1 = 0$ is interpreted as acceptance of BBL and rejection of $H_0: \gamma_2 = 0$ as acceptance of SR.

From an inspection of these columns, we see that in 10 out of the 13 replications SR is accepted and BBL is rejected. This is consistent with ranking of the models based on the psuedo- R^2 's reported in Table 3 and 4. As is intuitively plausible, this result occurs mainly in Game 1 without history and the Games 3 and 4, the divergent interest games. Both SR and BBL are rejected in 2 out of 13 replications ($H_0: \gamma_1 = 0$ is accepted and $H_0: \gamma_2 = 0$ is accepted); this occurs in Games 1 and 3, G1R2 and G3R1. In Game 1 with history the results are mixed. BBL was accepted and SR rejected ($H_0: \gamma_1 = 0$ is rejected and $H_0: \gamma_2 = 0$ is accepted) in G1R1 while the reverse is true in G1R3. Both are rejected in G1R2. The conclusions drawn from this testing exercise are similar to those produced by comparing the psuedo- R^2 's in Tables 3 and 4.

The tests were also conducted at the nominal 2.5% level, that is, using a critical value of 2.27. In this case, the probability of simultaneously rejecting both hypotheses is

approximately 0.05 using the Bonferroni inequality. The decrease in the nominal probability of making a Type I error only affected the test results for G3R3. Instead of only BBL being rejected, both BBL and SR are now simultaneously rejected.

There are two potential problems with artificial nesting using the hybrid model. The first is that if $H_0: \gamma_1 = 0$ and $H_0: \gamma_2 = 0$ is true, then γ_0 is not identified. The argument is similar to that previously discussed in connection with the BBL model. In this situation, the usual asymptotic distribution theory does not apply: the asymptotic distributions of the t statistics for testing $H_0: \gamma_1 = 0$ and $H_0: \gamma_2 = 0$ are not, in general, standard normal when the nulls are true. Hence, the finding that the null hypotheses $H_0: \gamma_1 = 0$ and $H_0: \gamma_2 = 0$ are simultaneously accepted in two cases when the absolute critical value is 2.00 and in three cases when the absolute critical value is 2.27 must be treated with caution.

Second, the explanatory variables may be highly collinear. Our results appear to be influenced by multicollinearity between the expected reward $\bar{X}_{ij}(t-1)$ and the actual reward $X_{ij}(t-1)$. The source of the multicollinearity is due to learning. If play converges to a pure strategy equilibrium, then the actual and expected reward variables take on identical values. Thus, after convergence to equilibrium there is exact collinearity between the two information variables. Multicollinearity tends to be more of a problem in the common interest games; in particular, the rapid convergence to equilibrium appears to explain the singular Hessian in G2R2 and G2R3. The degree of multicollinearity depends on the number of observations included in the data after convergence has been achieved. This matter is discussed in more detail in the following section.

8. CONVERGENCE BIAS

It is common practice to include all observations from a particular experiment in any statistical estimation or testing exercise based on that experiment. Yet information about learning can only come from observations where learning occurs. As noted earlier, once behavior has converged, observations have no further information about learning. Including such observations will make the model appear to “fit” better, while at the same time reducing the contrast between models, making it difficult to distinguish the models empirically. The data shown in Figures 1 and 2 indicate that convergence typically occurs within 5 to 7 periods, while observations are included in the estimation for the entire sequence, in these data up to 20 periods.

To illustrate the possible bias that this might cause we calculated the psuedo- R^2 and average log likelihood (the maximized log likelihood divided by number of observations used) by progressively removing observations from the later periods, that is, by removing observations that include convergence. Figure 5 illustrates this bias for the replication G1R1 for BBL and SR. Under the hypothesis of no convergence bias we would expect the psuedo- R^2 and the average log likelihood to be invariant to the number of periods included. Hence, under the hypothesis, the curves in panels (a) and (b) in Figure 5 would have zero slope. In fact, all four curves have positive slope, which is characteristic of convergence bias. However, the difference between the curves in each panel is approximately constant in these data, which suggests that convergence bias makes both models appear to fit the data better, but does not otherwise bias the comparison of SR and BBL.

To measure the amount of bias requires taking a position on when convergence has occurred, a classification that is better made on individual data. We define the convergence operator $T_p(y_{it})$ by

$$\begin{aligned} T_p(y_{it}) &= 1 \text{ if } y_{it} = y_{it-1} = \dots = y_{it-p} \\ &= 0 \text{ else} \end{aligned} \tag{9}$$

In other words a player's (pure) strategy is said to have converged if the same action is taken p times in a row.⁵ To eliminate convergence bias one simply excludes observations where $T_p = 1$. We used this procedure with $p = 3$ and $p = 4$ to assess the extent of this bias. We found that at least for these data, the extent of the bias was small.

9. RELATED LITERATURE

There is an extensive and growing literature in experimental economics on learning. For example, see Richard T. Boylan and Mahmoud A. El-Gamal (1993), Camerer and Ho (1999), Cheung and Friedman (1997), David J. Cooper and Nick Feltovich (1996), James C. Cox, Jason Shachat and Mark Walker (1995), Vincent P. Crawford (1995), Roth and Erev (1995).⁶ The literature generally focuses on two broad classes of learning models, stimulus-response and belief based play. A wide variety of games are considered with various designs, e.g., whether or not players are provided with the history of the game. The performances of the learning models are evaluated using simulation and various statistical techniques. Unfortunately, the findings are mixed at best. This could be due to statistical issues, Drew Fudenberg and David K. Levine's (1998)

conjecture that with convergence to Nash in the "short term," the models maybe indistinguishable, or a combination of the two.

Roth and Erev (1995) focus on the stimulus-response model. Their concern is high (super rationality) versus low (stimulus-response) game theory and intermediate (e.g., Fudenberg and Levine's "short term") versus asymptotic results. Their model is a simple individual reinforcement dynamic in which propensities to play a strategy are updated based upon success of past play. Using simulation, they compare the simulated results to their experimental results. The simulated outcomes are similar to observed behavior and, more importantly, vary similarly across the different games considered. They interpret this as robustness of the intermediate run outcomes to the chosen learning rule. Comparisons have also been made between the stimulus-response model of Roth and Erev or similar reinforcement dynamics (e.g., Robert R. Bush and Frederick Mosteller (1955) and John G. Cross (1983)) and other learning models. Using logit, simulation or other statistical techniques, the general conclusions of these papers are that stimulus-response works well and that additional information when part of the design makes a difference; see Dilip Mookherjee and Barry Sopher (1994, 1997), Ido Erev and Amnon Rapoport (1998) and Ido Erev and Alvin E. Roth (1997).

Using a variety of games and an information condition (with and without game history) in an extended probit, Cheung and Friedman's (1997) find that the belief based model has more support than stimulus-response learning and information matters. Boylan and El-Gamal (1993) find that fictitious play is the overwhelming choice when compared with Cournot learning in their evaluation. Using the standard logit model, Van Huyck et al. (1997) focus on symmetric coordination games and evaluate the performance of the

replicator dynamic, fictitious and exponential fictitious play. Exponential fictitious play does best. Models of reinforcement learning can be used to justify the replicator dynamic (e.g., Tilman Boergers and Rajiv Sarin (1997)).

Richard D. McKelvey and Thomas R. Palfrey's (1995) model of quantal response equilibria in normal form games deserves attention here. The quantal response model is Nash with error. Using the logit model, they find that the quantal response model wins when compared to Nash without error and random choice. Important for us is their conclusion that errors are an important component in explaining experimental results. This has been implicitly assumed in the previous studies when logits and probits are used and explicitly assumed in the Erev and Roth (1997) study with simulations.

The lack of general findings in these and other papers has prompted Camerer and Ho (1999) to essentially give up on the horse race and develop a general model, which has as special cases the principal learning models in the literature. The key that ties the SR models to the BBL models is the reinforcement used. In the SR model, only actions that were taken are updated based on the actual payoffs, and in the BBL model every action is updated based on its hypothetical payoff, that is, the payoff it would receive had the action been taken. When actual and hypothetical payoffs are the same so are the models. Using maximum likelihood estimation under the constraints of logit, Camerer and Ho evaluate the possible contribution of the general model across a variety of games. As one would hope, the general model explains more of the variation.

Reinhard Selten (1997) is the true agnostic. He claims there is not enough data to form any conclusions, either theoretical or statistical. The best we can do is very general

qualitative models (e.g., learning direction theory) in which there are tendencies that are distinct from random behavior but nothing more. This view brings us full circle to Fudenberg and Levine's conjecture about whether you can distinguish among the models if equilibrium play is observed in the "short term" or alternatively, the statistical issues make such comparisons moot. The resolution of this debate is ultimately an empirical one. It is an important debate to address because of its implications for understanding the roles of learning and communication in coordinating activities.

10. SUMMARY AND FUTURE AGENDA

In this paper, we focused on the two principal learning models used in the literature, SR and BBL, using a sender-receiver game. In the experiment, an extensive form game is played repeatedly among players who are randomly matched before each round of play. This population-game environment is particularly appropriate for a comparison of myopic learning rules, if we believe that it lessens the role of repeated game strategies. Sender-receiver games with costless and *a priori* meaningless messages have the advantage that no matter how we specify the incentives, coordination requires communication and communication requires learning.

One consequence of studying learning in extensive form games is that since players in the experiment observe only each others' actions, not strategies, the principal difference between the two learning models is in the roles they assign to own experience versus population experience. Another consequence is that there are different natural specifications even for a learning model as simple as SR; we chose the cognitively least

demanding one, in which experience at one information set does not transfer to other information sets. Our results show that both the SR and BBL models fit the data reasonably well in common interest games with history. On the other hand, the test results accept SR and reject BBL in the games with no history and in all but one of the divergent interest games.

One of our findings, the inability to distinguish between the SR and BBL models in common interest games with history, illustrates the problems with discriminating among proposed learning models. The issue raised by this problem and those of other is how to learn more about learning from experiments. The starting point for our agenda for the future is based on the fact that learning models in games specify the data generating process. As a consequence, the models can be simulated. This opens up the possibility of investigating problems of optimal experimental design in game theory. It is worth stressing that our treatment of testing with experimental data has been, from a statistical point of view, entirely conventional. We have assumed that standard asymptotic theory provides a reliable guide for inference in models with sample sizes encountered in experimental economics.

Approximations based on asymptotic theory may be poor for sample sizes typically used in practice. In particular, the true probability of making a Type I error may be very different than the nominal probability. The simulated data can be used to estimate the probability of making a Type I error for tests based on the conventional asymptotic critical values. If asymptotic critical values do not work, then the performance of other approximations can be investigated, for example, bootstrap-based critical values.

Moreover, as has been noted, if the nuisance parameter is not identified when the

null is true, then the usual asymptotic theory is not applicable. This situation arises when testing that the reward parameter is zero. The simulated data can also be used to investigate the distribution of the test statistics in this case.

Once the probability of making a Type I error is under control, the powers of the tests can be examined. This will tell the sample size needed to be able to discriminate between the models. Finally, considerations of power lead to an optimal design framework, a framework that will enable us to design our experiments so as to learn more about learning.

REFERENCES

- Andrews, D.K.K., "A Conditional Kolomorgov Test," *Econometrica*, September 1997, 65, pp. 1097-1128.
- Austen-Smith, David; "Information Transmission in Debate," *American Journal of Political Science*, February 1990, 34(1), pp. 124-152.
- Blume, Andreas; DeJong, Douglas V., Kim, Yong-Gwan and Sprinkle Geoffrey. "Experimental Evidence on the Evolution of the Meaning of Messages in Sender-Receiver Games," *American Economic Review*, December 1998, 88 (5), pp 1323-1340.
- Boergers, Tilman and Sarin, Rajiv. "Learning Through Reinforcement and Replicator Dynamics," *Journal of Economic Theory*, November 1997, 77, pp. 1-14.
- Boylan, Richard T. and El-Gamal, Mahmoud A. "Fictitious Play: A Statistical Study of Multiple Economic Experiments," *Games and Economic Behavior*, April 1993, 5 (2), pp. 205-222.
- Bush, Robert R. and Mosteller, Frederick. *Stochastic Models of Learning*, New York: Wiley, 1955.
- Camerer, Colin and Ho, Teck-Hua. "Experience-weighted Attraction Learning in Games: A Unifying Approach", *Econometrica*, July 1999, 67(4), pp. 827-874.
- Cheung, Ying-Wong and Friedman, Daniel. "Individual Learning in Normal Form Games: Some Laboratory Results," *Games and Economic Behavior*, April 1997, 19(1), pp. 46-79.
- _____ "A Comparison of Learning and Replicator Dynamics Using Experimental Data," University of California at Santa Cruz working paper, University of California at Santa Cruz, October 1996.
- Cooper, David J. and Feltovich, Nick. "Reinforcement-Based Learning vs. Bayesian Learning: A Comparison," University of Pittsburgh working paper, University of Pittsburgh. 1996,
- Cox, James C.; Shachat, Jason and Walker, Mark. "An Experiment to Evaluate Bayesian Learning of Nash Equilibrium," University of Arizona working paper, University of Arizona. 1995
- Crawford, Vincent P. "Adaptive Dynamics in Coordination Games," *Econometrica*, January 1995, 63(1), pp.103-143.

- Crawford, Vincent P. and Sobel, Joel. "Strategic Information Transmission," *Econometrica*, November 1982, 50(6), pp. 1431-1452.
- Cross, John G. *A Theory of Adaptive Economic Behavior*, Cambridge: Cambridge University Press. 1983.
- Davidson, Russell and MacKinnon, James G. "Convenient Specification Tests for Logit and Probit Models," *Journal of Econometrics*, July 1984, 25(3), pp. 241-62.
- _____. *Estimation and Inference in Econometrics*, Oxford, University of Oxford Press. 1993.
- Erev, Ido. "Signal Detection by Human Observers: A Cutoff Reinforcement Learning Model of Categorization Decisions under Uncertainty," Technion working paper, Technion. 1997.
- Erev, Ido and Rapoport, Amnon. "Coordination, "Magic," and Reinforcement Learning in a Market Entry Game." *Games and Economic Behavior*, May 1998, 23(2), pp. 146-175.
- Erev, Ido and Roth, Alvin E. "Modeling How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," University of Pittsburgh working paper, University of Pittsburgh. 1997.
- Farrell, Joseph and Gibbons, Robert. "Union Voice," Cornell University working paper, Cornell University, 1990.
- Fudenberg, Drew and Kreps, David M. "Learning Mixed Equilibria," *Games and Economic Behavior*, July 1993, 5(3), pp. 320-67..
- Fudenberg, Drew and Levine, David K. *Theory of Learning in Games*, Cambridge: MIT Press. 1998.
- Gigler, Frank. "Self-Enforcing Voluntary Disclosures," *Journal of Accounting Research*, Autumn 1994, 32(2), pp. 224-240.
- Hansen, Bruce. "Inference when a Nuisance Parameter is not Identified under the Null Hypothesis," *Econometrica*, March 1996, 64, pp. 413-430.
- McFadden, Daniel, "Conditional Logit Analysis of Qualitative Choice Behavior," in Paul Zarembka, ed., *Frontiers in Econometrics*, New York, Academic Press. 1974.
- McKelvey, Richard D. and Palfrey, Thomas R. "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, July 1995, 10(1), pp. 6-38.

- Milgrom, Paul and Roberts, John. *Economics, Organization and Management*, New Jersey, Prentice Hall, 1992.
- Mookherjee, Dilip and Sopher, Barry. "Learning Behavior in an Experimental Matching Pennies Game," *Games and Economic Behavior*, July 1994, 7(1), pp. 62-91.
- _____. "Learning and Decision Costs in Experimental Constant Sum Games. *Games and Economic Behavior*, April 1997, 19(1), pp. 97-132.
- Robinson, Julia. "An Iterative Method of Solving a Game," *Annals of Mathematics*, September 1951, 54(2), pp. 296-301.
- Roth, Alvin E. and Erev, Ido. "Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, January 1995, 8(1), pp. 164-212.
- Selten, Reinhard. "Features of Experimentally Observed Bounded Rationality," Presidential Address, European Economic Association, Toulouse. 1997.
- Stein, Jeremy. "Cheap Talk and the Fed: A Theory of Imprecise Policy Announcements," *American Economic Review*, March 1989, 79(1), pp. 32-42.
- Van Huyck, John B.; Battalio, Raymond C., and Rankin, Frederick W. "On the Origin of Convention: Evidence from Coordination Games," *Economic Journal*, May 1997, 107, pp. 576-596.
- Stinchcombe, Maxwell B. and White, Halbert, "Consistent Specification Testing with Nuisance Parameters Present Only under the Alternative," *Econometric Theory*, June 1998, 14(13), pp. 295-325.

Table 1. Payoffs of Games in Experiments

Panel (a)						
Actions						
Types	Game 1			Game 2		
	a ₁	a ₂	a ₁	a ₂	a ₃	
θ_1	0,0	700,700	0,0	700,700	400,400	
θ_2	700,700	0,0	700,700	0,0	400,400	

Panel (b)						
Actions						
Types	Game 3			Game 4		
	a ₁	a ₂	a ₃	a ₁	a ₂	a ₃
θ_1	0,0	200,700	400,400	0,0	200,500	400,400
θ_2	200,700	0,0	400,400	200,500	0,0	400,400

Table 2. Experimental Treatments

	Game	Session Data	Population History	Number of Periods
G1	Game 1	Session II	Senders	20
G1NH	Game 1	Session II	No history	20
G2	Game 2	Session II	Senders	20
G3	Game 3	Session II	Senders	40
G4	Game 4	Session II	Senders	40

Note: In the analysis of the experimental data only the first 20 periods are used for G3 and G4.

Table 3. Maximum Likelihood Estimates of SR Model

Model	ϕ_0	ϕ_1	lnL	Pseudo-R ²
G1R1 (N=108)	0.9456 (0.1538)	0.17403 (0.5652)	-47.80	.4253
G1R2 (N=108)	0.9172 (0.0181)	1.5872 (0.1708)	-34.40	.5864
G1R3 (N=108)	0.4736 (0.1267)	3.5681 (0.8870)	-46.93	.4358
G1NHR1 (N=108)	0.1851 (0.1190)	5.5810 (1.4533)	-57.30	.3112
G1NHR2 (N=108)	0.8273 (0.1004)	1.4381 (0.4030)	-63.59	.2355
G1NHR3 (N=108)	0.7205 (0.0763)	3.1341 (0.7084)	-45.65	.4511
G2R1 (N=108)	0.5426 (0.1535)	4.7313 (1.4677)	-20.60	.7523
G2R2 (N=108)	0.7280 (0.2150)	0.35060 (1.0627)	-19.71	.7631
G2R3 (N=108)	0.5124 (0.1518)	4.9880 (1.4986)	-20.50	.7536
G3R1 (N=138)	0.9851 (0.3504)	2.8820 (1.0404)	-28.95	.6519
G3R2 (N=138)	0.9491 (0.0947)	1.2066 (.3802)	-54.95	.3394
G3R3 (N=138)	0.8281 (0.1750)	1.0235 (0.4477)	-61.03	.2663
G4R1 (N=138)	1.6190 (0.2193)	-0.2497 (0.0408)	-72.38	.1299
G4R2 (N=138)	1.4765 (0.4789)	-.1514 (0.8147)	-72.18	.1322
G4R3 (N=138)	0.6861 (0.1056)	2.7581 (0.7122)	-34.43	.5860

Notes: N = 6 times the number of periods (20) minus 12.
Standard errors in parentheses beneath coefficient estimates.

Table 4. Maximum Likelihood Estimates of BBL Model

Model	β_0	β_1	lnL	Pseudo-R ²
G1R1 (N=108)	1.7284 (0.5678)	0.4638 (0.3163)	-36.41	.5622
G1R2 (N=108)	0.8849 (0.1213)	1.8064 (0.5685)	-33.74	.5944
G1R3 (N=108)	1.2281 (0.0976)	0.8545 (0.3665)	-62.39	.2499
G1NHR1 (N=108)	-0.9465 (0.3421)	-0.3521 (0.3924)	-82.64	.0065
G1NHR2 (N=108)	-0.8265 (0.5927)	-0.2220 (0.3790)	-83.04	.0017
G1NHR3 (N=108)	1.0711 (0.1475)	0.3866 (0.2022)	-67.83	.1833
G2R1 (N=108)	0.9873 (0.1532)	1.9969 (0.6476)	-25.84	.6894
G2R2 (N=108)	2.2454 (0.126)	3.7093 (8.951)	-15.81	.8099
G2R3 (N=108)	1.0200 (0.1551)	1.7909 (0.6039)	-30.97	.6277
G3R1 (N=138)	0.9814 (0.2108)	9.8583 (3.4109)	-28.79	.6539
G3R2 (N=138)	1.7229 (0.6331)	0.0666 (0.1855)	-77.79	.0648
G3R3 (N=138)	0.0000 (1.5019)	0.0000 (0.7414)	-83.17	.0000
G4R1 (N=138)	-1.4393 (0.5116)	0.0732 (0.2082)	-82.73	.0053
G4R2 (N=138)	-0.6046 (1.9623)	-0.4218 (1.2747)	-83.17	.0007
G4R3 (N=138)	1.0663 (0.1647)	0.4338 (0.3367)	-83.04	.0017

Notes: See Table 3.

Table 5. Tests of pooling for SR model

Game	$\chi^2(4)$	$\Pr(\chi^2 > x)$
Game 1 ^a	4.412	0.353
Game 2	0.288	0.991
Game 3	28.487	0.001
Game 4	22.857	0.001

^a Game 1 with history

Table 6. Maximum Likelihood Estimates of Hybrid Model

Model	γ_0	γ_1	γ_2	lnL	Pseudo-R ²
G1R1 (N=108)	0.8130 (0.3266)	1.1212 (0.3810)	1.6447 (1.2291)	36.0458	.5666
G1R2 (N=108)	0.7561 (0.1836)	1.2515 (0.6917)	0.9965 (0.6125)	31.9447	.6159
G1R3 (N=108)	0.3642 (0.1495)	1.1935 (0.7053)	3.8228 (1.0649)	42.3970	.4903
G1NHR1 (N=108)	0.2115 (0.1303)	-0.3655 (0.4755)	5.1816 (1.4836)	51.2332	.3841
G1NHR2 (N=108)	0.7881 (0.1394)	-0.0521 (0.1709)	1.5156 (0.4961)	63.8264	.2327
G1NHR3 (N=108)	0.6872 (0.1147)	0.3774 (0.2841)	2.7282 (0.6555)	44.6377	.4833
G2R1 (N=108)	0.4476 (0.1859)	1.7137 (0.9218)	4.0113 (1.4409)	17.7263	.7869
G2R2 ^a (N=108)	137.9155	309.4403	1.9873	12.3397	.8516
G2R3 ^a (N=108)	0.5430	2.1311	4.3204	17.3824	.7910
G3R1 (N=138)	4.6007 (7.2234)	10.2102 (5.7366)	0.6294 (1.1680)	25.6062	.6922
G3R2 (N=108)	0.9001 (0.1243)	-0.7904 (0.8868)	1.2961 (0.3973)	54.4720	.3451
G3R3 (N=108)	0.8643 (0.1771)	0.2473 (0.8698)	0.8550 (0.3894)	61.2352	.2638
G4R1 (N=108)	0.6502 (0.1547)	0.6751 (0.7011)	1.0102 (0.3152)	69.2994	.1669
G4R2 (N=108)	0.5190 (0.1449)	-0.4692 (1.2318)	1.4801 (0.3458)	66.4417	.2012
G4R3 (N=108)	0.6614 (0.1055)	-0.2407 (0.8962)	2.6258 (0.6044)	33.3678	.5988

Notes: See Table 3.

^a ML failed to converge; Hessian singular at maximum.

Table 7. 0.05 Symmetric Wald t Tests of Hybrid Model

Model	t statistic $H_0: \gamma_1 = 0$	t statistic $H_0: \gamma_2 = 0$	A $H_0: \gamma_1 = 0$ A $H_0: \gamma_2 = 0$	R $H_0: \gamma_1 = 0$ A $H_0: \gamma_2 = 0$	A $H_0: \gamma_1 = 0$ R $H_0: \gamma_2 = 0$	R $H_0: \gamma_1 = 0$ R $H_0: \gamma_2 = 0$
G1R1 (N=108)	2.94	1.34	0	1	0	0
G1R2 (N=108)	1.81	1.63	1	0	0	0
G1R3 (N=108)	1.69	3.59	0	0	1	0
G1NHR1 (N=108)	0.77	3.49	0	0	1	0
G1NHR2 (N=108)	0.30	3.05	0	0	1	0
G1NHR3 (N=108)	1.33	4.16	0	0	1	0
G2R1 (N=108)	1.86	2.78	0	0	1	0
G2R2 (N=108)						
G2R3 (N=108)						
G3R1 (N=138)	1.78	0.54	1	0	0	0
G3R2 (N=108)	0.89	3.26	0	0	1	0
G3R3 (N=108)	0.28	2.20	0	0	1	0
G4R1 (N=108)	0.96	3.21	0	0	1	0
G4R2 (N=108)	0.38	4.28	0	0	1	0
G4R3 (N=108)	0.27	4.34	0	0	1	0
Total			2	1	10	0

Notes: A denotes acceptance of H_0 using a nominal 0.05 symmetric asymptotic test and R the rejection of H_0 using this test. The critical values are ± 2.0 , and the nominal probability of simultaneously accepting both hypotheses when they are true is asymptotically at least 0.90 when both hypotheses are true.




Figure 1. Plots of the actual (jagged line) and the predicted fraction (smooth line) of players sending message 2 by type (Δ = type 1, 0 = type 2) when the SR model is true




Figure 2. Plots of the actual (jagged line) and the predicted fraction (smooth line) of players sending message 2 by type (Δ = type 1, 0 = type 2) when the BBL model is true




Figure 3. Log likelihood function for SR model using GIR1 data




Figure 4. Log likelihood function for BBL model using G1R1 data


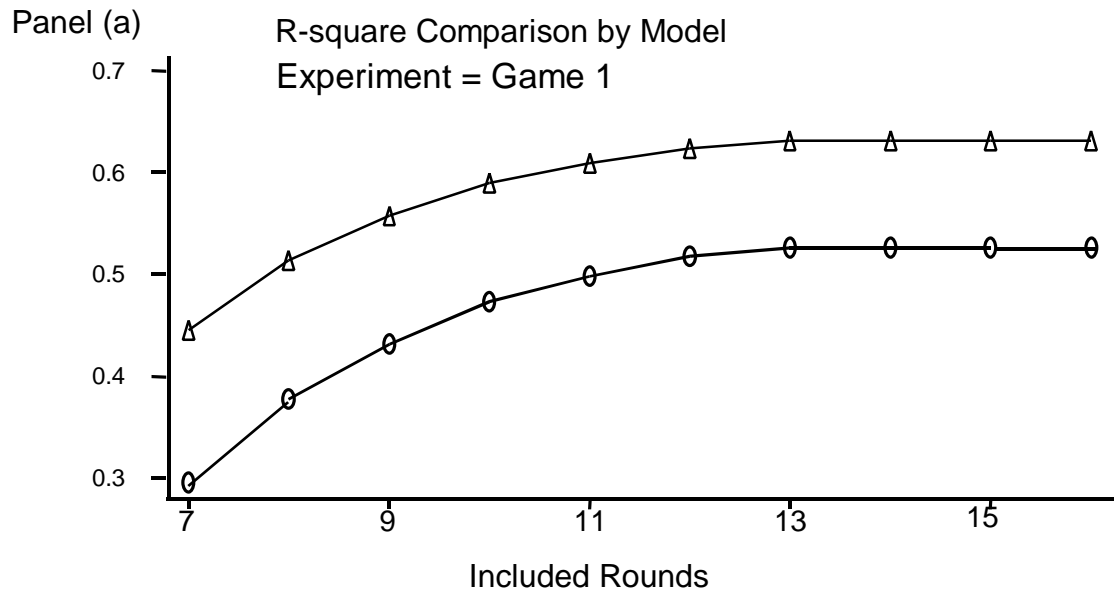


Figure 5. Convergence bias: Δ --BBL; 0--SR

ENDNOTES

¹ For example, See David Austin-Smith (1990), Vincent P. Crawford and Joel Sobel (1982), Joseph Farrell and Robert Gibbons (1990), Frank Gigler (1994) and Jeremy Stein (1989).

² All replications had a common session, which preceded the games described above. In particular, each cohort participated in 20 periods of a game with payoffs as in Game 1 and a message space of $M = \{A, B\}$. The common session provides players with experience about experimental procedures and ensures that players understand the structure of the game, message space and population history.

³ The specification of the logit function in (2) exploits the fact that all rewards, X , in the games that we examine are non-negative. Were this not the case, a transform that keeps the value of the payoffs non-negative, such as the exponential function, can be used.

⁴ The coefficients in the hybrid model can be normalized so that the coefficients on the information variables are δ and $(1 - \delta)$, respectively, which makes the hybrid model look similar to the model employed Camerer and Ho (1999).

⁵ Defining convergence for mixed strategies is conceptually the same as the pure strategy case; empirically identifying convergence is more difficult.

⁶ The number of studies is growing at an increasing rate. We have selected representatives from the set and apologize for any omissions.