

An Experiment to Evaluate Bayesian Learning of Nash Equilibrium Play

James C. Cox*

Department of Economics, University of Arizona, Tucson, Arizona 85721

Jason Shachat*

*Department of Economics, 9500 Gilman Drive, University of California,
San Diego, La Jolla, CA 92093-0508*

and

Mark Walker*

Department of Economics, University of Arizona, Tucson, Arizona 85721

Received March 24, 1997

We report on an experiment designed to evaluate the empirical implications of Jordan's model of Bayesian learning in games of incomplete information. A finite example is constructed in which the model generates unique predictions of subjects' choices in nearly all periods. When the "true" game defined by players' private information was one with a unique equilibrium in pure strategies, the experimental subjects' play converged to the equilibrium, as Jordan's theory predicts, even when the subjects had not attained complete information about one another. But when there were two pure strategy equilibria, the theory's predictions were not consistent with observed behavior. *Journal of Economic Literature* Classification numbers: D83, C72, C92. © 2001 Academic Press

The attempt to rationalize equilibrium analysis in games has largely shifted in recent years from arguments based on players' introspection and common knowledge to the idea that equilibrium play is *learned* by the players of a game through repeated play. Parallel to this theoretical research on learning in games, a growing body of experimental research has begun to investigate alternative models of the ways that players might

* E-mail: jcox@bpa.arizona.edu, jshachat@weber.ucsd.edu, mwalker@arizona.edu.

adjust their behavior as they play many repetitions of a game.¹ The learning models one encounters in this experimental research typically attribute quite limited rationality to the players, as in the best response, fictitious play, and reinforcement learning models. In this paper we undertake an experimental investigation of a learning model that lies at the opposite end of the rationality spectrum, the Bayesian learning model introduced by Jordan (1991), which attributes full Bayesian rationality to the players. Our experimental results suggest that behavior may be consistent with the Bayesian model when there is a unique equilibrium in pure strategies, but that behavior is not consistent with the model when there are multiple equilibria, or when the only equilibrium is in mixed strategies.

In Jordan's model, players are engaged in a noncooperative normal form game. Each player knows his own payoff function but he has only a probabilistic belief about the other players' payoff functions. In this game of incomplete information, the players are assumed to play a Bayesian Nash equilibrium, but this will not generally yield a Nash equilibrium (NE) of the "true game" they are playing—the game defined by the players' *actual* payoff functions. The question Jordan addresses is whether, by playing the game repeatedly, the players will learn to play a Nash equilibrium of the true game. Jordan assumes that at each stage of play the players play a Bayesian Nash equilibrium based upon their current beliefs, and that after each play has occurred each player updates his belief according to Bayes' Rule by incorporating the play he has just observed. Jordan proves that under certain restrictions on beliefs this process does indeed converge to a Nash equilibrium of the true game.²

In order to evaluate the empirical implications of Jordan's Bayesian learning model, we have devised and conducted an experiment based on the elegant example in Jordan (1991). We convert Jordan's example, which has a continuum of player types, to an example with a very small finite set of player types. This provides us with a small but diverse collection of 2×2 games: there are games that have a unique equilibrium which is in pure strategies; a coordination game with opposing interests (essentially, a Battle of the Sexes game); a game in which one equilibrium Pareto dominates the others; and two games that have a unique equilibrium in mixed strategies.

¹For example, Camerer and Ho (1996); Cheung and Friedman (1997); Roth and Erev (1995).

²A further interesting feature of Jordan's analysis is that while the players do eventually play a complete-information Nash equilibrium, in most cases they do so without actually becoming completely informed: each player simply learns enough about the others to lead him to play his part of the complete-information Nash equilibrium.

We find that when the true game the players are playing has a unique, pure-strategy Nash equilibrium, play generally moves to that equilibrium. Perhaps surprisingly (but as Jordan's model predicts), this convergence occurs even in cases in which the players clearly have not attained complete information about one another's types. However, in more difficult situations—when there are multiple equilibria or mixed strategy equilibria—we find that the behavior of the experimental subjects is not generally consistent with the theory. Based upon the prior beliefs we induce in our subjects, the Jordan theory predicts *which* equilibrium will be attained in the games with multiple equilibria, but we find that play generally does not converge to the predicted equilibrium.

Jordan's theory assumes that the players' behavior is myopic—that each player maximizes his current expected payoff, given his beliefs, without looking forward to future plays of the game. In order to assess the role of this assumption, we have conducted some of the experiments described above under two different protocols for matching subjects against one another. We first conducted experiments in which subjects were randomly rematched after each play, in an attempt to *induce* myopic behavior; and then we conducted the same experiments without rematching the subjects at all—i.e., each pair of subjects played against one another repeatedly. The form of the matching was found to have an effect: when subjects were rematched with a new opponent each period, they generally did not converge to any of the game's equilibria; when the subjects competed against one another in fixed pairs, they generally converged to a pure-strategy equilibrium, but not significantly more often to the one predicted by Jordan's model than to the other pure-strategy equilibrium.

The remainder of the paper is organized as follows. We describe Jordan's model and the experiment based upon it in Section 1. In Section 2 we describe the experimental procedures and some of the considerations behind the specific protocols we have adopted. The experimental results are reported in Section 3, and Section 4 presents concluding remarks and some implications of our results for further research.

1. AN EXPERIMENT BASED ON JORDAN'S THEORY

We begin with a brief description of Jordan's model of Bayesian learning in games. Although Jordan develops the theory for n -player games with arbitrary finite strategy sets, we describe the model for two-player 2×2 games (two strategies for each player), which is the framework for his example and for our experiment. Our convention throughout the paper will be to refer to the two players as the Row player and the Column player,

with strategy sets $\{\text{Top, Bottom}\}$ and $\{\text{Left, Right}\}$, or simply $\{\text{T, B}\}$ and $\{\text{L, R}\}$.

Jordan's model assumes that each player i has a payoff function π_i that gives his payoff at each of the four possible profiles of play that can occur, and that each player has a prior belief about his opponent's possible payoff function, or "type" (i.e., he has a probability distribution over his opponent's *possible* types). The two players play the game repeatedly, with their given, unchanging payoff functions, in each period playing a Bayesian Nash equilibrium of the stage game for their current probability beliefs about their opponents' types. (Thus, the players are assumed to play the game myopically, always attempting to maximize only the current period's payoff.) After each play, each player updates his belief according to Bayes' Rule, incorporating the play by his opponent that he has just observed. Jordan refers to a path of play that satisfies the conditions we have just described as a *Bayesian strategy process*, and he shows that for all finite games (except possibly for a set with prior probability zero) every Bayesian strategy process converges to a Nash equilibrium of the "true" game, that is, of the game defined by the players' actual payoff functions.

In Jordan's example the players' types (their payoff tables) are drawn from a continuum of possible payoff tables. In an experiment, however, the set of types must be finite. Indeed, in the interest of simplicity (in the theory, and for the experimental subjects), the number of types for each player should be kept as small as possible. An additional reason for keeping the number of types small is to minimize the total number of draws (and therefore also time and subject payments) required in order to generate sufficient sample sizes for each possible draw of a type for each player. We have therefore developed a simplified finite analogue of Jordan's example, in which Row's payoff table is drawn from a set \mathcal{T}_1 containing only four possible types, or payoff tables, and Column's payoff table is drawn from a set \mathcal{T}_2 containing only two possible types. The specific payoff tables, denoted \mathcal{A} , \mathcal{B} , \mathcal{C} , and \mathcal{D} which comprise the sets \mathcal{T}_1 and \mathcal{T}_2 are given in Fig. 1. Note that Column's two possible payoff functions, \mathcal{B} and \mathcal{D} , are two of the four that are possible for Row, so that when each player draws \mathcal{B} or each draws \mathcal{D} , the resulting true game is symmetric. Finally, the probabilities with which the payoff tables are drawn are also given in Fig. 1. There are thus eight possible 2×2 true games.

Table I provides a concise summary of the path of play, according to Jordan's theory, assuming that the players' initial beliefs are the probability distributions over types given in Fig. 1. Note in particular in Table I that in every one of the eight games that can be drawn from $\mathcal{T}_1 \times \mathcal{T}_2$ the theory makes a precise, unique prediction about the play that will occur in every period, with the sole exception of periods three and later in the mixed-equilibrium-only games \mathcal{B} vs \mathcal{D} and \mathcal{D} vs \mathcal{B} . The action pairs in the " $t = 1$ "




FIG. 1. Induced prior beliefs (types and probabilities).

column of Table I constitute the unique pure-strategy Bayesian Nash equilibrium of the incomplete-information game at period 1. Note that this Bayesian Nash equilibrium separates the Column player's two types (i.e., his play at period 1 reveals his type), and that it separates the Row player's types into the two sets $\{A, D\}$ and $\{B, C\}$. Then, at period 2, there is again (conditional on each of the four possible plays at $t = 1$) a unique Bayesian Nash equilibrium, which is in pure strategies, for each of the four incomplete-information games. This Bayesian Nash equilibrium reveals the

TABLE I
Equilibria and Predicted Path for Each Game

True game	Nash equilibria	Path of play predicted by Jordan model		
		$t = 1$	$t = 2$	$t = 3, 4, 5, \dots$
A vs B	(T,R)	(T,R)	(T,R)	(T,R)...
A vs D	(T,L)	(T,L)	(T,L)	(T,L)...
C vs B	(B,L)	(B,R)	(B,L)	(B,L)...
C vs D	(B,R)	(B,L)	(B,R)	(B,R)...
D vs B	$(1/3, 1/3)^a$	(T,R)	(B,R)	$(1/3, 1/3)^a \dots$
D vs D	(T,L), (B,R), $(1/3, 1/3)^a$	(T,L)	(T,L)	(T,L)...
B vs B	(B,L), (T,R), $(1/3, 1/3)^a$	(B,R)	(B,L)	(B,L)...
B vs D	$(1/3, 1/3)^a$	(B,L)	(T,R)	$(1/3, 1/3)^a \dots$

^aA fraction denotes a mixed-strategy equilibrium: Row's fraction is the probability with which he plays Top; Column's fraction is the probability with which he plays Left.

Row player's type (followed by Nash equilibrium play at all subsequent t), or else it is "pooling," in which case the Row player makes the same play whatever his type (followed by the same behavior at all subsequent periods, since the Column player learns nothing from it).

Note that in every one of the eight games in $\mathcal{F}_1 \times \mathcal{F}_2$, the Column player's Nash equilibrium play depends upon the Row player's type. Moreover, in half of the eight games the Column player, according to the theory, never becomes fully informed of the Row player's type. The players can therefore be said to have "learned" to play Nash equilibrium in each of the games, even though in some of the games they play Nash equilibria in every period.

2. EXPERIMENTAL PROCEDURES

The experiment consisted of 10 sessions of about 2 hours each, conducted on the networked computer facility at the University of Arizona's Economic Science Laboratory. There were 12 subjects in each session, 6 of whom played the role of Row player throughout, and 6 of whom always played Column. No subject participated in more than 1 session. The entries in the games' payoff matrices were converted into U.S. dollar rewards at the rate of one U.S. dollar for every 10 payoff units in Sessions 2, 3, and 4 and at a rate of one U.S. dollar for every 6 payoff units in the other sessions. A summary of the salient rewards the subjects actually earned in each session is provided in Table II. Subjects were paid an additional \$5 "show-up fee" upon signing in for the experiment.

Each experimental session consisted of several "regimes." Each regime began with a random drawing of a Row payoff table (revealed only to the Row players, by onscreen highlighting of that table, the same table

TABLE II
Salient Payoffs

	Low payoff (\$)	Mean payoff (\$)	High payoff (\$)
Session 1	17.67	21.99	26.00
Session 2	15.90	16.96	18.20
Session 3	12.70	15.73	19.10
Session 4	11.10	12.95	13.90
Session 5	17.33	19.14	20.83
Session 6	18.00	19.83	21.67
Session 7	18.00	20.47	22.67
Session 8	19.33	23.49	27.67
Session 9	19.33	24.89	28.33
Session 10	20.00	24.54	28.67

TABLE III
Experimental Regimes (Row Player Type vs Column Player Type)

Session 1	<i>C vs B</i>	<i>A vs B</i>	<i>D vs D</i>	<i>A vs D</i>		Rematched
Session 2	<i>A vs B</i>	<i>D vs D</i>	<i>B vs D</i>	<i>C vs B</i>	<i>D vs D</i>	Rematched
Session 3	<i>A vs D</i>	<i>C vs B</i>	<i>C vs B</i>	<i>A vs B</i>	<i>A vs B</i>	Rematched
Session 4	<i>B vs B</i>	<i>D vs B</i>	<i>C vs D</i>	<i>D vs D</i>	<i>B vs B</i>	Rematched
Session 5	<i>B vs B</i>	<i>D vs B</i>	<i>C vs D</i>	<i>D vs D</i>	<i>B vs B</i>	Rematched
Session 6	<i>B vs B</i>	<i>D vs B</i>	<i>C vs D</i>	<i>D vs D</i>	<i>B vs B</i>	Rematched
Session 7	<i>B vs B</i>	<i>D vs B</i>	<i>C vs D</i>	<i>D vs D</i>	<i>B vs B</i>	Rematched
Session 8	<i>B vs B</i>	<i>D vs B</i>	<i>C vs D</i>	<i>D vs D</i>	<i>B vs B</i>	Fixed pairs
Session 9	<i>B vs B</i>	<i>D vs B</i>	<i>C vs D</i>	<i>D vs D</i>	<i>B vs B</i>	Fixed pairs
Session 10	<i>B vs B</i>	<i>D vs B</i>	<i>C vs D</i>	<i>D vs D</i>	<i>B vs B</i>	Fixed pairs

for each Row player), and a random drawing of a Column payoff table (revealed only to the Column players, and the same table for each one). The resulting game was played for 15 periods, and then a new regime was begun with the drawing of new payoff tables. Table III indicates the sequence of regimes—the sequence of true games that were drawn—in each of the 10 experimental sessions. The games were drawn according to the probabilities that appear in Fig. 1. A version of Fig. 1 was present on each subject's computer screen, and it was emphasized to the subjects that the players' types would be drawn according to the probabilities shown onscreen under the associated payoff tables. These probabilities are therefore assumed to constitute the subjects' prior beliefs about one another's types.

After each period's play was complete, each subject was informed of the action his opponent had chosen (but no other subjects' actions) and his own resulting payoff. Each subject had the ability at any time to bring back onto his computer screen all the information he had ever been given, viz., his own and his opponents' past choices and his own past payoffs. Subjects were not allowed to communicate with one another. Note that under these information conditions, a subject who is playing repeatedly against a fixed opponent will have available to him at period t all the actions his opponent has chosen prior to t .

Jordan's theory assumes that a fixed pair of players is playing the same true game repeatedly. Some of our experimental sessions were therefore conducted in that way: subjects were matched into fixed pairs to play all 15 periods of a regime, and were then rematched at the beginning of the next regime. But Jordan's theory also assumes that the players play myopically, ignoring any effect their current choice might have upon future play. It is not clear whether actual players will in fact play myopically in this situation. Consequently, the experimental sessions conducted with fixed pairs in each regime must be regarded as evaluating the *joint* hypothesis of myopic play and Jordan's Bayesian learning model.

An alternative to testing this joint hypothesis is to attempt to control for the part of the joint hypothesis which is not of central interest, in our case the assumption of myopic play. A technique often used in experimental research to induce myopic play is to randomly rematch subjects prior to each new interaction, thereby eliminating, or at least minimizing, any repeated game effects. Some of our experimental sessions employed this technique, randomly rematching subjects against one another after every play instead of after every regime. While rematching subjects in this way ought to minimize the influence of repeated game effects, it also introduces a feature that is not present in Jordan's model by altering the experiment's information structure: instead of knowing, at any play, what his *current* opponent has done in all preceding plays of the current regime, a subject instead knows only the sequence of choices he has seen by the randomly selected subjects he has played against (all of whom have the same pay-off table as his current opponent). A subject now has to draw inferences from information about the population of his *potential* opponents, instead of from information about his *actual* opponent. Thus, in the sessions with rematching after each play, we are again testing a joint hypothesis: the hypothesis that subjects will indeed play myopically when rematched in this way, together with Jordan's model of Bayesian learning.³

If every subject plays precisely according to the theory, the matching protocol will have no effect (so long as the theory makes a unique prediction, as it generally does in our experiment; see Table I). Recognizing, however, that actual players will not always play exactly as the theory predicts, an important objective of our experiment is to determine whether Jordan's idealized model still makes generally accurate predictions about the path of play and about the behavior to which play converges when players' actions sometimes depart from the theory. When some play deviates in this way, the information structure may matter a great deal.

Thus, our experiment was divided into random rematching sessions and fixed-pairs sessions. In our first seven sessions we randomly rematched Row and Column opponents in every period. The subjects were informed of the random matching protocol with nontechnical wording in the experiment's instructions. The most interesting phenomena to appear in the first four sessions occurred in Session 4, in which three of the regimes involved a multiple-equilibrium game, either \mathcal{B} vs \mathcal{B} or \mathcal{D} vs \mathcal{D} , games in which the path of play is likely to be more sensitive to the informational differences in the matching protocols. The remaining six sessions were therefore focused

³Our random matching procedure reflects one of the two classes of naturally occurring, non-experimental settings in which players are indeed likely to play myopically: when interactions occur between players who meet randomly or when the number of players whose interactions affect one another is quite large.

upon the sequence of games in Session 4, with random rematching continued in Sessions 5, 6, and 7, and with fixed pairs matched against one another for the duration of a regime in Sessions 8, 9, and 10.

3. ANALYSIS OF THE DATA

We first try to determine whether our experimental subjects played the particular equilibrium that the Jordan model predicts. We find that for all but the most basic cases they did not. However, we also find that the alternative matching protocols made a substantial difference in whether play converged to *any* pure strategy equilibrium. We then build upon this observation to suggest a possible source of the Jordan model's failure: the players' information in the first few periods of a regime may not have been reliable enough to support the actions predicted by the theory. Although most play in the first two periods of each regime was in agreement with the theory, there was some play in these early periods that deviated from the theoretical predictions, and this deviant play may have been enough to significantly influence subsequent play in a regime, and in particular it may have been enough to lead to *different* play under the two alternative information conditions induced in the experiment.

3.1. *Learning the Nash Equilibrium*

In our experiment, Jordan's learning model predicts that all play will eventually be at a *particular* Nash equilibrium of whichever true game is being played; further, that this profile of play (or pair of mixtures) will be reached no later than period 3; and that once reached, play will remain at this equilibrium profile or mixture-pair in all subsequent plays during that regime. A first look at the data's consistency with this prediction is provided in Table IV, which reports, for periods three and later in each regime, the frequencies with which joint plays, or "profiles of play," fall into each of the four game cells. For each of the six true games that have pure strategy equilibria, the Jordan model makes a unique prediction about the profile that will be played, and that profile is indicated in Table IV as a boldface frequency.

We consider first the four games with a unique equilibrium in pure strategies, then the two games with multiple equilibria, and finally the two games with a unique equilibrium in mixed strategies.

Unique Equilibrium in Pure Strategies. The four games that have a unique pure strategy equilibrium, (A vs B , A vs D , C vs B , and C vs D), appear in the left half of Table IV. The frequency of play in the respective Nash equilibrium cells for each of these games was extremely high. This is

TABLE IV
Cell Frequencies for Periods Three and Later

Random matching sessions (%)								
Cell	A vs B	A vs D	C vs B	C vs D	B vs D	D vs D	B vs D	D vs B
(T,L)	3	91	0	0	17	63	14	15
(T,R)	97	5	0	0	18	14	23	32
(B,L)	0	5	96	9	29	15	32	16
(B,R)	1	0	4	91	36	9	31	37
# of Obs	312	156	312	312	624	546	78	312
Fixed pair sessions (%)								
(T,L)	-	-	-	1	10	77	-	13
(T,R)	-	-	-	0	35	3	-	23
(B,L)	-	-	-	9	45	4	-	31
(B,R)	-	-	-	90	10	16	-	33
# of Obs	-	-	-	234	468	234	-	234

Note. Boldface indicates a unique prediction by Jordan model.

not surprising, since the Row player has a dominant strategy in each game (as at least one of the players must in a 2×2 game with a unique pure strategy equilibrium). But this empirical result is perhaps slightly more informative than it seems: the Column player does *not* have a dominant strategy. The Column player's equilibrium action depends upon the Row player's type, which the Column player does not know at the outset of a regime, and which he never observes directly. It is therefore appropriate to say, in light of the extremely high frequencies of equilibrium play, that the players have indeed "learned" to play the NE in each of these games.

Multiple Equilibria. Each of the two games with multiple equilibria, \mathcal{B} vs \mathcal{B} and \mathcal{D} vs \mathcal{D} , has two pure strategy equilibria as well as a third equilibrium in mixed strategies. \mathcal{B} vs \mathcal{B} is a symmetric coordination game in which the players have opposing preferences for the two pure-strategy equilibria, i.e., a Battle of the Sexes game, and the Jordan model predicts that under the beliefs induced in our experiment the players will learn to play the profile (B,L). In \mathcal{D} vs \mathcal{D} the equilibrium profile (T,L) strictly Pareto dominates all other distributions of play over the four profiles, including the other pure-strategy equilibrium, (B,R); the Jordan model predicts that under our induced beliefs the players will learn to play (T,L).

In the eight rematching regimes in which \mathcal{B} vs \mathcal{B} occurred, Table IV reveals that the equilibrium predicted by the Jordan model, (B,L), was played on only 29% of the plays, and that the non-equilibrium profile

(B,R) was played more often, on 36% of the plays. In the seven rematching regimes in which \mathcal{D} vs \mathcal{D} occurred, the equilibrium (T,L) predicted by the Jordan model was played 63% of the time, a figure which also seems surprisingly low in light of the strong Pareto dominance of this equilibrium.

Table IV suggests that when the pairs were not rematched, but remained fixed during regimes, the Jordan model fared only slightly better. Here, play was at (B,L) 45% of the time in \mathcal{B} vs \mathcal{B} and was at (T,L) 77% of the time in \mathcal{D} vs \mathcal{D} . However, if we focus our attention on *pure strategy equilibrium* predictions, instead of the *specific* equilibrium selected by the Jordan model, Table IV suggests that the alternative information conditions induced by our two matching protocols had a significant effect upon play. In \mathcal{B} vs \mathcal{B} under rematching, fewer than half of the plays, only 47%, were at one of the two pure-strategy equilibria. But with fixed pairs, 80% of the plays were at one of these equilibria. In \mathcal{D} vs \mathcal{D} , only 63% of plays were at the Pareto dominant equilibrium, and 72% were at one of the two pure strategy equilibria. With fixed pairs, 77% of plays were at the Pareto dominant equilibrium and 93% of plays were at one of the two pure-strategy equilibria.

Table IV presents only the aggregate frequencies of play, aggregated over periods three and later and over all regimes for each true game. Tables V–VIII present period-by-period data for each \mathcal{B} vs \mathcal{B} and \mathcal{D} vs \mathcal{D} regime. This disaggregated view of the data reveals some important and systematic features of the subjects' play that are obscured by aggregation. In each of these tables the entries for both of the pure-strategy equilibria ((B,L) and (T,R) in \mathcal{B} vs \mathcal{B} ; (T,L) and (B,R) in \mathcal{D} vs \mathcal{D}) appear either in *italic* (if predicted by the Jordan model) or in **boldface** (if not the Jordan model's predicted outcome).

Table V presents the time series of cell (or profile) counts for every play of the \mathcal{B} vs \mathcal{B} game in the eight rematching regimes in which it occurred. In two of the regimes (Regime 5 of Sessions 4 and 7), one can see that play was predominantly at the Jordan equilibrium (B,L). In each of the remaining six regimes, however, play was more often at the non-equilibrium profile (B,R), and in none of these six regimes were the proportions of play in the four cells even remotely close to any of the three equilibrium distributions. We conclude that in the random rematching regimes, designed to induce myopic play, not only is there little support for the Jordan model's prediction, but there is no evidence that players are likely to learn to play any of the equilibria.

Things were quite different when pairs were not rematched but remained fixed, as Table VI reveals. Table VI presents the complete time series of play for every occurrence of the game \mathcal{B} vs \mathcal{B} in a fixed-pair session. There were three such sessions, and \mathcal{B} vs \mathcal{B} occurred in two regimes in each

TABLE V
Time Series of Cell Counts in Random Matching **B** vs **B** Regimes

Period	Session 5															
	Regime 1				Regime 5				Regime 1				Regime 5			
	TL	TR	BL	BR	TL	TR	BL	BR	TL	TR	BL	BR	TL	TR	BL	BR
1	1	1	2	2	1	2	1	2	1	0	2	3	0	2	1	3
2	1	1	0	4	1	2	0	3	2	0	1	3	2	1	1	2
3	1	1	3	1	2	1	2	1	1	1	2	2	2	1	1	2
4	1	1	2	2	0	1	3	2	2	1	2	1	2	1	0	3
5	0	1	3	2	2	0	2	2	2	1	1	1	3	0	2	2
6	4	1	1	0	1	0	2	3	1	2	2	1	0	1	3	2
7	1	2	0	3	1	0	3	2	3	1	0	2	1	0	2	3
8	2	0	1	3	1	0	5	0	1	1	1	3	2	1	0	3
9	0	0	1	5	0	1	5	0	1	1	3	1	2	1	1	2
10	0	1	2	3	1	0	4	1	1	1	2	2	2	0	0	4
11	1	2	1	2	2	0	3	1	1	2	1	2	0	2	2	2
12	2	1	1	2	1	0	4	1	2	0	0	4	0	2	2	2
13	0	3	1	2	1	0	5	0	1	0	2	3	1	2	2	1
14	2	1	0	3	2	0	4	0	0	2	2	2	3	0	1	2
15	2	2	1	1	1	0	4	1	0	3	1	2	1	1	1	3

Session 7

Session 6

Period	Regime 1				Regime 5				Period	Regime 1				Regime 5			
	TL	TR	BL	BR	TL	TR	BL	BR		TL	TR	BL	BR	TL	TR	BL	BR
1	0	2	0	4	0	1	1	4	1	1	0	4	0	2	3	1	
2	1	2	0	3	1	1	1	3	1	1	2	2	2	2	2	0	
3	1	3	0	2	1	1	0	4	3	0	2	1	0	1	3	2	
4	1	2	0	3	1	1	1	3	4	0	3	3	1	0	2	3	
5	1	3	0	2	0	3	0	3	5	0	1	2	1	0	4	1	
6	0	3	1	2	1	1	0	4	6	2	1	3	2	1	3	0	
7	3	2	0	1	1	2	1	2	7	2	1	3	2	0	1	3	
8	0	2	2	2	0	3	1	2	8	1	0	3	1	0	3	2	
9	2	1	0	3	1	2	1	2	9	2	1	2	1	1	2	2	
10	2	1	1	2	0	2	2	2	10	1	1	3	0	0	3	3	
11	0	1	1	4	0	3	1	2	11	0	1	3	0	0	3	3	
12	1	1	0	4	1	3	1	1	12	1	1	2	2	0	3	1	
13	1	1	1	3	1	3	0	2	13	1	1	0	4	0	5	1	
14	0	1	2	3	0	3	1	2	14	1	0	3	2	0	5	1	
15	1	1	2	2	1	2	0	3	15	0	2	2	2	0	3	1	

Note: Jordan equilibrium; BL, Non-Jordan equilibrium; TR.

of the sessions. Table VI thus contains 36 time series of play: 18 subject-pairs, each of whom played \mathcal{B} vs \mathcal{B} in two regimes. In most of these 36 regimes, play converged quickly to one of the two pure-strategy equilibria: in 15 of the regimes the pair of subjects played the Jordan profile (B,L) in nearly every period, and in another 9 regimes the subjects played the other pure-strategy equilibrium, (T,R), in nearly every period. (Moreover, play in perhaps as many as 5 additional regimes in Table VI appears to have been converging to the profile (T,R) in the last few periods.) We conclude that when a fixed pair plays the Battle of the Sexes game \mathcal{B} vs \mathcal{B} under these beliefs, they will generally reach one of the pure-strategy equilibria, and will generally reach it quickly. But there is no evidence that they are more likely to reach the equilibrium predicted by the Jordan model than the other equilibrium.

The other multiple equilibrium game, \mathcal{D} vs \mathcal{D} , occurred in seven rematching regimes and in three fixed-pairs regimes. Table VII reports the cell counts for each of the random rematching regimes. In three of the regimes play clearly converged quickly to the Pareto dominant profile (T,L), and in a fourth (Regime 4 of Session 4) play was very close to complete convergence to (T,L). In the remaining three regimes, there is no evidence of convergence to any of the three equilibria.

When pairs were fixed, however, Table VIII indicates that their play of \mathcal{D} vs \mathcal{D} converged far more consistently to the Pareto dominant equilibrium (T,L): only 4 of the 18 pairs did not converge quickly and completely to (T,L). Three of those 4 pairs converged, just as quickly and just as completely, to the other, Pareto dominated, pure-strategy equilibrium, (B,R). Thus, in \mathcal{D} vs \mathcal{D} , just as in the other multiple equilibrium game, \mathcal{B} vs \mathcal{B} , the different information conditions induced by our two matching protocols again appear to have had a dramatic effect on whether the experimental subjects reached equilibrium or not.

Unique Equilibrium in Mixed Strategies. The two remaining games, \mathcal{B} vs \mathcal{D} and \mathcal{D} vs \mathcal{B} , each has a unique equilibrium in mixed strategies. The equilibrium is the same in both games: Row plays Top with probability one-third and Column plays Left with probability one-third. Thus, mixed-strategy play in either game will typically yield observed cell frequencies that are approximately as follows: 11% for (T,L), 22% each for (T,R) and (B,L), and 44% for (B,R). All four profiles of play were indeed played with considerable frequency in these games, as described in Table IV, and the observed distributions of play are not dramatically different from the predicted distribution (0.11, 0.22, 0.22, 0.44). However, as Table IX describes, goodness-of-fit tests reject the hypothesis that the data were generated by this multinomial distribution. These results are consistent with both the the-

TABLE VI
Time Series of Each Pair's Play in Fixed-Pair **B** vs **B** Regimes

Pair	regime	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	1	BR	TR	TR	TR	TR	TR	TR	TR	TR	TL	BL	BL	TL	BR	TR
	5	TL	BL	BL	TL	BL	BL	BL	BL	BL	TL	BL	BL	BL	BL	BL
2	1	BR	BL	TL	BR	TL	BR	BR	BR	TR	TL	BL	TR	BL	TR	BR
	5	TR	TL	TR	BR	TR	BL	TR	TL	TR	BR	TR	BR	TR	BL	TR
3	1	TL	BR	BR	BR	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL
	5	BR	BR	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL
4	1	BR	TL	BR	BL	BR	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL
	5	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR
5	1	BL	BL	TL	BL	BL	TL	TL	TL	TR	TL	TR	TR	TR	TR	TL
	5	TL	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR
6	1	BR	BR	TR	BL	TR	BL	TL	BR	TR	BR	TR	BL	TL	BR	TR
	5	BR	BL	TL	BL	TL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL
7	1	BR	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL
	5	BR	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL
8	1	BR	BR	BR	TR	TL	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR
	5	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL
9	1	BR	BL	BL	TL	TR	BR	BL	BL	BL	BL	BL	BL	BL	BL	BL
	5	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL
10	1	BR	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL
	5	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR
11	1	BL	BR	BL	BL	TL	BR	BL	BL	BL	BL	TL	TR	TR	TR	TR
	5	TR	TR	BR	TL	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR
12	1	BR	BR	BR	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL
	5	BR	TR	TL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL
13	1	BR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR
	5	TL	BR	BL	BL	BL	TL	TL	BL	BL	BL	BL	BL	BL	BL	BL
14	1	BR	BR	TL	BR	BR	BR	BR	BR	BL	BL	BL	BL	BL	BL	TL
	5	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL	BL
15	1	BR	TR	TL	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR
	5	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR
16	1	TR	TL	TR	TR	TR	TR	TR	TR	TR	TR	TL	TL	BR	TL	TR
	5	TR	TR	TR	TR	TR	TR	TR	TL	TL	TR	TR	TR	TL	TL	TL
17	1	TL	TL	BR	BR	BR	TL	BR	BR	BR	BR	TL	BR	BR	BR	BR
	5	BR	BR	BL	TL	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR
18	1	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR	TR
	5	TR	TR	TR	TL	BL	BL	TL	BL	BL	BL	BL	TL	TL	TL	TL

Note. Jordan equilibrium, *BL*; Non-Jordan equilibrium, **TR**.

oretical and the experimental literature, which find serious difficulties for convergence to mixed strategy-equilibria.⁴

⁴For example, Crawford (1985), Fudenberg and Kreps (1993), and Jordan (1993) in the theoretical literature; Mookherjee and Sopher (1994), Ochs (1995), Bloomfield (1994), and Erev and Roth (1998) in the experimental literature.

TABLE VII
Cell Counts in Random Matching **D** vs **D** regimes

Period	Session 1			Session 2			Session 3			Session 4			Session 5			Session 6			Session 7							
	TL	TR	BL	BR	TL	TR	BL	BR	TL	TR	BL	BR	TL	TR	BL	BR	TL	TR	BL	BR	TL	TR	BL	BR		
1	1	2	2	1	5	1	0	0	2	1	2	1	2	1	3	0	2	2	2	1	2	1	3	2	0	
2	3	2	0	1	4	1	1	2	0	3	1	2	0	1	3	0	3	2	2	1	1	2	2	2	0	
3	2	2	2	0	4	1	1	0	5	0	0	1	1	4	1	0	2	2	1	2	3	0	3	1	2	0
4	4	0	2	0	6	0	0	0	5	1	0	0	5	0	1	0	3	3	1	0	2	3	2	1	0	
5	4	2	0	0	6	0	0	0	5	1	0	0	4	0	1	1	3	1	2	2	1	2	2	2	0	
6	6	0	0	0	6	0	0	0	6	0	0	0	4	1	1	0	2	1	0	3	1	1	2	3	0	
7	5	0	1	0	6	0	0	0	6	0	0	0	4	1	1	0	1	2	2	2	2	2	2	2	0	
8	6	0	0	0	6	0	0	0	6	0	0	0	5	0	1	0	2	1	3	0	2	4	1	0	1	
9	6	0	0	0	6	0	0	0	6	0	0	0	3	2	1	0	3	1	2	2	4	0	0	2	2	
10	6	0	0	0	6	0	0	0	6	0	0	0	4	1	1	0	3	1	0	3	2	3	1	1	1	
11	6	0	0	0	6	0	0	0	6	0	0	0	4	0	2	1	2	0	4	1	3	1	2	0	0	
12	6	0	0	0	6	0	0	0	6	0	0	0	5	0	1	1	1	1	1	2	3	1	2	0	0	
13	6	0	0	0	6	0	0	0	6	0	0	0	5	0	0	1	1	0	2	1	3	3	1	2	0	
14	6	0	0	0	6	0	0	0	6	0	0	0	5	0	1	2	1	2	0	2	3	3	1	2	0	
15	6	0	0	0	6	0	0	0	6	0	0	0	4	1	1	0	4	1	1	3	1	3	1	2	0	

Note. Jordan equilibrium, *TL*; Non-Jordan equilibrium, **BR**.

TABLE VIII
Time Series of Each Pair's Play in Fixed-Pair \mathbf{D} vs \mathbf{D} Regimes

Pair	Period														
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL
2	TR	TL	BL	TR	TR	TL	TL	TR	TR	BR	TL	BL	TR	TL	TL
3	TR	BL	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR
4	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL
5	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL
6	BR	BL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL
7	BR	BR	BL	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR
8	BL	TR	BL	BL	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR
9	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL
10	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL
11	TR	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL
12	TR	BR	BR	BL	BL	TL	BL	TL	TL	TL	TL	TL	TL	TL	TL
13	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL
14	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL
15	TR	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL
16	TR	TR	BL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL
17	BL	BR	TR	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL
18	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL	TL

Note. Jordan equilibrium, *TL*; Non-Jordan equilibrium, **BR**.

3.2. Off-the-Path Choices

Our experiment has cast some doubt on the Jordan model's predictions about which equilibrium will be learned when the true game the players are engaged in has multiple equilibria. We now try to gain some understanding of why the model fared so poorly in these situations.

Perhaps the most salient characteristic of Jordan's theory is the very delicate path-dependence of its predictions. At each stage of play, the model assumes that players condition their actions on inferences they have drawn from their opponent's past play. If the players' inferences are correct in the first few periods, then play can be expected (in our experiment, at any rate) to converge quickly to an equilibrium and then to remain at that equilibrium. But incorrect inferences in the earliest periods—whether due to “incorrect” play by one's opponent or to drawing incorrect inferences from correct play—might never be overcome. Certainly the theory itself has nothing to say about the consequences of such “off-the-path” play.

We therefore concentrate upon just the first two periods of play in each regime. In these first two periods, which always follow immediately after new payoff tables (types) have been drawn for both the Row and the Column players, a player is always able to follow the prescription of the

TABLE IX
Hypothesis Tests for Equilibrium Play in Regimes with a Unique Equilibrium in Mixed Strategies

Game	Null	Alternative	Test statistic	Distribution	<i>P</i> -value
<i>B</i> vs <i>D</i> Random sessions	Cell frequencies are generated by multinomial distribution implied by mixed strategy N.E.	Cell frequencies are not generated by multinomial distribution implied by mixed strategy N.E.	7.41	χ^2 with 3 <i>df</i>	0.06
<i>D</i> vs <i>B</i> Random sessions	Cell frequencies are generated by multinomial distribution implied by mixed strategy N.E.	Cell frequencies are not generated by multinomial distribution implied by mixed strategy N.E.	28.79	χ^2 with 3 <i>df</i>	0.00
<i>D</i> vs <i>B</i> Fixed-pair sessions	Cell frequencies are generated by multinomial distribution implied by mixed strategy N.E.	Cell frequencies are not generated by multinomial distribution implied by mixed strategy N.E.	20.13	χ^2 with 3 <i>df</i>	0.00

TABLE X
 Period 1 Actions and Hypothesis Tests

	Row subjects			Column subjects	
	Type B	Type D		Type B	Type D
Top	32	74	Left	59	69
Bottom	58	28	Right	115	51
Total	90	102	Total	174	120
$F(x_{NJ})^a$	0.004	0.000	$F(x_{NJ})$	0.000	0.060

Test of the hypothesis that Row subjects' choices of Top or Bottom are distributed the same in each type:		Test of the hypothesis that Column subjects' choices of Left or Right are distributed the same in each type:	
z-statistic	-5.144	z-statistic	-4.01
P-value	0.000	P-value	0.000

Note. Boldface entries for actions predicted by Jordan model.

^a $F(x_{NJ})$ is the probability, if choices were equally likely, of observing as few as x_{NJ} non-Jordan choices, where x_{NJ} is the non-boldface entry in this column.

theory, but in later periods this is not generally true. In the first period, for example, only one of the Bayesian ingredients of Jordan’s model comes into play: players are predicted to play a Bayesian Nash equilibrium for their initial beliefs; no updating of beliefs has yet taken place. At the second period, beliefs updated on the basis of first-period observation are always well defined, because each of the two actions a player could have observed in the first period is part of a first-period Bayesian Nash equilibrium for at least one of his opponent’s types. Recall, too, that the experiment was designed in such a way that both the first-period and second-period Bayesian Nash equilibria are unique and are in pure strategies (see Table I).

Table X presents a summary of all first-period play by subjects who did not have a dominant strategy.⁵ The Jordan model predicts that all play will be at the Bayesian Nash equilibrium, for which the corresponding entries in Table X are boldfaced: Row plays Bottom when type \mathcal{B} and Top when type \mathcal{D} ; Column plays Right when type \mathcal{B} and Left when type \mathcal{D} .

The numbers in Table X provide convincing evidence that although play was by no means *always* at the equilibrium prescription, both Row and Column subjects’ play was indeed conditioned upon their types: in each of

⁵When a player has a dominant strategy, incorrect inferences are not an issue for him, either theoretically or in the observed behavior of the subjects in our experiment: as we have already seen, subjects who had a dominant strategy (i.e., Row subjects when they were either type \mathcal{A} or type \mathcal{C}), nearly always played “correctly.”

the four types the equilibrium prescription was chosen substantially more than half the time.⁶ We therefore conclude that Bayesian Nash equilibrium does have some explanatory power for first-period play in our experiment, and that inferences drawn from first-period observations would have been correct approximately two-thirds of the time for our experimental subjects if their inferences were formed according to the Jordan model.

In the second period the Jordan model assumes that the players will combine Bayesian updating of beliefs with Bayesian Nash equilibrium play. Table XI presents a summary of observed second-period play, conditioned upon subjects' information.⁷ The entries in the left half of Table XI indicate, for Row subjects of a given type (\mathcal{B} or \mathcal{D}) whose first-period opponents played a given strategy (L or R), how many of those Row subjects played Top and how many played Bottom in the second period. The analogous figures for the Column subjects are the entries in the right half of Table XI. The boldface entries in Table XI indicate the actions in the unique Bayesian Nash equilibrium at period 2 under the Bayesian updated beliefs: Row plays Top if he has observed his opponent play Left at period 1, and he plays Bottom if he has observed Right; Column's play depends upon his own type as well as upon the first-period play he has observed by his opponent.

We see in Table XI that second-period play was not quite as consistent with the theory as was first-period play. In six of the eight information conditions the theoretically predicted action was chosen more often than the alternative action, but the "wrong" action was chosen more often in the other two information conditions (viz., when Row is type \mathcal{B} , whichever first-period action he has observed). Further, in three of the six information conditions where play was more often correct, the excess of correct play was not statistically significant.

Thus, the actual play by our experimental subjects in the first two periods exhibited substantial noise when compared to the pure-strategy predictions of the Jordan model. This is perhaps not so surprising: each individual subject played in only five regimes and therefore did not gain very much experience at either first- or second-period play. Experience was especially limited in each of the information conditions, since there are two possible information states for a player at period 1 (four states for Row players) and four at period 2. This significant noise in the first two periods upsets the

⁶Below the two panels in Table X we report for each type the probability, if both actions were equally likely, of observing as few non-Jordan (non-boldfaced) choices as the number, x_{NJ} , actually observed.

⁷Table XI includes only plays in random rematching regimes and does not include a player's own first-period play as part of his information state.

TABLE XI
Period 2 Actions and Hypothesis Tests

	Row subjects				Column subjects						
	Subject's information				Subject's information						
	(B ;L)	(B ;R)	(D ;L)	(D ;R)	(B ;B)	(B ;T)	(D ;B)	(D ;T)			
Top	4	20	23	18	Left	35	14	17	37		
Bottom	14	16	3	22	Right	26	45	23	7		
Total	18	36	26	40	Total	61	59	40	44		
$F(x_{NJ})^a$	0.996	0.797	0.000	0.318	$F(x_{NJ})$	0.153	0.000	0.215	0.000		
Equiprobable model hypothesis test for Row subjects:					Equiprobable model hypothesis test for Column subjects:						
Chi-square statistic					21.78	Chi-square statistic					38.97
Degrees of freedom					4	Degrees of freedom					4
<i>P</i> -value					0.00	<i>P</i> -value					0.00
Equiprobable model hypothesis test for Row type B :					Equiprobable model hypothesis test for Column type B :						
Chi-square statistic					6.00	Chi-square statistic					17.62
Degrees of freedom					2	Degrees of freedom					2
<i>P</i> -value					0.05	<i>P</i> -value					0.00
Equiprobable model hypothesis test for Row type D :					Equiprobable model hypothesis test for Column type D :						
Chi-square statistic					15.78	Chi-square statistic					21.35
Degrees of freedom					2	Degrees of freedom					2
<i>P</i> -value					0.00	<i>P</i> -value					0.00

Note. A subject's information is the pair (His own payoff table; His opponent's play at $t = 1$) The action predicted by the Jordan model has a boldface entry.

^a $F(x_{NJ})$ is the probability, if choices were equally likely, of observing as few as x_{NJ} non-Jordan choices, where x_{NJ} is the non-boldface entry in this column.

delicate Bayesian updating exercise at subsequent periods and is the likely source of the model's poor performance in games with multiple equilibria.

4. CONCLUDING REMARKS

In order to evaluate the empirical content of Jordan's model of Bayesian learning in games, we have devised and conducted an experiment with a simple game of incomplete information in which the theoretical implications of the model can be compared with observed play. Chief among our desiderata in developing the experiment were (1) minimization of the number of both multiple and mixed-strategy Bayesian Nash equilibria, espe-

cially in the early periods; (2) minimization of the number of “true” games with multiple or mixed-strategy Nash equilibria; and (3) inclusion of sufficiently many plays in each fixed regime to be able to assess whether convergence to the Nash equilibrium predicted by the Jordan model had occurred.

The experiment we have described provides some support for the idea that equilibrium play can be attained through repeated play in an environment of incomplete information, especially when the true game being played has a unique equilibrium in pure strategies. When the true game the subjects were playing had multiple equilibria, play was no more likely to converge to the equilibrium predicted by the Jordan model than to an alternative equilibrium, and randomly rematching subjects to induce myopic play significantly reduced the likelihood that *any* equilibrium would be attained. The explanation for the lackluster performance of the Jordan model when there are multiple equilibria, and of equilibrium theory in general when subjects are rematched, most likely lies in the noise that was present in the very early rounds of play in each regime. This noise undermines players’ ability to effectively update their beliefs.

The noise we have identified in the multiple-equilibrium data could have come from two sources: individual subjects might have played differently across regimes (for example, best responding in some regimes and following the Bayesian prescription in others, or perhaps following some other decision rule in some regimes), or there might have been unobserved heterogeneity across subjects—individual subjects playing consistently across regimes, but with different subjects using different decision rules. This suggests that one might attempt to identify the source of the noise by postulating alternative decision rules (i.e., alternative learning theories) and investigating whether the variation in play was due to variation in individual subjects’ use of the rules or to heterogeneity in their use across subjects. However, there are two features of our experimental design—a design not intended to choose among competing theories of learning—that stand in the way of such an analysis. First, each subject participated in only five regimes (five realizations of the incomplete-information environment), so there are far too little data on each realization to support any inferences about the decision rule a subject used, or whether he used different rules across regimes. And second, in the particular incomplete-information environment we constructed for the experiment, alternative theories of out-of-equilibrium learning make the same predictions in most of the information conditions that can occur in the experiment.

This suggests devising an experiment expressly designed to have subjects play many realizations of each information condition in an incomplete-information setting in which alternative learning theories make distinct predictions about play.

REFERENCES

- Bloomfield, R. (1994). "Learning Mixed Strategy Equilibrium in the Laboratory," *J. Econom. Behav. Org.*, **25**, 411–436.
- Camerer, C., and Ho, T.-H. (1996). "Experience-Weighted Attraction Learning in Games: A Unifying Approach," mimeo, California Institute of Technology.
- Cheung, Y.-W., and Friedman, D. (1997). "Individual Learning in Normal Form Games: Some Laboratory Results," *Games Econom. Behav.* **19**, 46–76.
- Crawford, V. (1985). "Learning Behavior and Mixed-Strategy Nash Equilibria," *J. Econom. Behav. Org.*, **6**, 69–78.
- Erev, I., and Roth, A. E. (1998). "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *Amer. Econom. Rev.* **88**, 848–881.
- Fudenberg, D., and Kreps, D. (1993). "Learning Mixed Equilibria," *Games Econom. Behav.*, **5**, 320–367.
- Jordan, J. S. (1991). "Bayesian Learning in Normal Form Games," *Games Econom. Behav.*, **3**, 60–81.
- Jordan, J. S. (1993). "Three Problems in Learning Mixed-Strategy Equilibria," *Games Econom. Behav.*, **5**, 368–386.
- Mookherjee, D., and Sopher, B. (1994). "Learning Behavior in an Experimental Matching Pennies Game," *Games Econom. Behav.*, **7**, 62–91.
- Ochs, J. (1995). "Games with Unique Mixed Strategy Equilibria: An Experimental Study," *Games Econom. Behav.*, **10**, 202–217.
- Roth, A. E., and Erev, I. (1995). "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Model in the Intermediate Term," *Games Econom. Behav.* **8**, 164–212.