Games and Economic Behavior **39,** 111–136 (2002)
doi:10.1006/game.2001.0885, available online at http://www.idealibrary.com on IDEAL®

Coordination and Learning Behavior in Large Groups with Asymmetric Players

Amnon Rapoport

Department of Management and Policy, Eller School of Business,
405 McClelland Hall, University of Arizona,
Tucson, Arizona 85721
E-mail: amnon@u.arizona.edu

Darryl A. Seale

University of Nevada at Las Vegas, Las Vegas, Nevada

and

Eyal Winter

The Hebrew University of Jerusalem, Jerusalem, Israel

Received October 13, 1998; published online February 7, 2002

We study a class of large-group, noncooperative, iterated market entry games with complete information, binary choices, and asymmetric players in which the incentive of each player to enter the market decreases the larger the number of entrants. Experimental results from two different studies show remarkable coordination on the aggregate level, which is accounted for successfully by the Nash equilibrium solution. The equilibrium solution is less successful in accounting for the differences among types of players with differential entry costs or differences among players of the same type. Rather, the behavioral patterns observed on the aggregate level are accounted for by a reinforcement-based learning model postulating an initial distribution of individual cutoff points. These cutoff points are assumed to change over time, at a decreasing rate, as a joint function of the decision and outcome of the preceding period. *Journal of Economic Literature* Classification Numbers: C7, D5, D8. © 2002 Elsevier Science (USA)

INTRODUCTION

Nash equilibria can be justified by assuming common knowledge of the game structure, common knowledge of rationality, and common knowledge

0899-8256/02 $35.00

of the agents' posterior beliefs of each other's strategy choices (Aumann & Brandenburger, 1995). This degree of common knowledge, particularly the common knowledge of beliefs, seems unreasonable in coordination games played by large groups (Canning, 1995). This has led several investigators to inquire how Nash equilibria can be justified as the steady states of some adaptive learning process. In the area of coordination, this investigation has been pursued, among others, by Crawford (1991, 1995), Crawford and Broseta (1995), and from a theoretical perspective by Kalai and Lehrer (1993), Marimon and McGrattan (1995), and Young (1993).

Most of this investigation has focused on coordination in noncooperative games played repeatedly by symmetric players (e.g., Van Huyck *et al.*, 1990, 1991, 1993). Our major objective is to extend this line of experimental investigation to a large number of *asymmetric* players interacting repeatedly. In particular, we wish to find out whether the coordination success achieved by symmetric players in the iterated market entry game (e.g., Rapoport *et al.*, 1998; Sundali *et al.*, 1995) changes with breakdown in symmetry. Another major purpose is to formulate and test a simple model of reinforcement-based learning proposed to account for the behavioral regularities observed in the data and changes in these regularities over time.

The Asymmetric Market Entry Game

We consider a group N of n players participating in the following market entry game. At the beginning of each period (stage game) t, $t = 1, 2, \ldots, T$, a different positive integer c, interpreted as the "known capacity of the market," is publicly announced ($1 \leq c \leq n$), and each player $i (i \in N)$ is privately told her entry fee, h_i, for that period. The values of n and T as well as the distribution of h_i are common knowledge. Once the value of c is publicly announced, each player i must decide privately whether to enter the market ($d_i = 1$) or stay out of it ($d_i = 0$). Communication before or during the game is prohibited. Individual payoffs, denoted by $H_i(\mathbf{d})$, are determined by

$$H_i(\mathbf{d}) = \begin{matrix} v, & \text{if } d_i = 0 \\ k + r(c - m) - h_i, & \text{if } d_i = 1, \end{matrix}$$

where v, k, and r are real-valued, commonly known constants that remain fixed across iterations, h_i is an individual entry fee that also remains constant during the game, m ($0 \leq m \leq n$) is the actual number of entrants in this period, and $\mathbf{d} = (d_1, d_2, \ldots, d_n)$ is the vector of individual (binary) decisions. At the end of each period, each player is informed of the value of m and, consequently, her payoff for the period. Information about the decisions and payoffs of other group members is not disclosed. The asymmetry in this market entry game is induced by charging differential entry fees.

In the experiment described below $n = 20$ and $c \in \{1, 3, 5, \ldots, 19\}$. There are exactly four players with the same entry cost, h_j. We refer to players with the same entry cost h_j as players of "type" $j (j = 1, \ldots, J)$; clearly, the players are symmetric within but not between types.

Asymmetry between players may affect coordination success in opposite directions. On the one hand, differential entry costs could facilitate coordination because, given the market capacity c and knowledge of the distribution of types, players with lower entry cost are expected to enter the market before players with a higher entry cost. Given the value of c, the entry cost could then serve as a signal that facilitates coordination. On the other hand, differential entry costs may disrupt coordination because they increase the "cognitive load" of the players. In the symmetric market entry game with zero entry cost, strategically sophisticated players must only consider the exogenously given market capacity c and the value of n when deciding whether to enter. In the asymmetric market entry game they must also consider the distribution of types.

EQUILIBRIUM ANALYSIS

Pure Strategy Equilibria

Table I lists all the pure strategy equilibria for the market entry game with $n = 20$ players divided into five types ($J = 5$) with four members each. The payoff parameters for the game—the ones actually used in our experiment—are $v = k = 1$, $r = 2$, and $h_j = j$ $(j = 1, 2, \ldots, 5)$. The first column on the left presents the market capacity value c, and the second shows the equilibrium total number of entrants, m^*, summed across the five player types. The next five columns present the equilibrium number of entrants of type j, $m_j^* (0 \leq m_j^* \leq 4)$, where $m^* = \sum_j m_j^*$. Finally, the right-hand column presents the group payoff, G, associated with equilibrium play. For each value of c the equilibria are ordered in terms of G.

To classify these equilibria, we refer to pure-strategy equilibria as *monotonic* (denoted by "#" in Table I) if $m_j^* \geq m_{j+1}^* (j = 1, 2, 3, 4)$. Note that most equilibrium profiles in Table I are nonmonotonic. An equilibrium is *compact* (denoted by "$*$" in Table I) if it is monotonic and it includes at most a single m_j^* that differs from 0 or 4. Compactness means that all of the four members of each type behave exactly in the same way (either enter or stay out) with at most one type whose members may behave differently from one another. An equilibrium is *efficient* if it maximizes the value of G across all the pure strategy equilibria for a given c value. Efficient equilibria are always compact. Considerations of symmetry within type and efficiency

TABLE I
Pure Strategy Equilibria for the Market Entry Game with $n = 20$ and Entry Costs $h_j = 1, 2, 3, 4, 5$

c	M^*	m_1	m_2	m_3	m_4	m_5	G		
1	0	0	0	0	0	0	20	#	*
3	2	2	0	0	0	0	22	#	*
3	2	1	1	0	0	0	21	#	
3	2	0	2	0	0	0	20		
5	4	4	0	0	0	0	24	#	*
5	4	3	1	0	0	0	23	#	
5	4	2	2	0	0	0	22	#	
5	4	1	3	0	0	0	21		
5	4	0	4	0	0	0	20		
7	5	4	1	0	0	0	34	#	*
7	5	4	0	1	0	0	33		
7	5	4	0	0	1	0	32		
7	6	4	2	0	0	0	24	#	
7	6	3	3	0	0	0	23	#	
7	6	2	4	0	0	0	22		
9	7	4	3	0	0	0	38	#	*
9	7	4	2	1	0	0	37	#	
9	7	4	1	2	0	0	36		
9	7	4	2	0	1	0	36		
9	7	4	0	3	0	0	35		
9	7	4	1	1	1	0	35	#	
9	7	4	0	2	1	0	34		
9	7	4	1	0	2	0	34		
9	7	4	0	1	2	0	33		
9	7	4	0	0	3	0	32		
9	8	4	4	0	0	0	24	#	*
11	9	4	4	1	0	0	41	#	*
11	9	4	3	2	0	0	40	#	
11	9	4	4	0	1	0	40		
11	9	4	2	3	0	0	39		
11	9	4	3	1	1	0	39	#	
11	9	4	1	4	0	0	38		
11	9	4	2	2	1	0	38	#	
11	9	4	3	0	2	0	38		
11	9	4	1	3	1	0	37		
11	9	4	2	1	2	0	37		
11	9	4	0	4	1	0	36		
11	9	4	1	2	2	0	36		
11	9	4	2	0	3	0	36		
11	9	4	0	3	2	0	35		
11	9	4	1	1	3	0	35		
11	9	4	0	2	3	0	34		

TABLE I—*Continued*

c	M^*	m_1	m_2	m_3	m_4	m_5	G		
11	9	4	1	0	4	0	34		
11	9	4	0	1	4	0	33		
13	11	4	4	3	0	0	43	#	*
13	11	4	3	4	0	0	42		
13	11	4	4	2	1	0	42	#	
13	11	4	3	3	1	0	41	#	
13	11	4	4	1	2	0	41		
13	11	4	2	4	1	0	40		
13	11	4	3	2	2	0	40	#	
13	11	4	4	0	3	0	40		
13	11	4	2	3	2	0	39		
13	11	4	3	1	3	0	39		
13	11	4	1	4	2	0	38		
13	11	4	2	2	3	0	38		
13	11	4	3	0	4	0	38		
13	11	4	1	3	3	0	37		
13	11	4	2	1	4	0	37		
13	11	4	0	4	3	0	36		
13	11	4	1	2	4	0	36		
13	11	4	0	3	4	0	35		
15	12	4	4	4	0	0	68	#	*
15	13	4	4	4	1	0	44	#	*
15	13	4	4	3	2	0	43	#	
15	13	4	3	4	2	0	42		
15	13	4	4	2	3	0	42		
15	13	4	3	3	3	0	41	#	
15	13	4	4	1	4	0	41		
15	13	4	2	4	3	0	40		
15	13	4	3	2	4	0	40		
15	13	4	2	3	4	0	39		
15	13	4	1	4	4	0	38		
17	14	4	4	4	2	0	72	#	*
17	14	4	4	4	1	1	71	#	
17	14	4	4	4	0	2	70		
17	15	4	4	4	3	0	44	#	
17	15	4	4	3	4	0	43		
17	15	4	3	4	4	0	42		
19	16	4	4	4	4	0	76	#	*
19	16	4	4	4	3	1	75	#	
19	16	4	4	4	2	2	74	#	
19	16	4	4	4	1	3	73		
19	16	4	4	4	0	4	72		

*, Compact equilibrium.

#, Monotonic equilibrium.

would seem to favor compactness. We would expect, then, that if the players converge to equilibrium it would be compact.

Compactness implies several testable predictions. Table I shows that if $3 \leq c \leq 5$ only players of type 1 will enter; if $c \leq 9$ only players of types 1 and 2 will enter; if $c \leq 13$ only players of types 1, 2, and 3 will enter; and for any $c \in \{1, 3, \ldots, 19\}$ no player of type 5 will ever enter. Tests of these and related predictions will be reported below.

Mixed Strategy Equilibria

We next consider completely mixed strategy symmetric equilibria in which players of the same type enter the market with the same probability. We have proved elsewhere (Rapoport *et al.*, 1997) that for our design with $J = 5$ types there exists no symmetric and completely mixed equilibrium because no more than two types can mix. However, there exist symmetric and partly mixed equilibria where at the most two types can mix. For example, if $c = 11$ there exists an equilibrium where types 1 and 2 enter, types 4 and 5 stay out, and type 3 players each enters with probability 1/6.

EXPERIMENT 1: THE MARKET ENTRY GAME WITH COMMON PRIOR

Method

Subjects. Forty subjects, divided into two groups of twenty players each, participated in the experiment (Experiment 1). The subjects consisted of University of Arizona students who, in response to advertisements in the local student newspaper, volunteered to take part in a two-hour experiment on economic decision making for monetary payoff contingent on performance.

Procedure. Experiment 1 was conducted in the Economic Science Laboratory at the University of Arizona. Upon arrival at the laboratory the subjects were seated at one of 20 computer terminals separated by partitions. Communication between the subjects was prohibited. The instructions were presented on individual computer screens in front of the subjects. Written copies of the instructions were made available for the entire duration of the experiment.

Each subject was given an endowment of 34 francs (a fictitious currency converted to US dollars at the end of the experiment at the rate of 10 francs = \$1.50). The subjects were instructed that payoffs gained or lost on each trial would be added or subtracted from their endowment. After all the subjects completed reading the instructions, they participated in four practice trials to acquaint them with the experimental procedure. These were followed immediately by 100 trials with payoff contingent on performance.

Each trial, or stage game, consisted of two parts. In the first part a value of c was displayed on the subject's monitor. The payoff function for entering or staying out had the form (in francs)

stay out: 1
enter: $1 + 2(c - m) - h_j$,

where m and h_j are as defined above. The subjects were divided into $J = 5$ types with $f = 4$ members each. Type j's entry cost was set at $h_j = j$, $j = 1, \ldots, 5$.

To reduce the burden of computation, the computer displayed on each trial a table listing the payoffs for all possible numbers of entrants ranging from 0 to 20, given the value of c for the trial. All together five different tables were presented, one for each type, to reflect the difference in the entry cost.

In the second part of the trial, the subjects were asked to type in their decision (Y for "enter," N for "stay out"). No time constraint was imposed. Once all 20 subjects entered their decisions, each was informed of the total number of entrants for the trial (m), the subject's payoff for the trial (which could be deduced from the knowledge of the value of m), and his/her cumulative payoff from the beginning of the experiment. Information about the decisions and payoffs of the other subjects, whether of the same or different type, was not disclosed.

All together 10 different values of c were presented, one on each trial. In each block of 10 trials these 10 values of c were randomly sampled without replacement from the set of odd integers $\{1, 3, \ldots, 19\}$. Thus, each subject observed the same value of c a total of 10 times during the experiment but in a different position within each block.

Each subject was given paper and pencil for taking notes, keeping record of his/her accumulated payoff, and recording his/her outcomes. The subjects were paid individually their accumulated earnings for the 100 trials, plus a \$5 show-up fee, and dismissed. The mean payoff across the two groups, excluding the show-up fee, was \$24.60.[1] The entire experiment lasted about 100 minutes.

Results

In presenting the results, we progress from the population to the individual level of analysis. We begin by testing several implications of

[1]A player deciding to stay out on all 100 trials would have earned a total of 134 ($100 + 34$) francs, or \$20.10.

the equilibrium solution on the aggregate level, disregarding differences between types or between players within types. Next, we examine the effects of type by studying the proportions of entry and of experience by studying both the total number of entries and the correlation between the predicted and observed numbers of entries across the 10 values of c. Finally, on the individual level, we examine and characterize the individual decision policies.

Group Behavior and Block Effects. Table II presents the number of entrants by block and the c value for each group separately. The 10 values

TABLE II
Total Number of Observed Entries by c Value, Block, and Group

Group 1														
	Block													
c	1	2	3	4	5	6	7	8	9	10	Total	Mean	SD	**Prediction**
1	2	0	0	0	0	1	0	0	0	0	3	0.3	0.67	**0**
3	0	1	2	2	3	4	2	1	2	2	19	1.9	1.10	**2**
5	7	6	2	3	5	3	3	4	5	4	42	4.2	1.55	**4**
7	1	9	3	7	6	7	7	5	4	5	54	5.4	2.32	**5 (6)**
9	9	3	11	9	8	7	8	8	7	7	77	7.7	2.06	**7 (8)**
11	9	7	16	9	11	8	10	9	10	9	98	9.8	2.44	**9**
13	11	10	10	8	13	11	11	10	12	12	108	10.8	1.40	**11**
15	7	16	15	10	13	12	13	13	15	12	126	12.6	2.63	**12 (13)**
17	12	15	12	11	15	15	15	15	14	16	140	14.0	1.70	**14 (15)**
19	13	18	14	14	16	16	17	17	17	16	158	15.8	1.62	**16**
Total:	71	85	85	73	90	84	86	82	86	83	825	82.5		**80 (84)**
r	0.86	0.90	0.86	0.95	0.99	0.98	0.99	0.99	0.99	0.99	1.00			
Group 2														
	Block													
c	1	2	3	4	5	6	7	8	9	10	Total	Mean	SD	**Prediction**
1	5	1	1	0	1	0	1	0	0	0	9	0.9	1.52	**0**
3	2	4	1	4	2	1	2	3	2	1	22	2.2	1.14	**2**
5	3	4	4	5	4	6	2	3	3	4	38	3.8	1.14	**4**
7	6	7	5	6	8	7	7	9	6	5	66	6.6	1.26	**5 (6)**
9	14	8	11	7	6	9	5	6	8	8	82	8.2	2.66	**7 (8)**
11	9	12	8	10	8	11	9	8	10	8	93	9.3	1.42	**9**
13	12	11	13	11	12	9	13	10	12	10	113	11.3	1.34	**11**
15	14	15	11	14	13	14	12	15	11	15	134	13.4	1.58	**12 (13)**
17	16	14	15	17	13	15	15	13	14	15	147	14.7	1.25	**14 (15)**
19	15	17	18	16	16	17	17	15	16	17	164	16.4	0.97	**16**
Total:	96	93	87	90	83	89	83	82	82	83	868	86.8		**80 (84)**
r	0.88	0.98	0.96	0.98	0.97	0.97	0.97	0.94	0.98	0.98	0.99			

of c in the left-hand column appear in an ascending order, not in the actual order of their presentation that varied from block to block. The total number of entries summed across the 10 blocks is presented in Column 12, and the mean and standard deviation are shown in Columns 13 and 14, respectively. The pure strategy equilibrium values, m^*, are presented in the right-hand column of the table.

Our first major finding is the remarkable coordination on the aggregate level. The null hypothesis that the mean number of entries is equal to m^* could not be rejected (t-test, $p > 0.05$) in all 20 cases (2 groups times 10 values of c). A statistic that measures the predictive power of the equilibrium solution across the 10 values of c is the correlation between m and m^* (for the efficient equilibrium). The product-moment correlations between m and m^* (denoted by r), each based on 10 pairs of values, are presented at the bottom row of Table II, one for each block. The correlations are seen to increase across blocks from about 0.87 to 0.99. For both groups, the correlation between the mean number of entries across blocks (Column 13) and m^* is practically equal to unity.

Changes across Blocks. Consider first the total number of entries across the 10 values of c for each block separately. These frequencies are shown in the next to the bottom row of Table II. The equilibrium number of entries across the 10 values of c ranges between 80 and 84 (see the right-hand column of Table II). Inspection of Table II indicates learning across blocks with convergence to a frequency of 82 or 83 reached after 5 to 6 blocks. It also shows small differences between the groups, with Group 1 initially having a relatively small number of entries in Block 1 and Group 2 a relatively large number of entries. The initial difference between the two groups, which reflects the different dispositions of the players at time $t = 0$, rapidly diminishes as the total number of entries approaches the predicted value. Thus, the total number of entries in Blocks 6–10 of Group 1 is between 82 and 86 and in Group 2 between 82 and 83.

Examination of the correlations between m and m^* across the 10 c values shows a similar pattern of convergence. In both groups, the coordination reached in Block 1 is not at all impressive; monotonicity of m in c is not maintained, and the correlations between m and m^* are below 0.90. However, coordination improves rapidly: in Blocks 5–10 there are hardly any violations of monotonicity; indeed, the correlations between m and m^* in these blocks are equal to or greater than 0.98 in Group 1 and (with a single exception) equal to or greater than 0.97 in Group 2.

These block effects in both the total number of entries and the correlation between observed and predicted number of entries suggest some sort of an adaptive learning process converging to the equilibrium total number of entries. Before studying adaptive learning, we turn next to a more detailed examination of the type effects and individual trial outcomes.

Differences between Types. Table III (upper panel) presents the proportions of entry across the 10 blocks of trials partitioned into the 5 types of players. The observed proportions are shown by the values of c and are summarized across the two groups in the second column from the left. For example, consider the eight (two groups by four players) type-2 players of Groups 1 and 2. Out of a possible total number of 80 entries (8 players by 10 iterations), these players entered twice on $c = 1$, resulting in a proportion of entry equal to $2/80 = 0.025$. For the same 8 type-2 players, the proportions of entry for $c = 3$, $c = 5$, and $c = 19$ are 0.338, 0.525, and 0.975, respectively. Across all five types, the proportion of entry for $c = 11$ is 0.478.

Comparison of Tables I and III rejects the (pure strategy) equilibrium as an explanatory concept for the differences between types. Whether the equilibria are monotonic or not, all type-1 players are predicted (Table I) to always enter if $c \geq 9$. The upper panel of Table III shows that the actual proportions of entry for type-1 players are considerably smaller than 1, particularly for $c = 9$ (0.375) and 11 (0.500). Similarly, for all the pure strategy equilibria in Table I, players of type 5 are not expected to enter if $c \leq 15$. In contrast, Table III shows a high proportion of entries by type-5 players even for c values as small as 7 (0.300), 9 (0.393), and 11 (0.478). A third implication of the equilibrium solutions in Table I is that the total proportion of entry should decrease in the entry fee. If the players achieve the efficient equilibrium, the total proportion of entry across the 10 values of c should be 0.85, 0.60, 0.40, 0.15, and 0 for types 1, 2, 3, 4, and 5, respectively. In contrast, in both groups type-2 players entered more often than type-1 players. Across the two groups, type-2 players entered 65.75% (526/800) of the time compared to 46.25% (370/800) for type-1 players. The difference of almost 20 percentage points is significant ($t_{14} = 2.024$, $p < .05$ by a one-sided test).

Table III (upper panel) shows that in both groups and for all five types of players the proportion of entry within type increases in c. There exist few violations of monotonicity, but all are mostly minor. Players of each type are seen to be sensitive to changes in the market capacity c, entering more often as c increases.

Switches in Decision. For each subject separately, we counted the number of times—out of 10—that she switched her decision for a given value of c from block $b-1$ to block b ($b = 2, 3, \ldots, 10$). If a subject does not change her decisions for all 10 values of c from block $b-1$ to block b, the resulting number of switches (alternations) is zero. If she changes her decision for each value of c, this number is 10. Therefore, across the 9 pairs

TABLE III
Observed and Simulated Proportions of Entry by c Value and Player Type across Groups

Observed	Type					Across Types	
c	1	2	3	4	5	**Experiment 1**	**Experiment 2**
1	0.000	0.025	0.087	0.013	0.025	0.030	**0.032**
3	0.088	0.338	0.050	0.013	0.050	0.102	**0.118**
5	0.200	0.525	0.150	0.100	0.025	0.200	**0.254**
7	0.312	0.600	0.288	0.138	0.163	0.300	**0.304**
9	0.375	0.725	0.400	0.200	0.287	0.397	**0.404**
11	0.500	0.850	0.500	0.263	0.275	0.478	**0.444**
13	0.650	0.813	0.613	0.312	0.375	0.552	**0.564**
15	0.825	0.800	0.675	0.525	0.425	0.650	**0.630**
17	0.850	0.925	0.862	0.538	0.412	0.717	**0.729**
19	0.850	0.975	0.888	0.800	0.513	0.805	**0.806**
Total	0.463	0.658	0.451	0.290	0.255	0.423	**0.429**
Simulated	Type					Across Types	
c	1	2	3	4	5	**Experiment 1**	**Experiment 2**
1	0.058	0.027	0.020	0.010	0.003	0.024	**0.017**
3	0.188	0.085	0.050	0.030	0.015	0.073	**0.071**
5	0.533	0.135	0.050	0.035	0.028	0.156	**0.180**
7	0.663	0.368	0.043	0.015	0.015	0.220	**0.294**
9	0.815	0.630	0.025	0.077	0.053	0.341	**0.398**
11	0.840	0.738	0.330	0.097	0.073	0.416	**0.466**
13	0.940	0.908	0.768	0.262	0.175	0.611	**0.597**
15	0.957	0.935	0.873	0.413	0.270	0.690	**0.667**
17	0.958	0.913	0.830	0.635	0.295	0.726	**0.791**
19	0.967	0.955	0.930	0.900	0.671	0.884	**0.915**
Total	0.692	0.570	0.403	0.247	0.159	0.414	**0.440**

of adjacent blocks and 10 values of c, the minimum and maximum number of switches per subject are 0 and 90, respectively. In actuality, the total number of switches ranged between 0 and 32 for the players in Group 1, with a median of 16, and between 2 and 36 for the players in Group 2, with a median of 14.5. We find no differences between the frequency distributions of switches of the two groups (Kolmogorov–Smirnov test, $p > 0.10$). Having substantial evidence for block-to-block changes in decision policies, we study next how these switches were affected by player type and experience.

The upper panel of Table IV shows the mean number of switches per group, computed across the subjects of both groups in terms of type and block. The frequencies in this table are summed across the 10 values of c.

TABLE IV
Observed and Simulated Mean Number of Switches between Adjacent Blocks by Player Type and Block across Groups

Observed										
	Block									Across
Type	2	3	4	5	6	7	8	9	10	Blocks
1	10.5	8.0	4.0	3.5	7.5	4.5	4.5	6.0	3.5	52.0
2	13.0	11.0	8.0	9.0	6.0	6.5	6.0	4.5	3.0	67.0
3	14.0	15.5	12.5	8.5	7.5	7.0	5.5	5.5	6.0	82.0
4	9.5	10.0	8.5	8.0	5.5	6.5	4.0	4.0	3.5	59.5
5	9.5	10.0	11.0	8.5	4.0	2.5	2.5	4.5	1.0	53.5
Total	56.5	54.5	44.0	37.5	30.5	27.0	22.5	24.5	17.0	314.0

Simulated										
	Block									Across
Type	2	3	4	5	6	7	8	9	10	Blocks
1	14.2	10.9	8.4	8.2	7.1	4.5	3.0	2.4	1.8	60.6
2	15.5	12.8	9.7	8.4	7.4	4.8	3.3	2.2	1.5	65.7
3	14.6	11.4	8.0	7.6	6.5	4.7	3.1	2.2	1.7	59.7
4	15.7	11.6	8.0	7.9	6.7	4.4	3.1	2.4	1.7	61.3
5	15.0	11.6	7.9	8.8	8.1	4.8	3.7	2.7	2.1	64.8
Total	75.0	58.3	42.1	40.8	35.8	23.2	16.2	11.9	8.8	312.1

Table IV shows that the number of switches decreases with experience from 56.5 in Block 2 to 17.0 in Block 10. Out of a total number of 400 decisions (40 subjects by 10 c values), the players alternated their decision between Blocks 1 and 2 a total of 117 times (56.5/2 = 29.25%). In contrast, out of a total of 400 decisions, the same players alternated their decisions for the same 10 values of c between Blocks 9 and 10 only 34 times (17.0/2 = 8.5%). The correlation between the block number b (top row of Table IV) and the total number of switches (bottom row of Table IV) is negative and highly significant: $r = -0.972$. Although slightly more erratic, with several violations of monotonicity, the same declining trend also appears for each of the five types.

Individual Decision Policies. Table IV provides strong evidence that the individual decision policies determining whether to enter or stay out become more stable with experience. Stability here means that the decision whether to enter, given a specific market capacity c, is no longer altered with further experience. This hypothesis is substantiated when we examine the individual decision policies, or *profiles*, and the way they change over time.

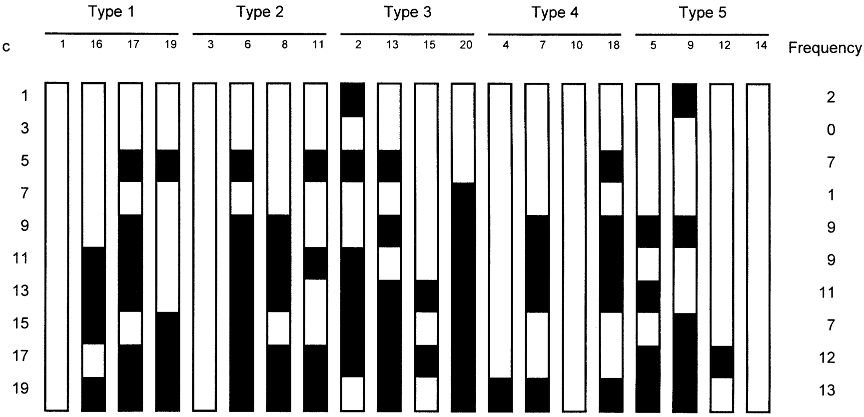

FIG. 1. Individual decision policies in Block 1 (Trials 1–10), Group 1 of Experiment 1.

Figure 1 displays the individual profiles in Block 1 of Group 1. An individual profile is presented by a column with 10 cells, each corresponding to a different value of c. A cell is empty if the decision is to stay out and is filled out otherwise. The 20 subjects are grouped by the 5 types. Similar results (not presented here) were obtained for Group 2 (see Rapoport *et al.*, 1997).

Figure 1 does not exhibit any discernible patterns in most of the individual profiles. There are considerable differences between members of the same type far exceeding any differences, if any, between types. In particular, there is only partial evidence for cutoff policies of the kind "enter if and only if $c \geq c^*$," where c^* is some positive real number. Note, for example, Subject 2 in Group 1 who, in violation of such a policy, entered on $c = 1, 5, 11, 13, 15,$ and 17. All together, across both groups, there are 13 (out of 40) subjects with deterministic decision policies in Block 1. Of the six subjects in Group 1 (Fig. 1), four never entered (Subjects 1, 3, 10, and 14), one always entered if $c \geq 7$ (Subject 20), and one if $c \geq 19$ (Subject 4). Across the two groups, the frequency of entry is distributed more or less equally across types, with players of lower types entering slightly more frequently than players of higher types. Common knowledge of the distribution of types and private information about entry fee seem to have at most minor effects on the individual profiles shown in Block 1.

Individual decision policies were substantially modified by experience. The upper panel of Table V shows the number of cutoff policies, both for individual groups and across the two groups of Experiment 1, for each block of 10 trials. The number of cutoff policies (out of a maximum of 40 across the two groups) increases over blocks from 13 to 34. Figure 2 exhibits the individual profiles of the 20 subjects in Group 1 in Block 10 (Trials 91–100). The individual profiles are presented separately for each

TABLE V
Frequency of Cutoff Decision Policies by Block and Group for Experiment 1 and Percentages of Cutoff Decision Policies for Experiments 1 and 2

	Observed Frequency of Cutoff Policies in Experiment 1									
	Block									
Group	1	2	3	4	5	6	7	8	9	10
1	6	10	9	11	14	16	17	18	18	17
2	7	9	8	10	11	8	13	12	15	17
Mean Exp. 1	6.5	9.5	8.5	10.5	12.5	12.0	15.0	15.0	16.5	17.0
	Observed Percentages of Cutoff Decision Policies for Both Experiments									
	Block									
	1	2	3	4	5	6	7	8	9	10
Experiment 1	32.5	47.5	42.5	52.5	62.5	60.0	75.0	75.0	82.5	85.0
Experiment 2	25.0	50.0	62.5	60.0	62.5	80.0	65.0	77.5	75.0	90.0

group, with the subjects in each group labeled in the same way as in Fig. 1. A comparison of Figs. 1 and 2 is, therefore, rather straightforward. Inspection of Fig. 2 verifies that 85% of the subjects in Group 1 converged to cutoff policies by the end of the experiment. (Similar results for Group 2 can be observed in Rapoport *et al.*, 1997.) There are only three exceptions in Group 1 (Subjects 8, 11, and 7), all of them minor. Figure 2 further shows that the individual subjects differ from one another in the value of the (inferred) cutoff point c^*, but that these differences are related to but not perfectly correlated with type. Adaptive learning models should account

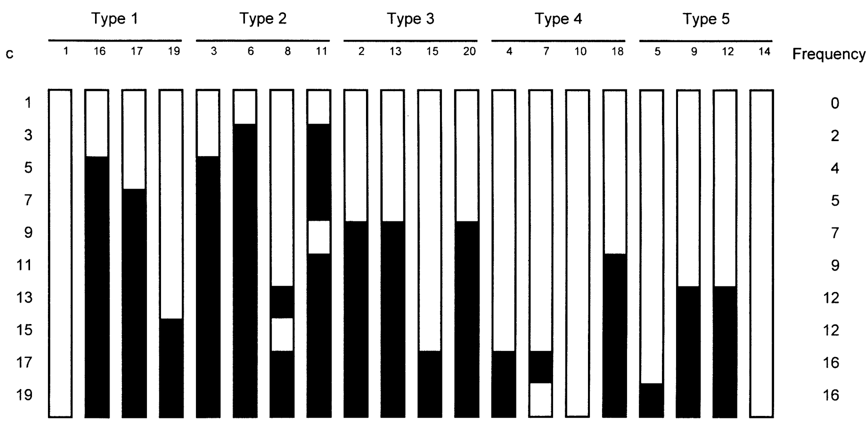

FIG. 2. Individual decision policies in Block 10 (Trials 91–100), Group 1 of Experiment 1.

for the convergence of the individual cutoff policies documented in Figs. 1 and 2 and the upper panel of Table V.

ADAPTIVE LEARNING

Taken together, the results presented in Table II show remarkable coordination with strong support for the Nash equilibrium on the aggregate level. However, the equilibrium solution is rejected as a descriptive concept when the entry data are broken down by type. Even more damaging is the large number of alternations, which provide no support for equilibrium play. Although there is substantial within-subject variability in the individual decisions for the same value of c across blocks, it is definitely not due to randomization. Rather, the results seem to support an evolutive defense of the equilibrium concept (Binmore, 1987) in which the players, boundedly rational to some degree or another, find their way to equilibrium by some process of trial-and-error adjustment. A few subjects seem to begin the experiment with deterministic decision policies, whereas the majority display no such consistent patterns. Coordination improves gradually through some form of an adaptive learning process in which each player modifies her behavior, in this case a decision policy specifying a unique decision for every value of c, in response to her own payoff. The rate of alternation between entering a market with a given capacity or staying out steadily decreases with experience, with almost all subjects in the final block adopting cutoff policies with different cutoff points that reflect partly, but not completely, the asymmetry between players induced by the differential entry fees.

Given these behavioral regularities, we propose to model the individual dynamics of play in the following way. Initially, at the beginning of the first trial ($t = 1$) before the market capacity for this trial is announced each player i ($i \in N$) is assumed to be characterized by a decision policy of the form

$$\text{“enter if and only if } c \geq c^*_{i,1}\text{,”} \qquad t = 1, 2, \ldots, T.$$

The values of $c^*_{i,1}$ are drawn randomly from a distribution defined over the real interval $[1, n]$; they are not conditioned on type. Rather than estimating this distribution from the data in Block 1 (and thereby improving the goodness of fit), we assume that this distribution is uniform.

Once they are informed of the value of m for trial t, players who *stayed out* on trial t compare their actual payoff for the trial (v) to the payoff that players of their type *would have received* if entering, $[k + r(c - m) - h_j]$.

Then, their cutoff point $c^*_{i,t}$ is revised as

$$c^*_{i,t+1} = \begin{cases} c^*_{i,t} - w^+_t[(k + r(c-m) - h_j) - v], & \text{if } [k + r(c-m) - h_j] > v, \\ c^*_{i,t}, & \text{otherwise,} \end{cases} \quad (1)$$

where w^+_t, the decision weight reflecting the effect of the subject's opportunity loss due to staying out on trial t, is a free parameter. This parameter is assume to be discounted in time; i.e.,

$$w^+_{t+1} = (1-d)w^+_t, \qquad t = 1, 2, \ldots, T-1,$$

where $0 \leq d < 1$ is a common discount factor.

Once they are informed of the value of m for trial t, players who *entered* the market on that trial are assume to compare their actual payoff to the payoff that players of their type *would have received* by staying out. Their cutoff point is then revised as

$$c^*_{i,t+1} = \begin{cases} c^*_{i,t} + w^-_t[v - (k + r(c-m) - h_j)], & \text{if } v \geq [k + r(c-m) - h_j], \\ c^*_{i,t}, & \text{otherwise,} \end{cases} \quad (2)$$

where w^-_t, the decision weight reflecting the effect of the player's opportunity loss due to entering on trial t, is a free parameter. This parameter, too, is discounted; i.e.,

$$w^-_{t+1} = (1-d)w^-_t, \qquad t = 1, 2, \ldots, T-1,$$

where d—the discount factor—is as defined above.[2]

Finally, to allow variability within the decision rule, particularly in the first few periods, we assume that each player i follows the decision rule above on period t with probability $1 - e_{i,t}$ and makes the opposite decision (error) with probability $e_{i,t}$, where

$$e_{i,t} = q_{i,t}\exp(-|c - c^*_{i,t}|), \qquad t = 1, 2, \ldots, T, \quad (3)$$
$$q_{i,t+1} = (1-d)q_{i,t},$$

and $q_{i,1}$ is an individual parameter randomly chosen from the uniform distribution defined over the interval [0, 0.5].

The probability of error $e_{i,t}$ satisfies two desiderata: (1) The smaller the absolute difference between c and $c^*_{i,t}$, the greater the probability of error. (2) The probability of error diminishes with time. As shown by Eq. (3),

[2] An alternative way to model the adjustment process is in terms of the opportunity loss due to switching, i.e., by comparing the player's payoff to the payoff this player would have obtained by switching her decision on trial t. The only difference caused by this change is replacing the term $(c - m)$ in Eq. (2) by the term $(c - m - 1)$. Our simulation of the model suggests that this model does not fit the results as well as the original model.

the probability of error on the first trial reaches its peak when $c = c^*_{i,1}$. If $q_{i,1} = 0.5$ and the difference $|c - c^*_{i,1}| = 1.0$, then $e_{i,1} = 0.184$, and if $|c - c^*_{i,1}| = 2.0$, then $e_{i,1} = 0.068$. As time progresses, $e_{i,t}$ approaches zero, with the rate of decline depending on the value of d.

In the present model all of the n players are characterized by cutoff policies (with individual error terms that diminish with time). Only players earning less than other players of their type who behaved differently are assumed to change their cutoff, and this change is assumed to be proportional to the difference in payoffs. The two decision weights w^+ and w^-, which are assumed to be common to all players regardless of type (an assumption that can be relaxed easily), and the probability of error, which differs from player to player, are discounted at the same rate.[3] With the two decision weights and the probability of error diminishing with experience, the decision policies converge to cutoff point policies with (possibly different) fixed cutoff points that are no longer affected by the payoff outcomes.

Model Testing

For each set of the three model parameters w^+, w^-, and d we simulated the decision behavior of 100 groups of 20 artificial players each playing the market entry game described above for 100 periods. Each group included five types with four players in each type, exactly as in our experiment. For each of the 20 artificial players, we randomly chose a cutoff point $c^*_{i,t}$ from a uniform distribution defined over the interval $[0, n = 20]$. The 20 artificial players were subjected to the same sequence of c values that was observed by the real subjects.[4] The cutoff points were then revised in accordance with the assumptions of the model. The 100 groups of artificial players differed only from one another only in the value of the seed number that was used to generate the random numbers, which, in turn, determined the initial individual cutoff points and probability of error. The parameter values of the model were set the same for all artificial players at the values $w^+ = 1.3$, $w^- = 1.8$, and $d = 0.02$, which were found through an extensive grid search for the best fitting parameters.[5]

[3]Of course, one could condition the two parameters w^+ and w^- on the type of player, or even let them vary from player to player within type. This change would have led to a proliferation of free parameters and, consequently, to a considerably less parsimonious model.

[4]To allow direct comparison between the two groups of real subjects, both were presented with the same sequence of the c values.

[5]Standard maximum likelihood methods do not apply in our case. The parameter values were determined by an extensive search of a 3-dimensional grid, with different starting points, for "best fitting" parameter values that minimize the absolute difference between observed and predicted total number of switches. Different parameter values might have been obtained by searching a finer grid or minimizing another objective criterion. In any case, the goodness of fit of the model could only improve.

Group Behavior and Block Effects. The mean number of entries was computed separately for each block of 10 trials across the 100 groups of artificial players. In addition, as we did with the real data (Table II), we computed the mean total number of entries, m, and the correlations, $r(m, m^*)$, between the simulation means and the 10 values of m^* associated with the efficient equilibrium. These computations were conducted separately for each of the 10 blocks. The results of the simulation and these computations are presented in Table VI using the same format as Table II.

Table VI shows considerable variability in the decisions of the artificial players on the first 10 trials (Block 1). The correlations between m and m^* equal 0.58 for Block 1 and 0.71 for Block 2, reflecting the frequent violations of the monotonicity of m in c. These correlations are considerably smaller than the ones computed for the real subjects (Table II). The correlations increase steadily across blocks, reaching the value of 0.97 in Block 7. The total number of entries across the 10 values of c also shows considerable oscillation, from a mean of 87.4 in Block 1 to 76.5 in Block 2 and 81.3 in Block 3. These frequencies, however, converge rapidly; in the last 5 blocks the total number of entries across the 10 c values varies between 82.6 and 85.3, mimicking the results of the real subjects in Table II.

Comparison of the mean number of entries across blocks with m^* shows that, similarly to the observed results, the simulation results converge to

TABLE VI
Mean Number of Entries by c Value and Block for 100 Groups of 20 Simulated Subjects

	Block													
c	1	2	3	4	5	6	7	8	9	10	Total	Mean	SD	**Prediction**
1	3.6	0.3	0.2	0.3	0.1	0.0	0.0	0.0	0.0	0.0	4.7	0.5	1.09	**0**
3	7.2	2.0	1.4	1.1	0.3	0.1	0.4	0.4	0.8	0.9	14.7	1.5	2.10	**2**
5	4.1	5.9	1.2	0.8	4.6	1.9	1.8	2.7	4.2	3.9	31.2	3.1	1.68	**4**
7	2.2	2.8	3.1	1.8	7.0	4.9	3.9	5.2	6.5	6.7	44.0	4.4	1.93	**5 (6)**
9	9.0	4.9	10.5	6.0	3.0	4.9	6.3	7.6	8.1	8.0	68.3	6.8	2.22	**7 (8)**
11	7.4	10.2	7.2	7.3	4.5	9.0	10.8	8.6	9.3	8.9	83.1	8.3	1.80	**9**
13	13.8	8.1	14.0	15.7	14.8	8.7	12.0	11.4	12.0	11.6	122.2	12.2	2.48	**11**
15	12.0	14.2	12.5	16.1	16.6	16.6	13.1	12.3	12.4	12.1	138.0	13.8	1.94	**12 (13)**
17	9.3	16.1	13.6	11.7	12.9	18.8	16.9	16.3	15.2	14.6	145.2	14.5	2.76	**14 (15)**
19	18.9	12.2	17.6	19.0	17.8	19.6	19.0	18.2	16.8	17.8	176.8	17.7	2.10	**16**
Total	87.4	76.5	81.3	80.0	81.6	84.5	84.2	82.6	85.3	84.7	828.0			**80 (84)**
r	0.58	0.71	0.86	0.85	0.84	0.93	0.97	0.97	0.98	0.98				

values very close to the equilibrium. Across the 10 blocks, the values of m (third column from the right) increase monotonically in c, as predicted. Monotonicity is seen to be reached around the seventh block. A comparison of the mean total number of entries in Blocks 7 to 10, after monotonicity has been achieved, and m^* shows only minor deviations.

Comparison of Tables II and VI shows that the learning model captures the major trends observed in the data. In both sets of data—real and artificial—we observe the total number of entries which are well accounted for by the equilibrium solution, convergence to monotonicity of m in c across the 10 blocks, and correlations $r(m, m^*)$ that approach unity. Even the standard deviations are comparable in their values; they range between 0.67 and 2.44 in Group 1 and 1.14 and 2.66 in Group 2, compared to a range from 1.09 to 2.76 in the simulation data. The only discrepancy is that the learning model overpredicts the frequency of entry for low values of c and underpredicts it for high values of c.

Differences between Types. We turn next to a more detailed analysis of the simulation outcomes by examining the proportions of entry player type. The lower panel of Table III presents the proportions of entry summed across blocks and broken down by value of c and by type. These results are directly comparable to the observed proportions in the upper panel. The lower panel shows that the proportions of entry decrease monotonically by type. The increase in the observed proportion of entry from Type 1 to Type 2 (upper panel) is not captured by the learning model, but the general decreasing trend from Type 2 to Type 5 is captured quite well. Similar to the observed results in the upper panel showing that m increases in c for each type, the proportions in the lower panel exhibit the same monotonic trend.

Switches in Decision. Yet a more detailed analysis, which focuses directly on the adjustment process postulated by the learning model (Eqs. (1) and (2)), considers the switching data. Figure 3 exhibits the simulated mean number of switches in terms of player type and c value and compares them to the means of the observed frequencies (see Rapoport *et al.*, 1997, for the actual frequencies). Figure 3 shows that the learning model describes the distributions of the number of switches across the 10 values of c rather well. For Types 1, 2, and 3, the observed distributions are unimodal and negatively skewed; the predicted distributions exhibit the same shape. Similarly, the general tendency of the mean number of switches in Types 4 and 5 to increase in c is also captured by the learning model. The goodness of fit across the five types (lower right panel) is particularly impressive.

Figure 4 compares the observed (Table IV) and simulated numbers of switches across the 10 blocks of trials. It depicts these means graphically

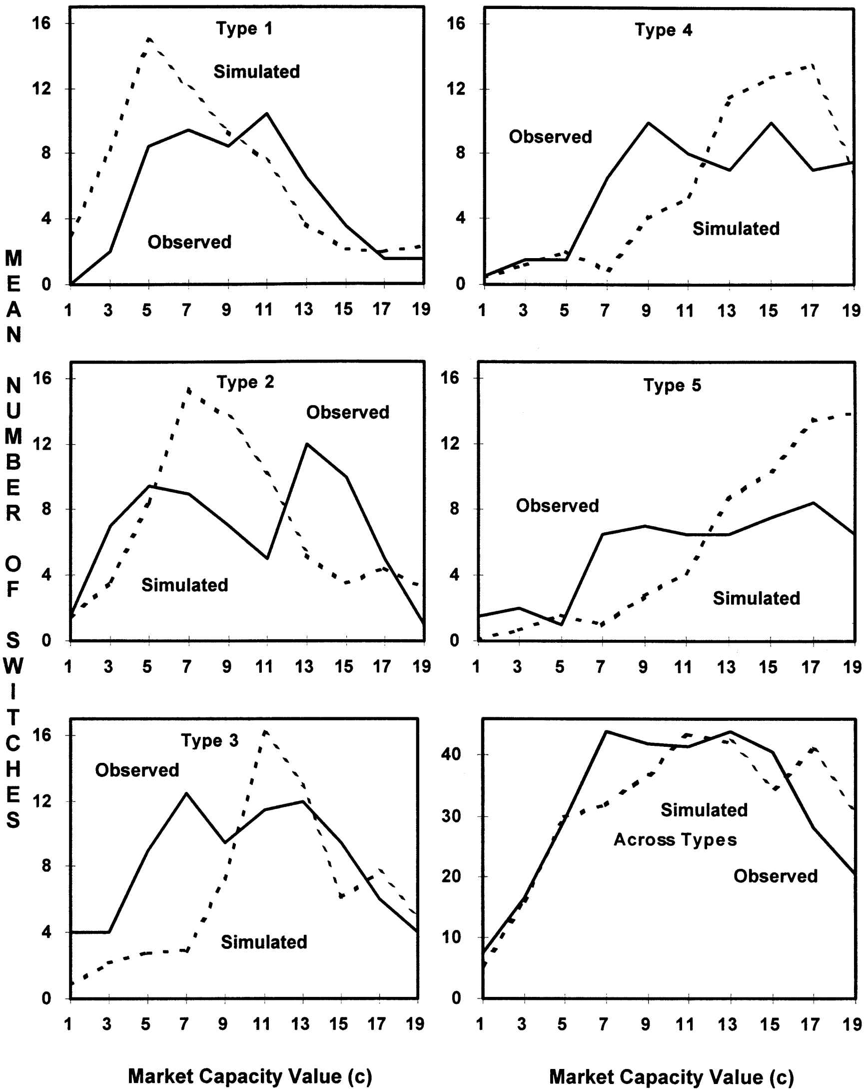

FIG. 3. Observed and simulated mean number of switches across market capacity values by player type in Experiment 1.

for each player type separately and also across types. The general tendency of the mean number of switches to decrease across blocks is captured by the learning model. However, for all five types the predicted mean number of switches exceeds the observed mean number of switches in Block 2 (i.e., in comparing the first and second presentations of the *c* values) and in four out of five cases also in Block 3. Similarly, the model predicts fewer switches in the last two or three blocks of trials. This consistent tendency to predict too many switches in Blocks 2–4 and too few in Blocks 8–10 is best seen by inspecting the observed and predicted mean numbers of switches across types in the lower right-hand panel of Fig. 4.

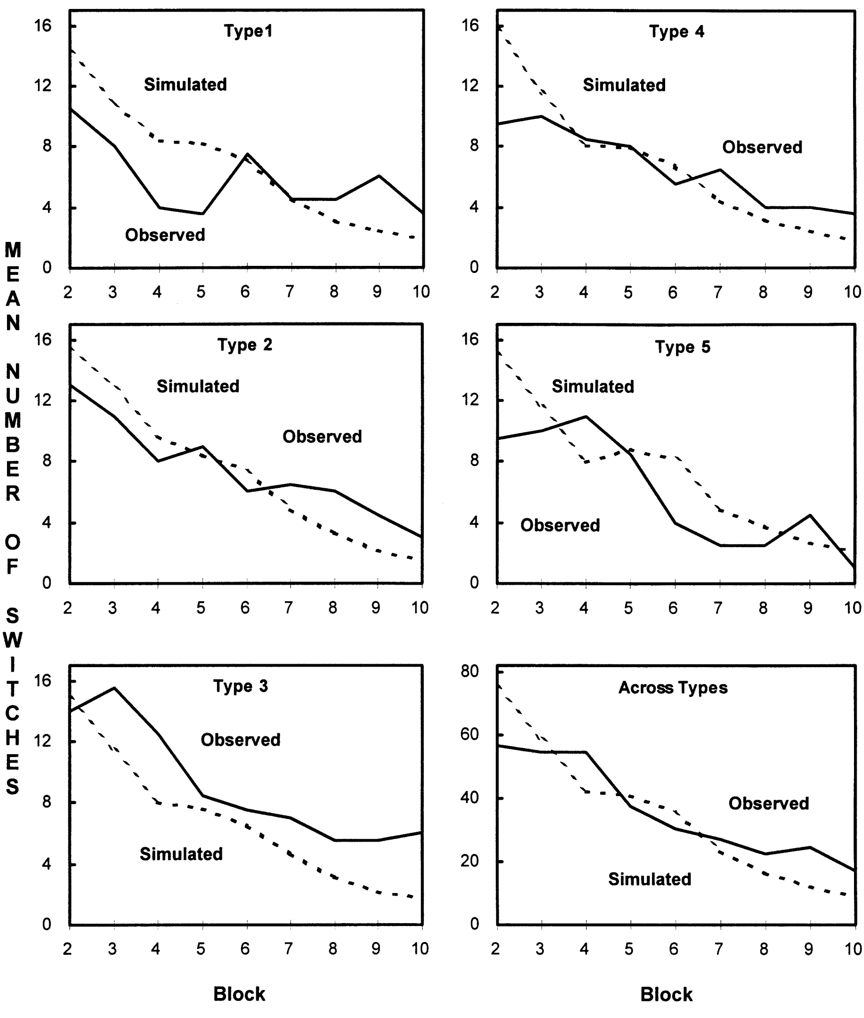

FIG. 4. Observed and simulated mean number of switches across blocks by player type in Experiment 1.

EXPERIMENT 2: THE MARKET ENTRY GAME WITH NO COMMON PRIOR

Although Experiment 1 includes 40 subjects, in actuality we have only two data points as the group rather than the individual player is the unit of analysis. Therefore, both to increase sample size and, in particular, to generalize the findings, we (Rapoport *et al.*, 2000, Exp. 1) replicated Experiment 1 with the only difference being that the subject instructions provided no information about the differential entry fees. Subjects were simply not informed that they were asymmetric. Although outcome feedback was provided at the end of each trial, exactly as in Experiment 1, this information could not be used by the subject to infer the parameters of the distribution of entry fees, nor could it even be used to infer the existence of different types of players. Presented below is only a brief summary of the major findings. For more details, consult Rapoport *et al.*, 2000.

Method

Similarly to Experiment 1, two groups of 20 subjects each participated in Experiment 2. With the exception of deleting information about the differential entry fees and the existence of different types, the instructions of the two experiments were identical. The conversion rate for the subjects of Group 1 was doubled (10 francs = \$3.00), whereas the one for Group 2 was the same as in Experiment 1 (10 francs = \$1.50). The mean payoff was \$51.02 for Group 1 and \$27.48 for Group 2.

Results

Table VII presents the means (across blocks) and standard deviations of entry decisions in Experiment 2. The results are presented separately for each group as well as across the two groups. The same format as that in Table II is used except that information about blocks is omitted. The results in Table VII exhibit the same remarkable coordination on the aggregate level observed in Experiment 1, which is accounted for surprisingly well by the equilibrium solution. Several statistical tests were conducted to compare the two groups, and none of them yielded significant differences (Rapoport *et al.*, 2000). Similarly (compare Tables II and VII), we observe no statistical differences in the mean frequency of entry for each value of c between Experiments 1 and 2.

When the entry decisions are analyzed separately by types, we observe patterns of behavior similar to those in Experiment 1 (Table III). The mean percentage of entry in Experiment 2 was 68.2, 57.6, 40.1, 26.2, and 17.6 for types 1, 2, 3, 4, and 5, respectively. Compare these results with the corresponding percentages for Experiment 1 shown in the bottom row of the

TABLE VII
Means and Standard Deviations of Observed Frequencies of Entry for Experiment 2

	Group I		Group II		Across Groups		
C	Mean	SD	Mean	SD	Mean	SD	**Prediction**
1	0.8	0.7	0.2	0.4	0.75	0.51	**0**
3	2.8	1.1	1.9	0.6	2.35	0.87	**2**
5	4.3	1.3	4.4	1.4	4.35	1.35	**4**
7	5.4	0.7	5.7	1.3	5.55	1.04	**5 (6)**
9	7.3	1.7	8.1	3.5	7.70	2.75	**7 (8)**
11	8.1	1.5	9.1	0.7	8.60	1.17	**9**
13	11.6	1.0	11.3	1.4	11.45	1.22	**11**
15	12.9	3.0	12.7	1.5	12.80	2.37	**12 (13)**
17	14.3	2.2	14.4	0.8	14.35	1.66	**14 (15)**
19	16.2	1.6	16.2	0.8	16.20	1.26	**16**
Total	83.8		84.0		83.90		**80 (84)**

upper panel of Table III. Note that the effect of type is more pronounced in Experiment 2 than in Experiment 1, and the anomalous result of type-2 players in Experiment 1 entering more frequently than type-1 players completely disappears.[6] We conclude from these comparisons that elimination of common knowledge of the distribution of types *improves* the coordination between types without affecting the overall coordination observed on the aggregate level.

Similarly to Experiment 1, we observe the same tendency of the percentages of (inferred) cutoff decision policies to increase with experience. The bottom panel of Table V compares these percentages (using t-test), separately for each block, between Experiments 1 and 2. None of the differences is statistically significant.

A major reason for conducting Experiment 2, rather than exactly replicating Experiment 1, was to test the learning model proposed above with a new set of data in an interactive situation where a common prior of the distribution of types can no longer be reasonably assumed. The learning model should apply to the data of Experiment 2, because it does not incorporate any assumptions about common knowledge of types or, indeed, common knowledge that the players are asymmetric. Extensive tests of the learning model reported by Rapoport *et al.* (2000) show that it performs as well as in Experiment 1. In particular, the learning model tends to underpredict the frequencies of entry for low values of c and overpredicts them for high values of c (compare the observed and simulated results for Experiment 2 reported in the right-hand column of Table III). With this exception, it

[6]We tentatively conclude that this anomalous result is due to the small sample size.

accounts for all the major regularities observed in the data. It accounts for the rate of convergence across blocks, for the differences between types, for the violations of the equilibrium predictions on the type level, and for the steady increase in the percentage of cutoff decision policies across blocks. The best fitting parameter values of w^+ and w^- are comparable in magnitude to those reported for Experiment 1. But the rate of convergence seems to be slightly faster, with values of the parameter d equal to 0.036 and 0.028 for Groups 1 and 2 of Experiment 2 compared to 0.020 across the two groups of Experiment 1.

DISCUSSION

A comparison of Experiments 1 and 2 of the present study with a previous market entry study with symmetric players (Rapoport *et al.*, 1998) shows that with asymmetry coordination slightly improves but not in any significant manner. In both studies, the pure strategy equilibrium solution accounts remarkably well for the aggregate results (Tables II and VII). This finding seems to be very robust, testifying to the descriptive power of the Nash equilibrium in tacit coordination games with relatively large groups with either symmetric or asymmetric players. However, when the players are asymmetric, we find no evidence for equilibrium behavior on the individual level and small but systematic departures from equilibrium play on the type level. Although the equilibrium solution serves as an excellent benchmark model, coordination is clearly achieved through learning. Our major findings are accounted for by a simple learning model that postulates changes in cutoff points due to the effects of positive and negative reinforcements whose magnitude decreases steadily over time. In testing this model, we followed the common procedure (e.g., see Roth and Erev, 1995, and Camerer and Ho, 1999) of finding the best-fitting parameter values, generating data for artificial subjects, and comparing simulated to actual behavior in terms of several statistics on both the aggregate and the individual levels. We have made no attempt to compare this model, which is customized to the present market entry game, to more general adaptive learning models that are reinforcement-based, belief-based, or combine principles of both reinforcement and belief updating (e.g., Camerer and Ho, 1999; Fudenberg and Levine, 1998; Erev and Roth, 1998; Stahl, 2000).

The learning model captures most of the basic behavioral regularities observed in the data. It accounts for the excellent coordination on the aggregate level (Table VI) as well as the monotonic increase in the correlations between m and m^*. It accounts for the tendency of the total number of entries to decrease with the entry fee. It also accounts for the monotonic increase in the proportion of entry across the values of the market capacity

c for each of the five types (Table III). It accounts for the total frequency of switches in decision between blocks; the nonmonotonic change in the total number of switches as a function of c, including the difference between type 5 and the other four types of players (Fig. 3); and the decrease in the number of switches across blocks (Table IV and Fig. 4). Finally, and perhaps most importantly, it accounts for the emergence of individual cutoff policies across time (Table V).

Compared to the observed results, the model postulates more erratic behavior in the first two or three blocks of trials and a faster learning rate. Although these discrepancies between simulated and observed results might be accounted for by increasing the number of parameters (e.g., assuming different discount rates for the two weight parameters, w^+ and w^-; conditioning the parameters w^+ and w^- on type; conditioning the parameters w^+ and w^- on player; and changing the assumptions about the effects of reinforcement), we opted not to do so. Parsimony of the learning model was judged to compensate for the potential gain in goodness of fit.

ACKNOWLEDGMENT

This work was supported in part by NSF Grant SBR-9512724 and in part by a grant from the Hong Kong Research Grants Council to the Hong Kong University of Science and Technology (Project CA98/99.BM01).

REFERENCES

Aumann, R., and Brandenburger, A. (1995). “Epistemic Conditions for Nash Equilibrium,” *Econometrica* **63**, 1161–1180.

Binmore, K. (1987). “Modeling Rational Players, I, II,” *Econ. and Philosophy* **3**, 9–55; **4**, 179–214.

Camerer, C. F., and Ho, T.-H. (1999). “Experience-Weighted Attraction Learning in Games: A Unifying Approach,” *Econometrica* **67**, 827–874.

Canning, D. (1995). “Learning and Social Equilibrium in Large Populations,” in *Learning and Rationality in Economics* (A. Kirman and M. Salmon, Eds.), pp. 308–323. Oxford: Blackwell.

Crawford, V. (1991). “An ‘Evolutionary’ Interpretation of van Huyck, Battalio, and Beil’s Experimental Results on Coordination,” *Games and Econ. Behav.* **3**, 25–59.

Crawford, V. (1995). “Adaptive Dynamics in Coordination Games,” *Econometrica* **63**, 103–143.

Crawford, V., and Broseta, B. (1995). “What Price Coordination? Auctioning the Right to Play as a Form of Preplay Communication,” unpublished manuscript. San Diego: University of California, San Diego.

Erev, I., and Roth, A. E. (1998). “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria,” *Amer. Econ. Rev.* **88**, 848–881.

Fudenberg, D., and Levine, D. K. (1998). *The Theory of Learning in Games*. Cambridge: MIT Press.

Kalai, E., and Lehrer, E. (1993). “Rational Learning Leads to Nash Equilibrium,” *Econometrica,* **61**, 1019–1045.

Marimon, R., and McGrattan, E. (1995). “On Adaptive Learning in Strategic Games, in *Learning and Rationality in Economics* (A. Kirman and M. Salmon, Eds.), pp. 63–101. Oxford: Blackwell.

Rapoport, A., Seale, D. A., Erev, I., and Sundali, J. A. (1998). “Coordination Success in Market Entry Games: Tests of Equilibrium and Adaptive Learning Models,” *Management Sci.* **44**, 119–141.

Rapoport, A., Seale, D. A., Parco, J. E. (2000). “Coordination in the Aggregate without Common Knowledge or Outcome Information,” Unpublished manuscript. University of Arizona, Department of Management and Policy.

Rapoport, A., Seale, D. A., and Winter, E. (1997). “Coordination and Learning Behavior in Large Groups with Asymmetric Players,” Report MKTG 97.098. Hong Kong University of Science and Technology, Department of Marketing.

Roth, A. E., and Erev, I. (1995). “ Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term,” *Games and Econ. Behav.* **8**, 164–212.

Stahl, D. O. (2000). “Rule Learning in Symmetric Normal-Form Games: Theory and Evidence,” *Games and Econ. Behav.* **32**, 105–138.

Sundali, J. A., Rapoport, A., and Seale, D. A. (1995). “Coordination in Market Entry Games with Symmetric Players,” *Organizational Behav. Human Decision Processes* **64**, 203–218.

Van Huyck, J., Battalio, R., and Beil, R. (1990). “Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure,” *Amer. Econ. Rev.* **80**, 234–248.

Van Huyck, J., Battalio, R., and Beil, R. (1991). “Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games,” *Quart. J. Econ.* **106**, 885–910.

Van Huyck, J., Battalio, R., and Beil, R. (1993). “Asset markets as an Equilibrium Selection Mechanism: Coordination failure, Game Form Auctions, and Tacit Communications,” *Games Econ. Behav.* **5**, 485–504.

Young, P. H. (1993). “The Evolution of Conventions,” *Econometrica*, **61**, 57–84.