An Experimental Study of Coordination and Learning in Iterated Two-Market Entry Games

Amnon Rapoport
Darryl A. Seale
Eyal Winter

Series No. MKTG 98.116
June 6, 1998

An Experimental Study of Coordination and Learning in Iterated Two-Market Entry Games

Amnon Rapoport
University of Arizona

Darryl A. Seale
University of Alabama
in Huntsville

Eyal Winter
Washington University

June 6, 1998

Please address all correspondence to

Amnon Rapoport
Hong Kong University of Science and Technology
Department of Marketing
Clear Water Bay
Kowloon, Hong Kong

Fax: 852-2358-2429
Phone: 852-2358-7719
email: mkamnon@usthk.ust.hk

Coordination paper

Abstract

Tacit coordination in large groups is studied experimentally in an iterated market entry game with complete information and multiple market capacities that are varied randomly from period to period. On each period, each player must decide independently whether to enter any of the markets, and if entering, which of the markets to enter. Across symmetric and asymmetric markets, we find remarkable coordination on the aggregate level, which is accounted for by the Nash equilibrium, together with considerable individual differences in frequency of entry and decision rules. With experience, the decisions of most of the subjects converge to cutoff rules with cutoff values on the combined market capacity that determine whether or not to enter but not which of the markets to enter. Both the aggregate and individual results are accounted for quite well by a reinforcement-based learning model that combines deterministic and probabilistic elements.

1. Introduction

Strategic interactions in which people or firms attempt to avoid congestion (e.g., choosing a driving route, picking a vacation spot, entering a new market) often display an incredible level of tacit coordination. However, the process by which such coordination is achieved cannot be determined by casual observation alone nor by sheer speculation. The present paper addresses this issue by studying tacit coordination in large groups playing an n-person noncooperative game called a two-market entry game (TMEG). We present and then analyze experimental results in which individuals face a decision whether to enter one of two markets or stay out. We iterate the game, changing the capacities of the two markets from period to period, and then develop and test an adaptive learning model to account for the dynamics by which symmetric players learn to coordinate their actions in order to solve the congestion problem.

Similar to the game of pure coordination studied by Ochs (1990) and the binary allocation game investigated by Meyer, Van Huyck, Battalio, and Saving (1992), the market entry game (MEG) was originally devised to investigate experimentally the coordination problem, which is at the heart of the inefficiencies that plague decentralized markets. In its original form, the MEG is played by a group N of n players. It begins with a public announcement of a positive integer c ($1 \leq c \leq n$) interpreted as the "commonly known capacity of the market". To enter the market, each player i ($i \in N$) is charged an entry fee h_i, which is private knowledge. Individual decisions are made simultaneously and anonymously. Once the value of c becomes publicly known, each player i must decide whether to pay his entry fee h_i and enter the market ($d_i = 1$) or stay out of it ($d_i = 0$). Communication between players is strictly forbidden. Individual payoffs are determined according to the payoff function $H_i(\mathbf{d})$:

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

where v, k, and r are real-valued constants, m ($0 \leq m \leq n$) is the (endogenously determined) number of entrants, and $\mathbf{d} = (d_1, d_2, \ldots, d_n)$ is the vector of individual (binary) decisions.

When $h_i = 0$ for all $i \in N$ (Erev & Rapoport, 1998; Rapoport, 1995; Rapoport, Erev, Seale, & Sundali, 1998a; Sundali, Rapoport, & Seale, 1995), the MEG is characterized by symmetric players, complete information, simultaneous play, zero entry fees, and an incentive to enter the market which decreases linearly in the number of entrants. If the values of h_i differ across players (Rapoport, Seale, & Winter, 1997), the effects of asymmetry between types of players (where types are determined by the entry fee) can be studied. In yet another variation (Camerer & Lovallo, in press), asymmetry is introduced by conditioning the individual payoff of an entrant on his rank relative to other entrants, where relative ranks are determined by skill at solving puzzles or answering trivia tests. In all these cases, the individual choice on each period is between taking a costless action which determines his payoff with certainty ("staying out"), or taking another action ("entering") whose outcome becomes less favorable the larger the number of players taking this same action.

When the n players are symmetric, the MEG has multiple pure-strategy Nash equilibria in which a subset of the players enter and the remaining players stay out, as well as a symmetric equilibrium in mixed strategies (see e.g., Sundali et al., 1995; Rapoport et al., 1998a). There are good reasons to expect failure in coordination particularly (as is the case in most of these studies as well as in the present experiment) if the MEG is iterated in time with the values of v, k, and r kept fixed across iterations and the value of c randomly changing from period to period. First, because the multiple equilibria in pure strategies are not Pareto rankable, there are no cues which could have helped the players to determine the particular subset of entrants. Focal points (Schelling, 1960) are irrelevant for achieving coordination in this class of games. Secondly, the equilibria in the MEG do not maximize group (and, subsequently, individual) payoff. Consequently, like in the iterated Cournot oligopoly game, our design allows for the emergence of tacit coordination in the repeated game which may increase considerably the payoffs, if <u>all</u> the players lower their frequency of entry. Finally, the coordination failure reported by Van Huyck, Battalio, and Beil (1990, 1991), who investigated coordination in a different class of coordination games with a relatively large number of symmetric players, provides no reasons to expect coordination success in the present MEG.

Previous Research. Given this prior expectation for coordination failure, the results reported by Rapoport (1995) were quite surprising (see Ochs, 1998, for a recent review). In his preliminary study, Rapoport reported a very high degree of coordination success, with the aggregate frequencies of entry converging to the equilibrium solution, together with considerable individual differences in decision rules and total frequency of entry that showed no sign of diminishing with experience. Subsequent experiments were designed to improve the methodology, replicate the results with different populations of subjects, and, most importantly, extend the original MEG in several different directions in an attempt to delineate the conditions under which coordination success in the MEG may or may not be achieved. Beside improving on the methodology of Rapoport, Sundali et al. used a between-subjects design to compare the original MEG either with or without trial-to-trial outcome information about the number of entrants and individual payoff. Using a between-subjects design, Rapoport et al. (1998a) manipulated the payoff for staying out, v, while keeping the values of the two other payoff parameters k and r fixed, in order to compare coordination behavior in the domains of gains ($v > k$) and losses ($v < 0 < k$). Erev and Rapoport (1998) went a step further by 1) using an experimental design in which the same value of c was presented repeatedly, and 2) manipulating the outcome information given to the players at the end of each period. Finally, Rapoport et al. (1997) introduced asymmetry into the game by charging different subsets ("types") of players differential entry fees; Seale and Sundali (1998) allowed for "cheap talk", where players can signal their intention whether or not to enter; and Camerer and Lovallo (in press) conditioned the individual payoff on pre-determined and commonly known ranks. The major findings of remarkable coordination at the aggregate level, which strongly support the Nash equilibrium solution, coupled with considerable individual differences in frequency of entry that do not diminish with practice and individual decision rules that vary considerably from one player to another, have proved to be very robust to these systematic manipulations of the payoff function, information structure, costless vs. costly signaling, and symmetry vs. asymmetry of players.

The present study extends this line of research in yet a different and major direction by introducing multiple markets. Simultaneously faced with multiple capacities, one for each market, each player in this new class of games must decide whether or not to stay out of the market, thereby avoiding competition with the other group members and ensuring a fixed (positive or negative) payoff, and then, if she decides to enter, determine which market to enter. Games of this type capture a coordination problem that commuters face when determining which mode of transportation to take in order to arrive as quickly as possible at their destination. For example, when trying to reach point B from point A when the roads between these two points are expected to be congested, a commuter may face a choice between taking a train ("stay out") and arriving at point B in a fixed and known amount of time, or driving her car ("enter") in which case the length of the drive may vary considerably depending on the road she chooses to take and the decisions made by the other commuters. Couples on Friday evenings face a coordination problem when having to decide between eating at home ("stay out") or going out ("enter"), in which case they must further choose between several restaurants of similar quality, when one of their major objectives is to minimize the combined driving and waiting time. Firms face a similar problem when deciding whether or not to enter one of several emerging markets, which may or may not be equally attractive in terms of their profitability which, in turn, depends on their capacity. Tacit coordination in this class of games is considerably more difficult than in the MEG, which includes a single market and, therefore, only a binary choice.

When the capacities of the multiple markets are varied randomly from period to period, as in the present experiment, the players are faced with a sequence of one-play coordination games with no information to form expectations that can support any of the multiple pure strategy Nash equilibria. A major objective of the present study is to determine if, and to what extent, the symmetric mixed-strategy Nash equilibrium can account for the behavioral regularities in this sequence of one-play coordination games. If, as our previous results of the simpler MEG studies lead us to believe, individual players do not mix their strategies in a way consistent with symmetric equilibrium play, a second major objective is to determine if coordination in this class of games, or any other consistent pattern of behavior that emerges with repeated play, can be achieved by some process of adaptation.

Section 2 presents the multiple-market entry game and characterizes its pure and mixed Nash equilibria. We begin the investigation of this class of games with the simplest case of two markets; extensions to a larger number of markets are obvious. Section 3 describes the experimental procedure and reports the results from three different groups of twenty players each. These results show remarkable coordination by each group and provide strong support for the mixed-strategy Nash equilibrium solution on the aggregate but not individual level. Section 4 presents and tests a reinforcement-based learning model, which incorporates the assumptions that the decision whether or not to enter is determined solely by the combined capacity of the two markets, whereas the decision of which market to enter is determined probabilistically. The major behavioral regularities observed in the experiment and reported in Section 4 are captured quite well by this learning model. Section 5 concludes with a discussion of the major findings.

2. A Multiple-Market Entry Game

Consider a group N of n symmetric players who participate in the following TMEG. At the beginning of each period t, $t = 1, 2, \ldots, T$, a pair of positive integers (c_A, c_B) is publicly announced. The values of c_A and c_B are interpreted as the "market capacities" of two independent markets A and B, respectively. In our experiment the capacities satisfy the constraint $1 \leq c_A + c_B \leq n$. Once the values of c_A and c_B are publicly announced, each player i ($i \in N$) must decide whether to stay out ($d_i = 0$), enter market A ($d_i = A$), or enter market B ($d_i = B$). Decisions are made anonymously, with no communication between players. The TMEG is iterated for T periods with different capacity pairs presented on each period. On each block of ten periods these capacity pairs are chosen randomly and without replacement from a prespecified set. Individual payoffs are determined by the payoff function $H_i(\mathbf{d})$, which remains unaltered for all T periods:

$$H_i(\mathbf{d}) = \begin{cases} v, & \text{if } d_i = 0 \\ k_A + r_A(c_A - m_A), & \text{if } d_i = A \\ k_B + r_B(c_B - m_B), & \text{if } d_i = B \end{cases}$$

where k_j, r_j (j = A, B), and v are real-valued commonly known constants that do not depend on the period t, m_j is the actual number of entrants in market j (j = A, B), and $\mathbf{d} = (d_1, d_2, \ldots, d_n)$ is the vector of individual decisions.

The experiment reported below examines a special case of the TMEG where $k_A = k_B \equiv k$ and $r_A = r_B \equiv r$. Therefore, with the exception of the possible difference between the capacities c_A and c_B, the two markets are indistinguishable. The pair of market capacities presented on each period t are sampled without replacement from the set

$$(C_A, C_B) = \{(2,2), (4,4), (6,6), (8,8), (10,10), (2,4), (2,8), (6,8), (6,12), (4,16)\}.$$

We shall refer to the game as <u>symmetric</u> if $c_A = c_B$, and <u>asymmetric</u>, otherwise. In the experiment reported below, five games are symmetric and five are not.

Equilibrium Solutions

In characterizing the pure-strategy equilibria, we distinguish between two cases depending on whether or not $v - k$ is a non-negative integer multiple of r, i.e., whether or not $v - k = qr$ for some non-negative integer q. If there exists no such integer, then there are $n!/m_A^*!m_B^*!m_N^*!$ pure-strategy equilibria, where m_A^* players enter market A, m_B^* enter market B, and $m_N^* = n - m_A^* - m_B^*$ players stay out. The values of m_j^* are determined from

$$m_j^* = \lfloor (rc_j + k - v)/r \rfloor, \quad j = A, B$$

where $\lfloor x \rfloor$ is the integral part of x. If there exists such an integer q, then there exist $n!/m_A^*!m_B^*!m_N^*!$ pure-strategy equilibria with m_j^* entrants and, in addition, $n!/(m_A^* - 1)!(m_B^* - 1)!(m_N^* - 2)!$ equilibria with $m_j^* - 1$ entrants, where

(1) $$m_j^* = (rc_j + k - v)/r, \quad j = A, B.$$

In the first class of equilibria, players who enter are indifferent between staying in or moving out, whereas in the second class of equilibria, non-entrants are indifferent between staying out or entering one of the markets.

There also exists a symmetric mixed-strategy equilibrium, where each player either enters market A, enters market B, or stays out with respective probabilities p_A, p_B, and $1 - p_A - p_B$. To establish this equilibrium, note that the expected number of entrants to market j, conditional on player i entering this market, is $Ej = 1 + (n - 1)p_j$. Player i's expected payoff from entering market j is, therefore, $k + r(c_j - Ej)$. In equilibrium this expression must equal v. This yields

$$(2) \qquad p_j = (r(c_j - 1) + k - v)/r(n - 1), \qquad j = A, B.$$

as the equilibrium probability of entering market j. As expected, np_j is very close to m_j*.

Two comments about the equilibria are in order. First, the values of m_j* and p_j only depend on the capacity of market j, not on the capacity of the other market. Thus, for example, the equilibrium probability of entering market A when $(c_A, c_B) = (4, 4)$ is the same as when $(c_A, c_B) = (4, 16)$. This implication will be tested below. Second, the equilibria—whether pure or mixed—do not maximize total group payoff. The players can increase the total group payoff considerably by lowering the frequency of entry. For example, suppose that k = 0, r = 2, v = 1, and n = 20, as in our experiment, and that $(c_A , c_B) = (8, 8)$. Then (Eq. 1), the pure-strategy equilibrium calls for seven players to enter market A, seven players to enter market B, and six players to stay out, for a total group payoff of 34. If, instead, only three players enter each market, the total group payoff increases to 76. Similarly, the mixed-strategy equilibrium probability of entering either of the two markets in this case (Eq. 2) is 0.342, and the associated expected total group payoff is vn = 20. However, if each player enters with probability $p_j/2 = 0.171$, the expected total group payoff is more than doubled. If they actually mix their strategies, there is a strong incentive for all the n players to tacitly cooperate by lowering their probability of entry, and thereby increasing their payoffs. This is a testable implication which we also examine below.

3. The Experiment

Method

Subjects. Sixty subjects divided into three groups of n = 20 subjects each participated in the experiment. Subjects in Groups 1 and 2 consisted of undergraduate and graduate students at the University of Arizona, who volunteered to take part in a two-hour computer-controlled experiment on economic decision making with payoff contingent on performance. The subjects in Group 3 were all MBA students who volunteered to participate in the experiment for the same monetary incentive. None of the subjects had participated in previous experiments on the MEG.

Procedure. The computer-controlled experiment was conducted at the Group Decision Making Laboratory at the University of Arizona. The procedure followed closely the one used in previous market entry experiments conducted by Rapoport et al. (1998a) and Sundali et al. (1995). Upon arriving at the laboratory, the subjects were seated at one of twenty computer terminals located in a big conference room. Communication between the subjects was strictly forbidden. The instructions were presented on individual screens in front of the subjects. Hard copies of these instructions were made available to the subjects.

Each subject was endowed with 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 money gained or lost during the experiment would be added or subtracted from their initial endowment. After they completed reading the instructions, the subjects played 100 iterations of the TMEG with payoff contingent on performance.

Each trial (period) consisted of two parts. In the first part, the two market capacity values c_A and c_B were displayed on the subjects' individual screens. To reduce the burden of calculation, the computer displayed on each period two tables, one for each market, listing the payoffs for all possible numbers of entrants, for the given values of c_j, ranging from 0 to 20.

In the second part of the trial, each subject was requested to type in his decision (A for entering market A, B for entering market B, and N for staying out). Once all the twenty subjects entered their decisions, each was informed of the values of m_A and m_B for the trial, his payoff for the trial, and his accumulated payoff (in francs) from the beginning of the experiment.

The ten pairs of market capacities (c_A, c_B) were sampled randomly and without replacement from the set (C_A, C_B) that included ten elements (trials 1 – 10 or block 1). The same sampling procedure was repeated for ten blocks. Thus, the subjects were presented with the same pair of market capacities (c_A, c_B) a total of ten times in a possibly different random order on each block.

At the end of the experiment, the subjects were paid their accumulated payoff for the entire experiment plus a \$5.00 show-up fee. The mean payoff across the three groups, excluding the show-up fee, was \$19.80. (Individual payoff would have been exactly \$20.00 if the subjects were to play the mixed equilibrium solution). The experiment lasted about 100 minutes.

Results

As the unit of analysis is the group, not the individual player, the study includes only three independent data points. All of the analyses that we conducted indicated no significant differences among the three groups. In particular, we observed no differences between the participants in Groups 1 and 2, most of them undergraduate students, and the MBA students of Group 3, who had acquired some knowledge of individual decision behavior and game theory in their MBA classes. We conducted several statistical tests comparing the three groups to one another in terms of the distribution of individual total number of entries across the 100 periods, frequency distributions of number of entries for each game summed across the ten blocks, rate of learning across blocks measured in terms of the difference between the (pure strategy) equilibrium and observed frequency of entry, and frequency of switches (see below) between the two markets on two successive presentations of the same game. All of these tests, which we report below, did not result in the rejection of the null hypothesis of no difference among the three groups.

In presenting the asymmetric games to the subjects, the low and high capacity markets were interchanged, with the labels A and B in five (randomly chosen) blocks, and B and A in the other five blocks. To facilitate presentation of the results, we shall henceforth identify A with the low capacity and B with the high capacity market.

Frequency of Entry: Aggregate Results. Table 1 presents the mean number of entries across blocks by game, and within each game by market. The results are shown separately for the symmetric (upper panel) and asymmetric (lower panel) games. In each panel, the top row labels the market (A or B), the second row lists the market capacity (value of c_j), and the next three rows present the observed mean number of entries for each group separately. The means and standard deviations of the observed number of entries computed across the three groups are presented in the next two rows. The bottom two rows of each panel present the mean and standard deviation of the number of entries predicted by the mixed-strategy equilibrium solution. Thus, for example, when $c_j = 6$, $p_j = 0.237$ (Eq. 2) and the mean and standard deviation for a group of 20 subjects are 4.74 and 1.90, respectively. The pure-strategy equilibria—which are not presented in the table--are always equal to $c_j - 1$.

--Insert Table 1 about here—

Inspection of rows 3, 4, and 5 in each of the two panels of Table 1 shows no differences between the three groups in mean number of entries across the ten blocks. We tested this hypothesis with the median test. For each of the five symmetric games, we counted for each subject the number of entries across the ten blocks and then compared the three frequency distributions in terms of their medians. We conducted the same test for each of the low capacity markets in the five asymmetric games and each of the high capacity markets, for a total of ten tests for this class of games. None of the 15 tests yielded significant differences among the three groups.

The major finding (Table 1) is remarkable coordination on the aggregate level. The mean observed frequencies of entries track the (pure or mixed) equilibrium values very closely across a wide spectrum of the combined market capacity values spanning the range from $c_A + c_B = 4$ to $c_A + c_B = 20$. Moreover, for the same value of c_j, there seem to be no significant differences between the two markets A and B. Also, as predicted by the equilibrium solution, for the same value of c_j there are no significant differences between symmetric and asymmetric markets.

The only systematic, though very small, deviations between observed and theoretical frequencies of entry (Table 1) are that low theoretical frequencies are slightly over-weighted and high theoretical frequencies are slightly under-weighted. Similar but larger discrepancies were reported in another coordination game experiment conducted by Ochs (1990). The discrepancies are very small; we noticed them only because they had been observed (in a more exaggerated form) in all previous studies of the MEG (Rapoport, Seale, & Ordonez, 1998b). Following the theoretical arguments of prospect theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992), and the ensuing experimental evidence (Tversky & Fox, 1995; Tversky & Wakker, 1995), Rapoport et al. (1998b) proposed a transformation of the mixed-strategy equilibrium probabilities of entry in the MEG by a nonlinear decision weight function $\pi(p)$, which is concave near zero, convex near 1, and satisfies the conditions $\pi(0) = 0$ and $\pi(1) = 1$. In particular, they proposed to apply the specific two-parameter subjective probability function proposed and tested experimentally by Lattimore, Baker, and Witte (1992), that provides a continuous approximation to the decision weight function of prospect theory. .

When only the probability of entry $p_j = p_j(c)$ and its complementary probability $1 - p_j$ are considered, the subjective probability function proposed by Lattimore et al. (1992) assumes the form

(3) $\pi(p_j) = \delta p_j^{\gamma}/(\delta p_j^{\gamma} + (1 - p_j)^{\gamma})$,

with positive parameters δ and γ. The parameter γ determines whether there is an inflection in the decision weighing functional and the direction of the inflection, whereas the parameter δ provides additional weight on the outcome probability p_j. In particular, if $\delta < 1$, p_j is under-weighted, and if $\delta > 1$, p_j is over-weighted. In addition to satisfying the basic features of the decision weight function proposed by prospect theory, the function π is sufficiently flexible to be consistent with alternative nonlinear expectation models of individual decision making under risk proposed by Quiggin (1982), Yaari (1987), and Karmarkar (1978).

Rapoport et al. (1998b) suggested combining the concept of mixed-strategy Nash equilibrium under its interpretation as an equilibrium in beliefs (see, e.g., Osborne & Rubinstein, 1994) with the nonlinear transformation of the probability scale proposed by prospect theory in order to enhance the descriptive power of the mixed-strategy equilibrium solution. Following this suggestion, we converted the observed frequencies of entry in Table 1 (by dividing each frequency by $n = 20$) to proportions, and then estimated for each group separately the values of the two parameters δ and γ in Eq. (3) which minimize the root mean square error (RMSE) between observed and predicted proportions. The results are displayed in Fig. 1. There are two plots for each group, each including all twenty data points (ten games by two markets). Each data point—a maximum likelihood estimate of the probability of entry—is based on 200 observations. The left side of each panel plots the observed proportions, denoted by P, against the equilibrium probabilities of entry $p_j = p_j(c)$. The right side of each panel displays the transformed values of p_j, $\pi(p_j)$, against the observed proportions.

--Insert Fig. 1 about here—

Inspection of the left-hand parts of the three panels exhibits the tendency we observed earlier of over-weighing low equilibrium probabilities and under-weighing high equilibrium probabilities of entry. This very small but systematic deviation appears in all three groups. Inspection of the right-hand panels shows that the transformed probabilities $\pi(p_j)$ remove this bias and improve the goodness of fit: the RMSE scores for Groups 1, 2, and 3 are 0.010, 0.009, and 0.008. These scores are considerably smaller, by a factor of about four, than the RMSE scores associated with the linear fit of the equilibrium probabilities in the left-hand panels in Fig. 1; the latter scores assume the values 0.045, 0.044, and 0.037 for Groups 1, 2, and 3, respectively.

If $\delta = \gamma = 1$, then Eq. (3) reduces to the frequency interpretation of probability and is, therefore, consistent with the expected value model. In the present study, the pair of best-fitting parameters (δ, γ) assumed the values (0.942, 0.794), (0.876, 0.783), and (0.911, 0.845) for Groups 1, 2, and 3, respectively. For all three groups, $\delta < 1$, consistent with the subcertainty effect ($\pi(p) + \pi(1 - p) < 1$) postulated by prospect theory.

<u>Frequency of Entry: Individual results.</u> Because the results presented in Table 1 are averaged across subjects, they conceal information about individual differences. Our first attempt to assess

individual differences was by counting the number of times (out of a maximum of 10) that each subject stayed out of the market on the ten iterations of the same TMEG. The frequency distributions of the individual total number of staying out decisions are presented in Table 2 for each of the five symmetric games separately.

--Insert Table 2 about here--

Table 2 describes the reason for the over-entry in markets with relative low capacities and under-entry in markets with relative high capacities. Whereas the majority of the subjects (32 out of 60) never entered either of the two markets in Game (2, 2), 12 subjects (20%) entered it at least 4 times. In equilibrium, practically no subject should enter either of the two markets in this game more than four times. With the number of entries truncated from below by zero, we observe a negatively skewed distribution with a mean (1.68) that exceeds the median. Similar results hold for Game (4, 4). The frequency distribution for Game (10,10) is also skewed but in a positive direction, with a similar difference between the mean (8.67) and median (10). Similar results hold for Game (8, 8).

Perhaps more important is the failure of the mixed-equilibrium solution to account for individual results. Table 2 also shows the mixed-strategy equilibrium frequencies of staying out in the ten iterations of the game. In each of the five symmetric games, the difference between the observed and predicted frequency distributions—tested by the Chi Square test—is highly significant ($p < 0.01$). Whereas the aggregate results are described extremely well by the mixed-strategy equilibrium solution, particularly after the equilibrium probabilities of entry are transformed non-linearly by the function π, individual subjects as a rule do not mix their strategies with the equilibrium probability. Rather, compared to this solution, there are many subjects who either enter too frequently or too infrequently, depending on the values of c_j. Similar results, not presented here, also obtain for the asymmetric markets.

<u>Individual Total Number of Entries</u>. Additional information about the magnitude of individual differences observed in all three groups is presented in Table 3. This table shows the individual total number of entries across all ten games and all ten blocks (with a minimum of 0 for a player who

never enters and a maximum of 100 for one who always enters). The individual differences in this table are much too large to be accounted for by the symmetric mixed-strategy equilibrium; across the three groups, the total number of entries varies from 0 to 100. However, in agreement with the aggregate results in Table 1, the three means of the individual total number of entries in Groups 1, 2, and 3 (56.30, 54.65, and 53.75) do not differ significantly from one another ($F_{93, 20} < 1$), nor does any one of them differs significantly from the equilibrium (pure or mixed) total number of entries.

--Insert Table 3 about here—

Block Effects in Frequency of Entry. Two additional analyses, each intended to explore the dynamics of play, are described next. Table 4 presents the total number of entries for each of the two markets by block. The results are organized by symmetry or asymmetry of the game, and for each game by the particular market (A or B). Thus, for example, when Game (2, 2) was presented for the first time (in block 1), 10 of the 60 players entered market A and 3 entered market B (47 players stayed out). The corresponding frequencies in Game (10, 10) of block 1 are 23 and 32 (with 5 players staying out). The mean numbers of entries across blocks (third column form the right) are the same as those displayed in Table 1, once they are divided by 3.

--Insert Table 4 about here—

With the exception of block 1, there is hardly any indication in the table for systematic improvement in coordination due to experience. For each of the two types of game (symmetric and asymmetric) and each of the two markets, we measured the deviation between predicted and observed frequencies by a root mean square error:

$$RMSE = [\Sigma(f_o - f_p)^2/5]^{1/2},$$

where f_o and f_p denote the observed and predicted frequencies, respectively. The RMSE scores are displayed in Table 4 below the ten columns; they do not show any systematic tendency to decrease across blocks. Spearman correlations (not shown here) were computed for each market and game separately between the ten observed frequencies and the block number (1 through 10). Only 3 of the 20 correlations were significant ($p < 0.05$, two-sided t-test).

Switches in Decision. On the individual level, we observe no such consistency in the decisions. Rather, across iterations of the same game, individual subjects continue changing their decisions whether or not to enter and, if deciding to enter, which market to enter. To measure this block-to-block change in decision, consider any of the ten games presented on block t and code the subject's decision as N, if he stays out; A, if he enters market A; and B, if he enters market B. When the same game is presented (after ten periods on the average) on block t + 1, the subject is faced with the same decision task. We say that he switches his decision, if, for the same game, he responds k on block t and k' on block t + 1, where $k \in \{N, A, B\}$, $k' \in \{N, A, B\}$, and $k \neq k'$. We can construct a 3 by 3 (k by k') transition matrix separately for each subject and each pair of adjacent blocks. An analysis of the transition matrix between blocks t and t + 1 allows measuring the consistency of the subject's decision rule over iterations. If, for example, a subject uses the same deterministic decision rule for the entire duration of the game, no switches will be recorded for this subject; rather, only the cells on the major diagonal of the 3 by 3 transition matrices will include positive elements.

Table 5 presents the 3 by 3 transition matrices with the frequencies summed across the 60 subjects. The matrices are presented by game, and for each game they are summed across the sixty subjects in the three groups and across the nine possible transitions from block t to block t + 1 (t = 1, 2, … , 9). Consequently, there are 9 x 60 = 540 observations in each transition matrix. The frequencies of switches appear in the six off-diagonal cells of the transition matrix. For example, in Game (2, 2) the subjects switched their decisions from staying out to entering market A between successive presentations of the same game on blocks t and t + 1 a total of 18 times, whereas in 21 times they switched their decision from entering market A to staying out. Below each table is the proportion of switches in the game, which is simply the total frequencies in the off-diagonal cells divided by 540.

--Insert Table 5 about here—

The analysis of the frequency of switches reveals a few patterns. In the symmetric games, the frequency of switches between markets A and B increases in the value of c_j: the combined frequencies of switches between these two markets are 21, 38, 59, 109, and 140 for Games (2, 2), (4, 4), (6, 6), (8, 8), and (10, 10), respectively. The number of switches between N and A is about the same as from N to B in all five symmetric games, showing no particular bias in favor of one of the markets. In the asymmetric games, we observe more switches in Game (6, 8) than in any other game. Particularly surprising is the high rate of switches in Games (6, 12) and (4, 16)), where one market is considerably more attractive than the other. These patterns are of particular importance for assessing the performance of any learning model that attempts to capture the dynamics of play.

Table 6 presents the same transition frequencies by block (rather than by game) summed across the ten games (rather than blocks) and all sixty subjects. As there are only nine transitions, the frequencies in each transition matrix sum up to 600. The major finding in this table is a gradual and very slow decline in the proportion of switches across blocks from 0.42 in the transition from block 1 to 2 and 0.33 in the transition from block 2 to 3 to 0.267 to the final transition from block 9 to block 10. This decline is mainly due to the transitions between N and either A or B. The number of switches between the two markets, which assumes the values 73, 70, 79, 73, and 71 in the last five transition matrices, hardly changes across blocks. These results suggest that with experience, more subjects might have converged to using decision rules which allow for no switches between entering and staying out. We shall test this hypothesis below. We have observed the same convergence to decision rules formulated in terms of a fixed cutoff on the market capacity dimension in the study of Rapoport et al. (1997) that included only a single market. This finding of gradual decline in the frequency of switches over time should also be accounted for by any learning model that attempts to capture the dynamics of play.

--Insert Table 6 about here—

Analysis of the frequency of switches on the individual level shows considerable differences between subjects. Recall that the maximum number of switches per subject is 9 x 10 = 90. In actuality, the ranges of total number of switches per subject were [3, 57], [4, 62], and [0, 43] for Groups 1, 2, and 3, respectively. The means of the number of switches per subject assumed the values 30.40, 29.05, and 28.15 with respective standard deviations of 5.88, 5.81, and 5.29, for Groups 1, 2, and 3, respectively. As in all of our previous comparisons of the three groups, the hypothesis of equal means could not be rejected by a one-way ANOVA ($F_{3, 20} < 1$).

Cutoff Decision Rules. Each subject must first decide whether or not to stay out and then, if deciding to enter, which market to enter. Switches in decision between presentations of the same game on two successive blocks might be due to the first, the second, or both decisions. As mentioned earlier, the results of our previous MEG study (Rapoport et al., 1997) suggest that most subjects slowly converge to using a deterministic decision rule as far as the first choice is concerned. However, if they decide to enter in the present study, then their choice of which market to enter may change as a function of their previous gains or losses in the same or different games. We tested the hypothesis that individual decision policies converge to a cutoff decision rule of the type "enter if and only if $c_A + c_B \geq c^*$", where c^* is an individual cutoff value that may vary from one subject to another. This rule does not specify which market to enter, nor does it distinguish between symmetric and asymmetric games.

Even if the subject adheres to a cutoff decision rule, the ten games in the present experiment do not allow to determine c^* uniquely. If, for example, the subject always stayed out when $c_A + c_B \leq 6$ (Games (2, 2) and (2, 4)) but entered on all other games for which $c_A + c_B \geq 8$ (e.g., Games (2, 8), (4, 4), (6, 6)), then c^* may take any value in the interval (6, 8]. We estimated c^* to be the upper limit of the interval, i.e., 8 in the present example.

All ten decisions of 35 of the 60 subjects (58.3%) in block 10 were accounted for, without a single violation, by this cutoff decision rule. In particular, the number of subjects (out of 20) whose last ten decisions could be accounted for by this rule was 11, 11, and 13, in Groups 1, 2, and 3, respectively. The decisions of ten more subjects violated this rule only once (e.g., entering on Games (6, 12) and (10, 10), but staying out on Game (4, 16)). Table 7 presents the frequency distribution of the estimated cutoff values c* for these 35 subjects. Thus, 7 of these 35 subjects always entered one of the two markets in all ten periods of block 10, 10 subjects entered if $c_A + c_B \geq 6$, and so on. Three of the 35 subjects never entered; the estimated cutoffs for these subjects were arbitrarily set at 21.

--Insert Table 7 about here—

As if by some magical hand, the estimated cutoffs in block 10 are distributed more or less uniformly on the set of values that the combined market capacity $c_A + c_B$ assumed in the present study. The null hypothesis of a uniformly distributed cutoff values for this subset of subjects could not be rejected by a Chi Square test ($p > 0.25$). If all the sixty subjects were to use cutoff rules with cutoff values sampled from this distribution, the Nash equilibrium solution would have been supported.

Taken together, the results reported above leave us with essentially the same puzzle that we encountered in the previous MEG studies. On the one hand, the mixed-equilibrium solution describes the aggregate results extremely well There is no indication for deviations from equilibrium play in the direction of fewer entry decisions which would result in considerably higher group payoffs. On the other hand, we observe considerable individual differences in the total number of entries (Table 3), frequency of switches (Tables 5 and 6), and estimated number of cutoff policies in the last block (Table 7) that cannot be accounted for by the symmetric mixed-strategy equilibrium, together with evidence for learning reflected in the decrease in number of switches across blocks.

4. **Reinforcement-Based Adaptive Learning**

The behavioral regularities recorded above present a challenge, which we try to meet by specifying and then testing an adaptive learning model based on the principles of reinforcement learning. This model incorporates the basic principles underlying a previous adaptive learning model (Rapoport et al., 1997) proposed to account for the results of the MEG and extends them to the TMEG. We emphasize that our goal here is not to find the best fitting model; as there is an obvious trade-off—which many learning models opt to ignore—between goodness of fit and number of free parameters incorporated by the model, goodness of fit can always be improved by adding more parameters. Rather, our goal is to identify the minimum cognitive sophistication which is required to account simultaneously for the considerable individual differences in decision rules and coordination success on the aggregate level that we reported above. In pursuing this goal, strong emphasis is placed on parsimony of the model and psychological interpretability of its parameters. We accept the premise (e.g., Cooper & Feltovich, 1996) that the simplest model in terms of cognitive sophistication should be tried first, and that models postulating higher levels of cognitive sophistication should only be employed as the first ones fail.

We make no attempt to explicate the concept of "cognitive sophistication" (see Cooper & Feltovich, 1996, for such an attempt). Rather, we only make the intuitive claim that reinforcement-based learning models are cognitively less demanding than belief-based learning models. The latter typically require the players to set up a probability distribution over the opponents' actions, update their beliefs, and choose a best-response strategy which maximizes their expected payoff, given their beliefs (Camerer & Ho, 1996). In contrast, reinforcement-based learning models, regardless of the way the reinforcement mechanism is assumed to work (e.g., probabilistically vs. deterministically), make weaker demands on the rationality and reasoning ability of the players (Roth & Erev, 1995). They assume that players only care about the payoffs strategies yielded in the past, not about the history of play that resulted in these payoffs. Subjects are simply supposed to recall what worked well for them in the past, and then do it more frequently in the future.

Model Specification

The model that we propose to account for individual learning in the TMEG combines elements from the reinforcement-based probabilistic approach to learning in games advocated by Roth and Erev (1995) and the deterministic approach to learning in the MEG (Rapoport et al., 1997). Whereas the decision whether or not to enter any of the markets in the TMEG is assumed to be made deterministically in terms of a cutoff point on the combined market capacity $c_A + c_B$ (see evidence in Table 7), the decision which of the two markets to enter is assumed to be determined probabilistically. Specifically, the model assumes that the decision on each period is actually broken into two stages. In the first stage, the player decides whether or not to stay out: he enters if and only if the combined market capacity exceeds some individual cutoff value. This cutoff value is updated individually in a deterministic manner: it moves up if both markets are overcrowded, moves down if both are undercrowded, and remains unchanged, otherwise. In the second stage, if he decides to enter, a decision is made as to which market to enter. For asymmetric markets, the model assumes that this decision is made probabilistically in terms of two propensities that reflect the attractiveness of the high capacity and low capacity markets. These propensities are updated when one market is overcrowded and the other is undercrowded. In these cases, the propensity of the market actually entered increases if it was undercrowded and decreases if it was overcrowded. When the two markets are symmetric, each is assumed to be entered with equal probability. These assumptions are formalized below.

Preliminaries. We maintain the distinction between symmetric ($c_A = c_B$) and asymmetric ($c_A \neq c_B$) TMEG. For the latter class of games, we distinguish between low (L) and high (H) markets, where $c_L < c_H$, $L \in \{A, B\}$, $H \in \{A, B\}$, and $L \neq H$.

The Decision Process. At the beginning of each period t, t = 1, 2, ... , T, each player $i \in N$ is assumed to be characterized by three non-negative numbers (the subject index i is omitted): c_t, $P_t(L)$, and $P_t(H)$.

The decision whether or not to enter a market is made as follows:

If $c_t \leq c_{L,t} + c_{H,t}$, enter on period t,

If $c_t > c_{L,t} + c_{H,t}$, stay out on period t.

If condition $c_t \leq c_{L,t} + c_{H,t}$ is satisfied, then the decision which market to enter depends on whether or not the game is symmetric.

If the game is symmetric, each player i is assumed to enter either of the two markets with equal probability. If the market is asymmetric, he is assumed to enter either market L or market H with respective probabilities $q_t(L)$ and $q_t(H)$, where

$$q_t(L) = P_t(L)/(P_t(L) + P_t(H))$$

and

$$q_t(H) = P_t(H)/(P_t(L) + P_t(H)).$$

Note that $q_t(L) + q_t(H) = 1$.

Outcome. If the game is symmetric, then set arbitrarily A = L and B = H.

At the end of period t exactly one of four outcomes obtains:

Case 1: $c_{H,t} \le m_{H,t}$ and $c_{L,t} \le m_{L,t}$ ("both markets are overcrowded").

Case 2: $c_{H,t} > m_{H,t}$ and $c_{L,t} > m_{L,t}$ ("both markets are undercrowded").

Case 3: $c_{H,t} \le m_{H,t}$ and $c_{L,t} > m_{L,t}$ ("market H is overcrowded and market L is undercrowded").

Case 4: $c_{H,t} > m_{H,t}$ and $c_{L,t} \le m_{L,t}$ ("market H is undercrowded and market L is overcrowded").

In these four expressions, $m_{H,t}$ and $m_{L,t}$ denote the actual frequencies of entry into markets H and L on period t, respectively, where $0 \le m_{H,t} + m_{L,t} \le n$.

Updating of the Propensities. In updating the two propensities $P_t(H)$ and $P_t(L)$ and the individual cutoff point c_t, the model distinguishes between two cases, depending on whether or not player i stayed out on period t.

Supposing that player i entered one of the two markets on period t. Denote the market he entered by X(t) and the one he did not enter by Y(t), where $X(t) \in \{H, L\}$, $Y(t) \in \{H, L\}$, and, of course, $X(t) \neq Y(t)$. The three parameters that characterize each player are assumed to be updated according to which of the four cases specified above obtained on period t.

If Case 1 obtains, then

$$c_{t+1} = c_t + w^+_t[v - (k + r(c_{X(t)} - m_{X(t)}))],$$

$$P_{t+1}(H) = P_t(H),$$

$$P_{t+1}(L) = P_t(L).$$

In words, if both markets were "overcrowded" on period t (and, consequently all subjects who entered loss money), then the individual cutoff increases in proportion to the difference in payoff between the outcome for staying out and the (negative) payoff for entering, whereas the two propensities $P_t(H)$ and $P_t(L)$ remain intact. Without specifying upper limit on c_{t+1}, this updating rule allows for the possibility that $c_{t+1} > c_H + c_L$, resulting in no entry for any pair of capacities on period t + 1.

If Case 2 obtains, then

$c_{t+1} = \max\{c_t + w^-_t[v - (k + r(c_{X(t)} - m_{X(t)}))], 0\}$,

$P_{t+1}(H) = P_t(H)$,

$P_{t+1}(L) = P_t(L)$.

If both markets are "undercrowded" on period t (and, consequently, all subjects who entered made money), then the individual cutoff value is assumed to decrease, whereas the two propensities remain unaltered. The equation for updating c_t has the same form as in Case 1, with the only exception that c_{t+1} is assumed to be bounded from below by zero.

If either Case 3 or Case 4 ("mixed" market) obtains, then

$c_{t+1} = c_t$,

$P_{t+1}(X(t)) = \max\{1, P_t(X(t)) + z_t[r(c_{X(t)} - m_{X(t)}) - r(c_{Y(t)} - m_{Y(t)})]\}$,

$P_{t+1}(Y(t)) = P_tY(t)$.

In these two latter cases, where one market is overcrowded and the other is undercrowded, only the propensity of the market actually entered is assumed to change; the other propensity and the cutoff value remain unchanged. The term in square brackets is the difference between the payoffs associated with the two markets.

Supposing that the player <u>stayed out</u> on period t. Then,

$$c_{t+1} = \begin{cases} c_t - w^+_t G_t, & \text{if } G_t < 0, \\ \max\{c_t - w^-_t G_t, 0\}, & \text{if } G_t \geq 0, \end{cases}$$

$P_{t+1}(H) = P_t(H)$,

$P_{t+1}(L) = P_t(L)$,

where G_t is the expected payoff (positive or negative) *foregone* by not entering, which is determined from

$G_t = q_t(H)[k + r(c_{H,t} - m_{H,t-1}) - v] + q_t(L)[k + r(c_{L,t} - m_{L,t-1}) - v]$.

In words, if player i stayed out on period t, the propensities for entering either of the two markets are assumed not to change, whereas the cutoff value changes according to the sign of the expected payoff foregone by not entering on the previous period. The max operator for the case of $G_t \geq 0$ assures that the cutoff value remains non-negative.

Discounting of the Weight Parameters. The effects of the three weight parameters w^+, w^-, and z, are assumed to be discounted with time:

$$w^+_{t+1} = (1 - d)w^+_t,$$

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

$$z_{t+1} = (1 - d)\, z_t,$$

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

In summary, the learning model has four parameters. Two of these (w^+ and w^-) determine the proportional change in the cutoff point for overcrowded and undercrowded markets, respectively; a third parameters (z) determines the rate of change in the propensity to reenter the same market when the markets on the previous period were "mixed", and the fourth determines the discounting of the effects of the former three parameters. In its present form, the model assumes that the three parameters w^+, w^-, and z have the same values for all the n players and that all three of them are discounted at the same rate; these assumptions could be relaxed. The discounting assures that, after considerable experience with the game, the decision whether or not to enter one of the two markets (rather than which market to enter) will be determined by a cutoff rule on the combined market capacity.

Initialization. The model allows for no individual differences in the rate of learning. The only differences between individual players are due to the initial cutoff values $c_{t=1}$. To express the fact that individual differences exist before the experiment commences, we assume that the initial cutoff values, one for each of the n players, are randomly drawn from a prior distribution. The initial cutoff values can be either drawn from some prior distribution which can be justified on prior grounds or estimated from the actual decisions of the players on the first block of trials. As for the initial propensities for entering the markets, we set them up at $P_1(H) = P_1(L) = 1$. Because of the evidence reported above for slow learning, we restricted the discount factor by setting $d \leq 0.05$.

Model Testing

Testing of the learning model proceeded as follows. For each of the sixty subjects in Groups 1, 2, and 3, we first estimated the initial cutoff value $c_{t=1}$ from his decisions in the first ten periods of block 1. To do so, for a fixed value of c_1 a deviation score of $|c_1 - (c_A + c_B)|$ was computed for each period in block 1 in which the subject (erroneously) entered when $c_1 > (c_A + c_B)$ or stayed out when $c_1 \leq (c_A + c_B)$. A simple search procedure was used to determine the best fitting value of c_1, namely, the one that minimizes the sum of absolute deviations form c_1 across the ten periods of block 1. When the estimation procedure yielded a range of values rather than a single value—this was the case for most of the subjects-- the median integer value of this range was selected. This estimation procedure yielded three sets of twenty initial cutoff values, one for each group. The empirical distribution of these initial cutoff values computed across the sixty subjects of the three groups is displayed in Fig. 2. With the exception of six cases of $c_1 = 0$, all the other cutoff values are between 8 and 20. If all the subjects were to use cutoff rules with these values consistently, their decisions would not have resulted in equilibrium play.

--Insert Fig. 2 about here—

Next, for each group of subjects separately, we estimated the best fitting values of the four parameters $w^+_{t=1}$, $w^-_{t=1}$, z_1, and d in the following way. With the aid of a high speed computer, we conducted a systematic search of a four-dimensional grid of parameter values, and for each combination of the four parameter values we simulated the results of each group member across the 100 periods using the leaning model equations for determining a decision and updating the cutoff value and the two propensities. At the end of each simulation, across the twenty players of each group, we computed the criterion measure of goodness of fit

$$Q = \sum_{t=1}^{T} \left[(m_{H,t}^{(s)} - m_{H,t})^2 + (m_{L,t}^{(s)} - m_{L,t})^2\right],$$

where $m_{H,t}^{(s)}$ and $m_{L,t}^{(s)}$ denote the simulated number of entrants for the small and large markets on period t, respectively, $m_{H,t}$ and $m_{L,t}$ are as defined above, and summation is taken across all the 100 periods. After finding combinations of parameter values with relatively small values of Q, the search proceeded with a "fine grain" four-parameter grid. We varied the range or precision level of each of the four parameters systematically and independently until we identified the best fitting combination of parameter values when no further reduction of Q was possible. Random numbers were generated and used to determine which of the two markets to enter according to the values of the two propensities $P_t(H)$ and $P_t(L)$ when the game was asymmetric. When it was symmetric, each market was chosen with equal probability of 0.5. The resulting parameter values, one set for each group, are shown in Table 8.

--Insert Table 8 about here—

Using this search procedure, we generated data for three groups each including twenty "simulated" players (called "stat-rats" by Bush & Mosteller, 1955), with parameter sets that vary from one group to another but not from one group member to another. Due to the probabilistic nature of the decision process when determining which of the two markets to enter (but not whether or not to enter), different simulation runs which start with the same twenty initial cutoff values and use the same set of parameter values may yield different results. To simplify the comparison between simulated and observed decisions, we generated only one set of simulated results for each group.

The comparison between observed and simulated decisions was, therefore, conducted separately for each group. Table 9 presents the means and standard deviations of the simulated number of entries. The means and standard deviations are computed across all ten blocks. The results are presented in the same format as in Table 1. Inspection of Table 9 suggests that the learning model yields remarkable coordination for each group and simulates the observed mean number of entries rather closely. This is evident from Fig. 3, which plots for each group separately and for the total across the three groups the simulated against the observed mean number of entries for each of the two markets in each of the ten games. The linear correlations between simulated and observed means for Groups 1, 2, 3, and for the total are 0.95. 0.95, 0.93, and 0.96, respectively ($p < 0.01$ in each case), and the corresponding slopes of the regression line of the simulated on the observed means (0.845, 1.140, 0.917, and 0.974) do not differ significantly from unity ($p < 0.01$ in each case). These results provide strong support for the learning model. Comparison of Tables 1 and 9 together with a careful inspection of the four plots in Fig. 3 shows that the only major discrepancy between simulated and observed mean number of entries is due to Game (6, 8), where the learning model underestimates the frequency of entry in market A ($c_A = 6$) and overestimates it in market B ($c_B = 8$). Comparison of Tables 1 and 9 further shows that the standard deviations of the simulated number of entries are only slightly higher than the observed values: whereas the standard deviations for the observed entry decisions range between 0.90 and 1.84, those for the simulated entry decisions range between 0.79 and 2.86.

--Insert Table 9 and Fig. 3 about here—

Our next comparison between observed and simulated number of entries concerns the effect of practice on decision. Table 10 presents the simulated number of entries for each group separately by game and block. The format of the table is the same as Table 4. Comparison of the standard deviations in Tables 4 and 10 shows larger block-to-block variations in the simulated than the observed frequencies of entry. A line-by-line comparison of these two tables shows that the simulated results follow the observed results rather closely. For both sets of data, in most of the markets there seem to be no systematic changes in frequency of entry across blocks. When the observed frequencies increase across blocks (market B of Games (2, 4) and (6, 12)), so do the simulated frequencies, and when they decrease across blocks (market A of Games (2, 4) and (6, 12)) again so do the simulated results. The only notable difference is in market A of Game (4, 4) where the trends in observed and simulated frequencies of entry across blocks have opposite directions.

--Insert Table 10 about here—

Yet, another comparison between the two sets of data concerns the individual total number of entries across the 100 periods of the experiment. One would expect the learning model to account not only for the mean trends but also for the individual differences. Figure 4 displays side by side the frequency distributions of observed (presented in Table 3) and simulated individual total number of entry decisions. Although the simulated players seem to be slightly more heterogeneous than the observed subjects, with more simulated players having either very low (0 – 20) or very high (81 – 100) individual total number of entries, the difference between the two frequency distributions is not significant ($\chi^2(8) = 3.09$, $p > 0.8$).

--Insert Fig. 4 about here—

The major difference between the observed and simulated data sets is in the frequency of switches in decision between adjacent blocks. The general finding is that the learning model underestimates the frequency of switches. The ten proportions of switches for the simulated subjects, which are shown in parentheses in Table 5, are in all but one case smaller that the observed proportions. Table 6, which presents the proportions of switches by block across games (rather than game across blocks), shows that the learning model captures the observed trend of the proportion of switches to decrease across blocks. However, in all nine cases the proportions of the simulated players are smaller than the corresponding proportions of the real subjects.

Table 7 shows that the major reason for this discrepancy in frequency of switches between real and simulated players is that the simulated players converged faster to cutoff decision rules (which allow for switches between markets A and B but not for switches between N and either A or B). Table 7 (bottom row) shows that all the sixty simulated players converged to cutoff rules by block 10 (with no violations) in comparison with only 35 of the real subjects. As the proportion of subjects using cutoff rules consistently gets higher, the frequency of switches between N and either A or B in the 3 by 3 transition matrices in Tables 5 and 6 approaches zero. Consequently, the discrepancy between the proportions of switches of real and simulated subjects is considerably reduced when the effect of the difference in the proportion of real and simulated players using cutoff decision rules consistently is accounted for.

Although the simulated subjects started playing on period 1 with estimated cutoff values (Fig. 3) which are not distributed uniformly and which keep being updated across trials, Table 7 shows that the learning mechanism postulated by the model caused these values to converge to a distribution which is uniform across the values of the combined market capacity. This distribution provides a sufficient condition for equilibrium play. As noted earlier, the same results were obtained for the 35 real subjects who used cutoff rules consistently in block 10 (see Table 7). The null hypothesis that the frequency distribution of the sixty cutoff values of the simulated players in Table 7 is uniform could not be rejected ($\chi^2(9) = 12$, $p > 0.25$).

5. Discussion and Conclusions

Independent sets of twenty stationary agents each repeatedly played a sequence of one-play coordination games with two markets and capacity pairs that were varied randomly from period to period. Our results show that across all games, the symmetric Nash equilibrium hypothesis organizes the aggregate entry decisions remarkably well. The degree of coordination success achieved by our subjects exceeds the one reported by Meyer et al. (1992) and Ochs (1990), who devised different games to study coordination, and is of the same order of magnitude as reported by previous MEG studies with either symmetric or asymmetric players. We have observed a small but consistent deviation from equilibrium frequencies of entry, where low (high) equilibrium frequencies are over-weighted (under-weighted), which we accounted for successfully by a nonlinear transformation of the equilibrium probabilities of the kind advocated by prospect theory (Kahneman & Tversky, 1979). Similar deviations of the same order of magnitude were also found in all previous MEG studies with symmetric players; in another market entry game under uncertainty (Rapoport et al. 1998b); and in a study by Ochs (1990), who reported (Proposition 5), but did not account for, a similar but larger deviation between relative frequencies of entry and the equilibrium probability weights in his high turnover condition. The generality of this finding and its explanation are topics that warrant further investigation.

We have hypothesized that each player perceives the TMEG as consisting of two rather than a single stage, where in the first stage he must decide whether or not to enter the market, and if deciding to enter, in the second stage he must determine which market to enter. Our results support this hypothesis. With experience gained in playing the TMEG, larger proportions of the subjects learn to solve the decision problem on the first stage by using cutoff decision rules in which the cutoff is stated in terms of the combined capacities of the two markets (Table 7). Although the players are symmetric, the individual cutoffs converge through some process of adaptation to different values whose distribution guarantees coordination success. These results are in line with those reported earlier by Rapoport et al. (1997) in their study of the MEG with asymmetric players in which 85 percent (compared to 58.3 percent in the present study) of the subjects used cutoff values consistently in the last block of trials. We attribute this discrepancy to the difference in the complexity of the two games, and conjecture that with more experience with the TMEG the proportion of players using cutoff decision rules will increase. In contrast, the results (Table 6) seem to suggest that the solution of the decision problem on the second stage is made probabilistically as some function of the relative capacities of the two markets.

The adaptive learning model incorporates this two-stage decision process hypothesis and organizes the basic behavioral regularities that we have observed quite well. In particular, it accounts for the coordination success on the aggregate level, the distribution of individual total number of entries, the convergence to cutoff decision rules stated in terms of the combined capacities of the two markets, and the dynamics of play captured by the analysis of switches in decision. Notwithstanding these results, our position is that claims for the simplicity of the model and for its goodness of fit can only be evaluated properly by competitively testing the present model against alternative models. No satisfactory criteria for choosing viable alternative learning models or for comparing competitive learning models with different number of parameters presently exist. The main difference between the behavior of the real and simulated players is due to a faster learning rate postulated by the model. Although we could have reduced this difference by increasing the number of parameters, introducing random error in the individual cutoff decision rules that diminishes with experience, or both, we have opted not to do so. Simplicity of the model's assumptions and its parsimony have been judged at this stage to compensate for the possible gain in goodness of fit.

References

Bush, R. R. and Mosteller, F. (1955). Stochastic Models for Learning. New York: Wiley.

Camerer, C. F., and Ho, T. --H. (1996). Experience-weighted attraction learning in games: A unifying approach. California Institute of Technology, Unpublished manuscript.

Camerer, C. F. and Lovallo, D. (in press). Overconfidence and excess entry: An experimental approach. American Economic Review.

Cooper, D. J. and Feltovitch. N. (1996). Selection of learning rules: Theory and experimental evidence. University of Pittsburgh, Department of Economics, Working paper No. 305.

Crawford, V. (1995). Theory and experiment in the analysis of strategic interaction. University of California, San Diego, Discussion Paper 95-37.

Erev, I. And Rapoport, A. (1998). Coordination, "magic", and reinforcement learning in a market entry game. Games and Economic Behavior, 23, 145 – 176.

Kahneman D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica, 47, 263 - 291.

Karmarkar, U. S. (1978). Subjectively weighted utility: A descriptive extension of the expected utility model. Organizational Behavior and Human Performance, 21, 61 – 72.

Lattimore, P. M., Baker, J. R., and Witte, A. D. (1992). The influence of probability on risky choice. Journal of Economic Behavior and Organization, 17, 377 – 400.

Meyer, D. L., Van Huyck, J. B., Battalio, R. C.. and Saving, T. R. (1992). History's role in coordinating decentralized allocation decisions. Journal of Political Economy, 100, 292 – 316.

Myerson, R. B. (1991). Game Theory: Analysis of Conflict. Cambridge, MA: Harvard University Press.

Ochs, J. (1990). The coordination problem in decentralized markets: An experiment. Quarterly Journal of Economics, 105, 545 – 558.

Ochs, J. (1998). Coordination in market entry games. In D. V. Budescu, I. Erev, and R. Zwick (Eds.), Games and Human Behavior: Essays in Honor of Amnon Rapoport. Hillsdale, NJ: L. Erlbaum.

Osborne, M. J. and Rubinstein, A. (1994). A Course in Game Theory. Cambridge, MA: The MIT Press.

Quiggin, J. (1982). A theory of anticipated utility. Journal of Economic Behavior and organization, 3, 323 - 343.

Rapoport, A. (1995). Individual strategies in a market entry game. Group Decision and Negotiation, 4, 117 133.

Rapoport, A., Seale, D. A., Erev, I., and Sundali, J. A. (1998a). Equilibrium play in large group market entry games. Management Science, 44, 129 – 141.

Rapoport, A., Seale, A., and Ordonez, L. (1998b). Weighted probabilities in large group coordination: Theory and experimental evidence from market entry games. Unpublished manuscript.

Rapoport, A., Seale, D. A., and Winter, E. (1997). Coordination and learning behavior in large groups with asymmetric players. Hong Kong University of Science and Technology, Department of Marketing. Working Paper 97.098.

Roth, A. E. and Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. Games and Economic Behavior, 8, 164 – 212.

Schelling, T. (1960). The Strategy of Conflict. Cambridge, MA: Harvard University press.

Seale, D. A. and Sundali, J. A. (1998). Cheap talk in large group coordination games. In D. V. Budescu, I. Erev, and R. Zwick (Eds.), Games and Human Behavior: Essays in Honor of Amnon Rapoport. Hillsdale, NJ: L. Erlbaum

Sundali, J. A., Rapoport, A, and Seale, D. A. (1995). Coordination in market entry games with symmetric players. Organizational Behavior and Human Decision Processes, 64, 203 - 218.

Tversky, A. and Fox, C. R. (1995). Weighing risk and uncertainty. Psychological Review, 102, 269 – 283.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and Uncertainty, 5, 297 - 323.

Tversky, A. and Wakker, P. (1995). Risk attitudes and decision weights. Econometrica, 63, 1255 – 1280.

Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. American Economic Review, 80, 234 – 249.

Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1991). Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. Quarterly Journal of Economics, 106, 885 – 910.

Yaari, M. (1987). The dual theory of choice under risk. Econometrica, 55, 95 – 115

Acknowledgement

This research has been supported by NSF Grant SBR-9512724 on market entry games awarded to Amnon Rapoport and Kenneth Koput. We wish to thank Lisa Ordonez for her help.

Table 1

Observed and Predicted Means and Standard Deviations of Total Number of Entries by Game

Symmetric Games

Game		(2, 2)		(4, 4)		(6, 6)		(8, 8)		(10, 10)	
	Market	A	B	A	B	A	B	A	B	A	B
	Capacity	2	2	4	4	6	6	8	8	10	10
Group 1	Mean	2.0	1.6	3.5	3.2	6.0	3.5	7.1	6.2	8.3	9.7
Group 2	Mean	2.0	1.5	2.5	4.1	4.6	4.6	6.5	6.2	8.3	8.8
Group 3	Mean	1.4	1.6	2.5	3.2	4.3	5.0	7.1	7.4	8.2	8.7
Total	Mean	1.80	1.57	2.83	3.50	4.97	4.37	6.90	6.60	8.27	9.07
Total	SD	1.22	0.90	1.26	1.20	1.45	1.68	1.52	1.84	1.39	1.41
Predicted	Mean	0.53	0.53	2.63	2.63	4.74	4.74	6.84	6.84	8.95	8.95
Predicted	SD	0.72	0.72	1.51	1.51	1.90	1.90	2.12	2.12	2.22	2.22

Asymmetric Games

Game		(2, 4)		(2, 8)		(6, 8)		(6, 12)		(4, 16)	
	Market	A	B	A	B	A	B	A	B	A	B
	Capacity	2	4	2	8	6	8	6	12	4	16
Group 1	Mean	1.2	4.0	1.7	7.4	5.8	7.1	5.8	10.9	3.1	14.5
Group 2	Mean	2.2	3.2	2.1	6.3	5.6	6.7	4.9	11.5	3.6	14.1
Group 3	Mean	2.1	3.2	1.0	5.6	5.9	6.5	5.8	10.8	2.5	14.7
Total	Mean	1.83	3.47	1.60	6.43	5.77	6.77	5.50	11.07	3.07	14.43
Total	SD	1.32	1.33	1.11	1.53	1.62	1.55	1.72	1.45	1.46	1.54
Predicted	Mean	0.53	2.63	0.53	6.84	4.74	6.84	4.74	11.05	2.63	15.26
Predicted	SD	0.72	1.51	0.72	2.12	1.90	2.12	1.90	2.12	1.51	1.90

Table 2

Observed and Predicted Frequency Distributions of Total Number of Staying out Decisions by Individual Subjects across Blocks for the Five Symmetric Games

	Game (2,2)		Game (4,4)		Game (6,6)		Game (8,8)		Game (10,10)	
Frequency of staying out	**Obs.**	**Pred.**	**Obs.**	**Pred.**	**Obs.**	**Pred.**	**Obs.**	**Pred.**	**Obs.**	**Pred.**
10	32	34.94	22	4.06	10	0.10	3	0	2	0
9	9	19.41	7	12.54	3	0.88	1	0.01	0	0
8	3	4.85	4	17.45	8	3.57	4	0.18	0	0
7	4	0.72	4	14.39	6	8.56	4	0.72	2	0
6	3	0.07	3	7.78	3	13.48	2	2.74	1	0.01
5	3	0	3	2.89	5	14.56	6	7.12	1	0.11
4	1	0	3	0.74	6	10.92	4	12.86	0	0.79
3	2	0	4	0.13	3	5.62	6	15.92	6	3.86
2	1	0	4	0	2	1.90	5	12.93	4	12.29
1	1	0	2	0	7	0.38	10	6.23	9	23.21
0	1	0	4	0	7	0.03	15	1.35	35	19.73
Mean	1.68		3.17		4.67		6.75		8.67	
Median	0		2		4.5		7.5		10	

Table 3

Observed Frequency Distribution of Total Number of Entries by Group

Class	Group 1	Group 2	Group 3
0- 9	1	0	1
10-19	0	1	0
20-29	1	0	2
30-39	2	6	2
40-49	3	1	4
50-59	5	4	3
60-69	3	2	2
70-79	2	3	3
80-89	2	2	1
90-100	1	1	2
Mean	56.30	54.65	53.75
SD	22.96	22.12	24.13
Predicted (Pure strategy)	54	54	54

Table 4

Observed Number of Entries by Block and Game

Symmetric Games: Market A

Game	Block 1	2	3	4	5	6	7	8	9	10	Mean	SD	Pr.
2, 2	10	4	7	5	5	8	3	3	5	4	5.4	2.27	1.58
4,4	10	11	12	10	8	8	5	10	5	6	8.5	2.51	7.89
6,6	15	13	12	10	14	23	15	15	12	20	14.9	3.90	14.21
8,8	22	23	24	16	24	17	19	22	19	21	20.7	2.83	20.52
10,10	23	23	25	23	29	27	28	26	25	19	24.8	2.94	26.84
RMSE	4.30	2.75	2.69	3.72	2.37	5.12	1.71	1.40	2.48	4.57	1.98		

Symmetric Games: Market B

Game	Block 1	2	3	4	5	6	7	8	9	10	Mean	SD	Pr.
2,2	3	5	6	3	3	5	5	7	4	6	4.7	1.42	1.58
4,4	6	11	10	10	16	14	9	11	9	9	10.5	2.80	7.89
6,6	17	14	9	22	13	5	15	11	15	10	13.1	4.70	14.21
8,8	16	18	19	20	17	24	18	26	18	22	19.8	3.22	20.52
10,10	32	23	30	29	23	26	26	27	27	29	27.2	2.90	26.84
RMSE	3.55	2.71	3.56	3.80	4.39	5.41	2.03	3.97	1.68	2.95	1.92		

Asymmetric Games: Market A

Game	Block 1	2	3	4	5	6	7	8	9	10	Mean	SD	Pr.
2,4	10	6	9	4	8	3	4	0	9	2	5.5	3.41	1.58
2,8	8	9	4	4	5	2	5	5	0	6	4.8	2.62	1.58
6,8	24	16	14	20	19	13	15	20	15	17	17.3	3.40	14.21
6,12	29	18	19	18	18	10	14	14	14	11	16.5	5.38	14.21
4,16	9	13	6	9	8	10	11	8	11	7	9.2	2.10	7.89
RMSE	9.08	4.18	3.40	4.25	2.27	2.36	3.09	3.68	2.78	2.91			

Asymmetric Games: Market B

Game	Block 1	2	3	4	5	6	7	8	9	10	Mean	SD	Pr.
2,4	6	13	8	13	10	8	10	12	7	17	10.4	3.20	7.89
2,8	19	15	20	21	27	15	19	17	19	21	19.3	3.29	20.52
6,8	17	26	20	20	18	26	20	19	19	18	20.3	3.00	20.52
6,12	26	33	35	30	35	33	37	33	35	35	33.2	2.99	33.16
4,16	48	39	50	39	45	42	41	44	42	43	43.3	3.41	45.79
RMSE	3.50	4.71	2.08	4.07	3.35	3.87	2.99	2.64	2.15	4.49	1.68		

Table 5

Observed (and Simulated) Number of Switches across Blocks by Game

Symmetric Games

	(2, 2)			(4, 4)			(6, 6)			(8, 8)			(10, 10)		
	N	A	B	N	A	B	N	A	B	N	A	B	N	A	B
N	**417**	18	14	**313**	26	26	**226**	27	37	**106**	35	37	**40**	14	14
A	21	**17**	12	16	**38**	25	27	**75**	27	27	**100**	59	16	**142**	71
B	14	9	**18**	37	11	**48**	39	32	**50**	40	50	**86**	19	69	**155**
		.163			.261			.350			.459			.376	
		(.060)			(.152)			(.237)			(.285)			(.395)	

Asymmetric Games

	(2, 4)			(2, 8)			(6, 8)			(6, 12)			(4, 16)		
	N	A	B	N	A	B	N	A	B	N	A	B	N	A	B
N	**343**	23	34	**252**	17	57	**138**	26	35	**52**	13	24	**41**	4	20
A	29	**11**	13	16	**10**	16	34	**71**	52	19	**72**	63	7	**36**	42
B	25	11	**51**	58	13	**101**	33	52	**99**	27	51	**219**	24	43	**323**
		.250			.328			.430			.365			.259	
		(.055)			(.098)			(.148)			(.193)			(.187)	

Table 6

Observed (and Simulated) Number of Switches across Games by Block

Block

	1			2			3			4			5		
	N	A	B	N	A	B	N	A	B	N	A	B	N	A	B
N	**183**	26	41	**206**	32	29	**204**	13	44	**209**	28	37	**207**	25	23
A	43	**63**	54	16	**69**	51	28	**63**	42	14	**64**	41	29	**66**	43
B	41	47	**102**	39	31	**127**	42	43	**121**	32	46	**129**	45	30	**132**
	.420			.330			.353			.330			.325		
	(.262)			(.223)			(.218)			(.195)			(.198)		

Block

	6			7			8			9		
	N	A	B	N	A	B	N	A	B	N	A	B
N	**222**	22	37	**226**	24	31	**232**	12	26	**239**	21	30
A	24	**62**	35	17	**61**	41	29	**62**	32	12	**62**	41
B	35	35	**128**	27	38	**135**	29	41	**137**	26	30	**139**
	.313			.297			.282			.267		
	(.163)			(.182)			(.168)			(.180)		

Table 7

Observed and Simulated Frequency Distribution of Estimated Cutoff Values c*

in Block 10

Estimated c*	4	6	8	10	12	14	16	18	20	21	**Sum**
Observed	7	3	1	5	3	4	3	4	2	3	35
Simulated	8	3	7	4	4	11	5	2	8	8	60

Table 8

Best Fitting Parameter Values of the Learning Model

Parameter	Group 1	Group 2	Group 3
w^+	0.85	0.19	0.45
w^-	0.88	0.18	0.18
z	3.30	1.98	3.80
d	0.05	0.04	0.04
Q	1210	1271	1028

Table 9

Simulated Means and Standard Deviations of Total Number of Entries

Symmetric Games

		(2, 2)		(4, 4)		(6, 6)		(8, 8)		(10, 10)	
	Market	A	B	A	B	A	B	A	B	A	B
	Capacity	2	2	4	4	6	6	8	8	10	10
Group 1	Mean	1.9	1.4	3.0	3.0	4.7	3.8	5.8	6.8	8.1	8.8
Group 2	Mean	0.9	1.1	3.2	2.1	4.5	4.3	8.3	8.0	9.2	9.8
Group 3	Mean	1.0	1.0	3.3	2.1	4.5	5.1	7.5	7.0	8.0	8.4
Total	Mean	1.27	1.17	3.17	2.40	4.57	4.40	7.20	7.27	8.43	9.00
Total	SD	0.94	0.79	1.42	1.38	1.25	1.40	2.16	2.23	1.83	1.98

Asymmetric Games

		(2, 4)		(2, 8)		(6, 8)		(6, 12)		(4, 16)	
	Market	A	B	A	B	A	B	A	B	A	B
	Capacity	2	4	2	8	6	8	6	12	4	16
Group 1	Mean	1.6	3.9	2.0	5.0	2.9	8.2	4.5	9.8	4.3	12.7
Group 2	Mean	0.8	1.2	1.3	5.7	4.1	9.0	5.8	11.7	4.8	14.2
Group 3	Mean	0.9	1.8	1.5	6.0	3.3	10.0	5.1	10.0	3.7	12.5
Total	Mean	1.10	2.30	1.60	5.57	3.43	9.07	5.13	10.50	4.27	13.03
Total	SD	0.80	1.47	1.28	1.41	1.65	1.57	2.86	2.45	2.21	2.26

Table 10

Simulated Number of Entries by Block and game

Symmetric Games: Market A

	Block											
Game	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**Mean**	**SD**
2,2	3	3	3	5	5	4	7	4	2	2	3.8	1.55
4,4	4	8	13	6	7	12	14	8	11	12	9.5	3.34
6,6	14	16	14	16	11	12	19	9	14	12	13.7	2.87
8,8	23	16	17	15	26	24	32	23	23	17	21.6	5.34
10,10	18	32	24	21	29	33	26	28	21	21	25.3	5.12

Symmetric Games: Market B

	Block											
Game	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**Mean**	**SD**
2,2	2	2	4	3	3	4	1	4	6	6	3.5	1.65
4,4	7	7	3	11	11	6	4	10	7	6	7.2	2.74
6,6	16	12	13	11	16	14	7	17	12	14	13.2	2.94
8,8	29	29	26	27	16	18	10	19	19	25	21.8	6.34
10,10	37	20	28	31	23	19	26	24	31	31	27.0	5.66

Asymmetric Games: Market A

	Block											
Game	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**Mean**	**SD**
2,2	3	5	4	3	2	4	3	1	4	4	3.3	1.16
2,8	3	12	6	3	6	2	2	4	4	6	4.6	2.97
6,8	16	13	8	13	8	3	9	7	12	14	10.3	3.95
6,12	35	20	22	16	10	13	9	10	9	10	15.4	8.33
4,16	19	4	13	16	21	16	14	5	12	8	12.8	5.67

Asymmetric Games: Market B

	Block											
Game	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**Mean**	**SD**
2,4	2	5	7	7	9	7	8	10	7	7	6.9	2.18
2,8	17	8	15	19	16	20	20	18	18	16	16.7	3.50
6,8	26	25	30	24	27	34	28	30	25	23	27.2	3.36
6,12	18	30	26	31	39	32	36	34	35	34	31.5	5.93
4,16	34	49	39	36	31	36	38	47	40	44	39.4	5.74

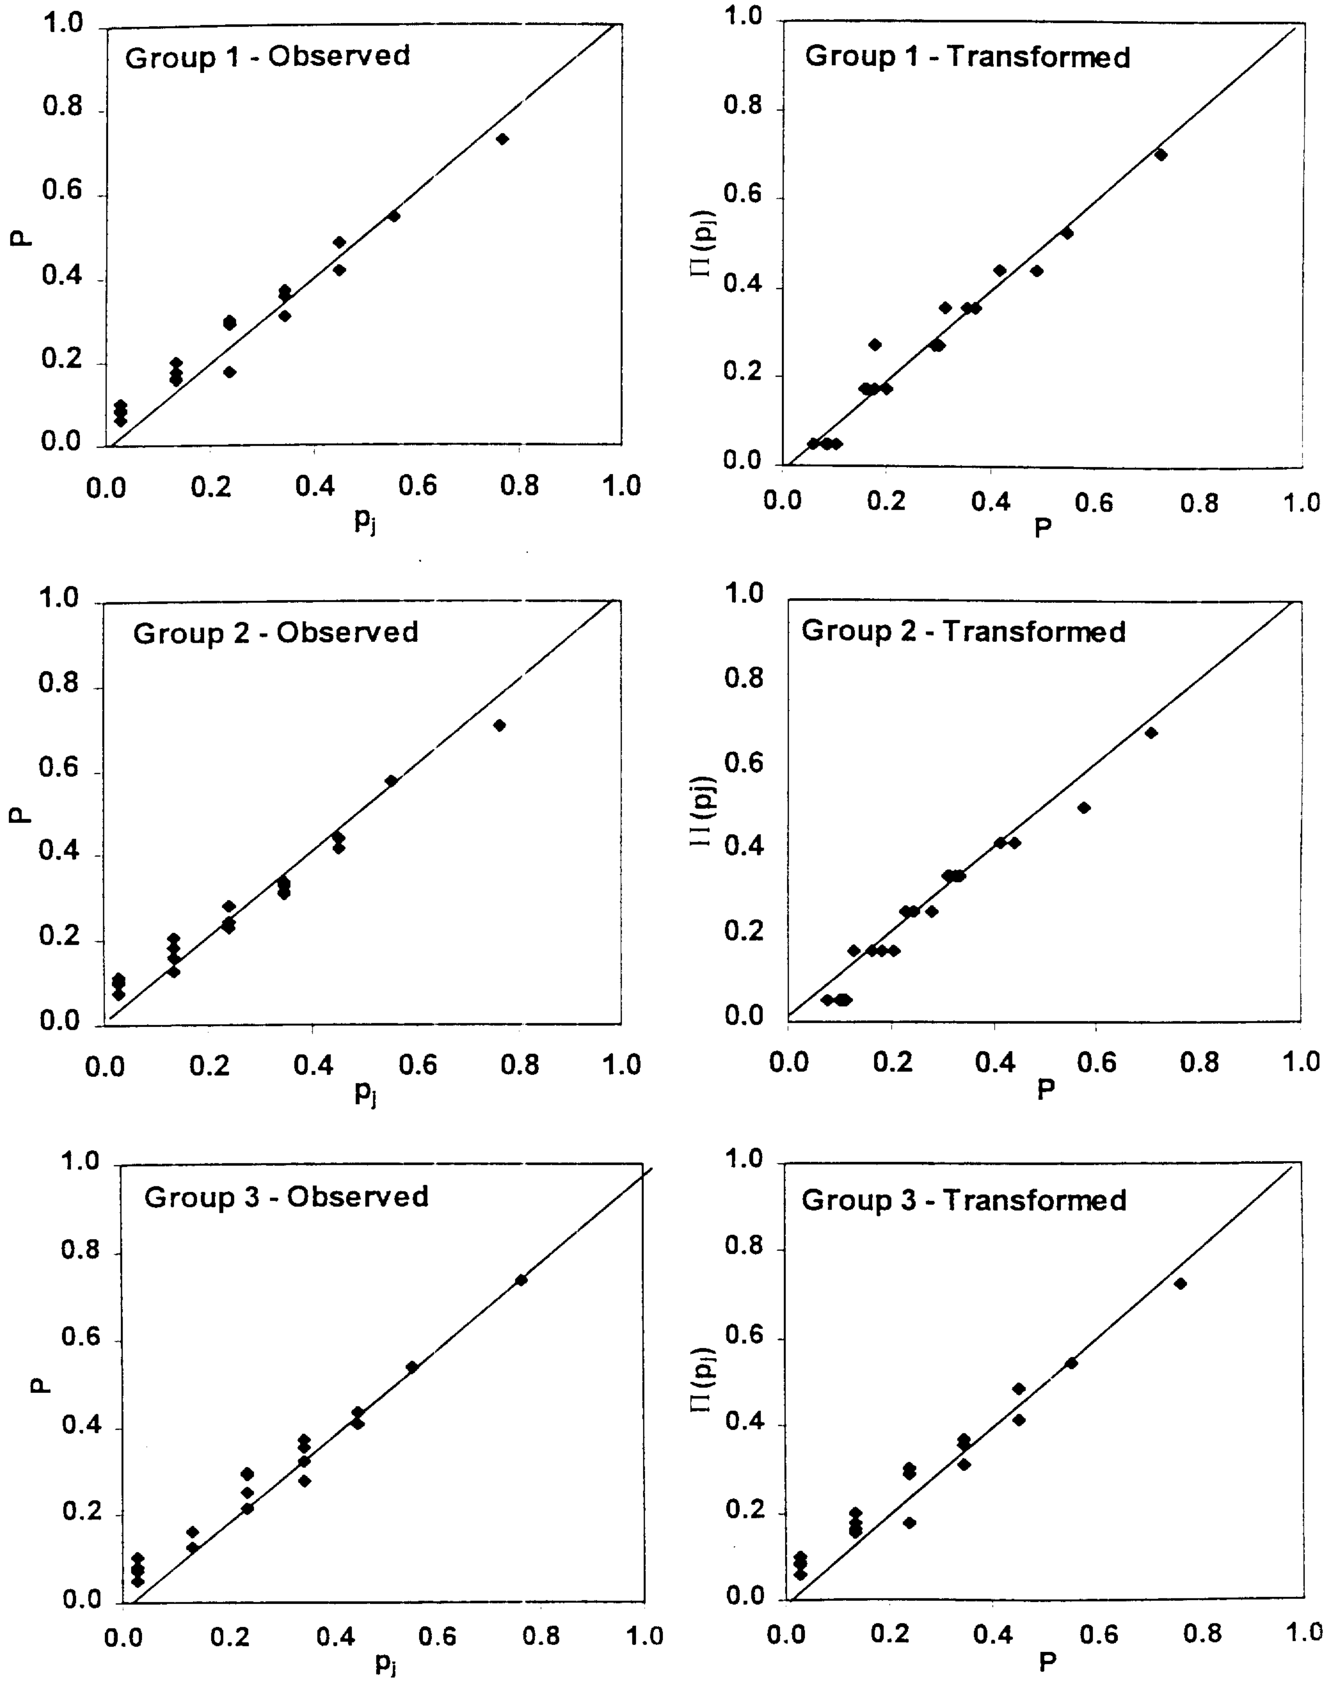

Fig. 1. Observed proportions of entry plotted against equilibrium probabilities of entry (left panel) and transformed probabilities of entry plotted against the observed proportions (right panel) across symmetric and asymmetric games.

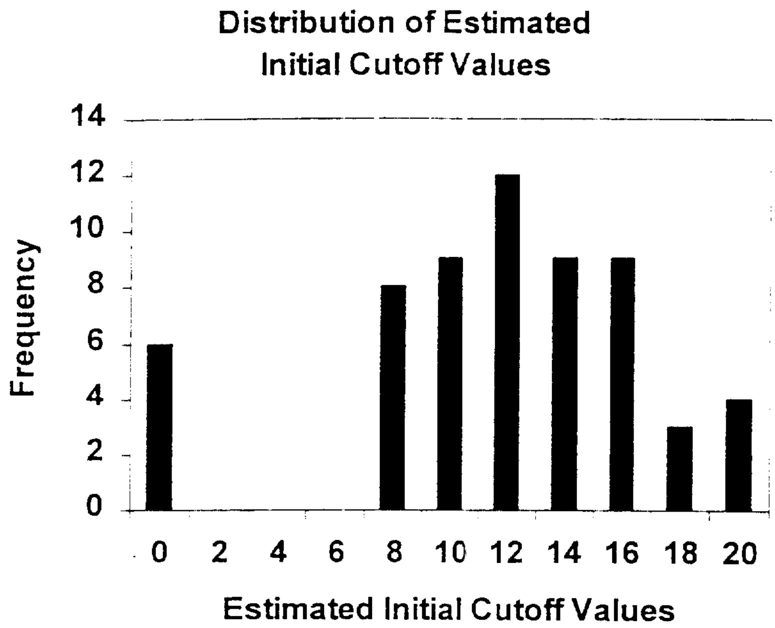

Fig. 2. Frequency distribution of estimated initial cutoff values for the simulated players.

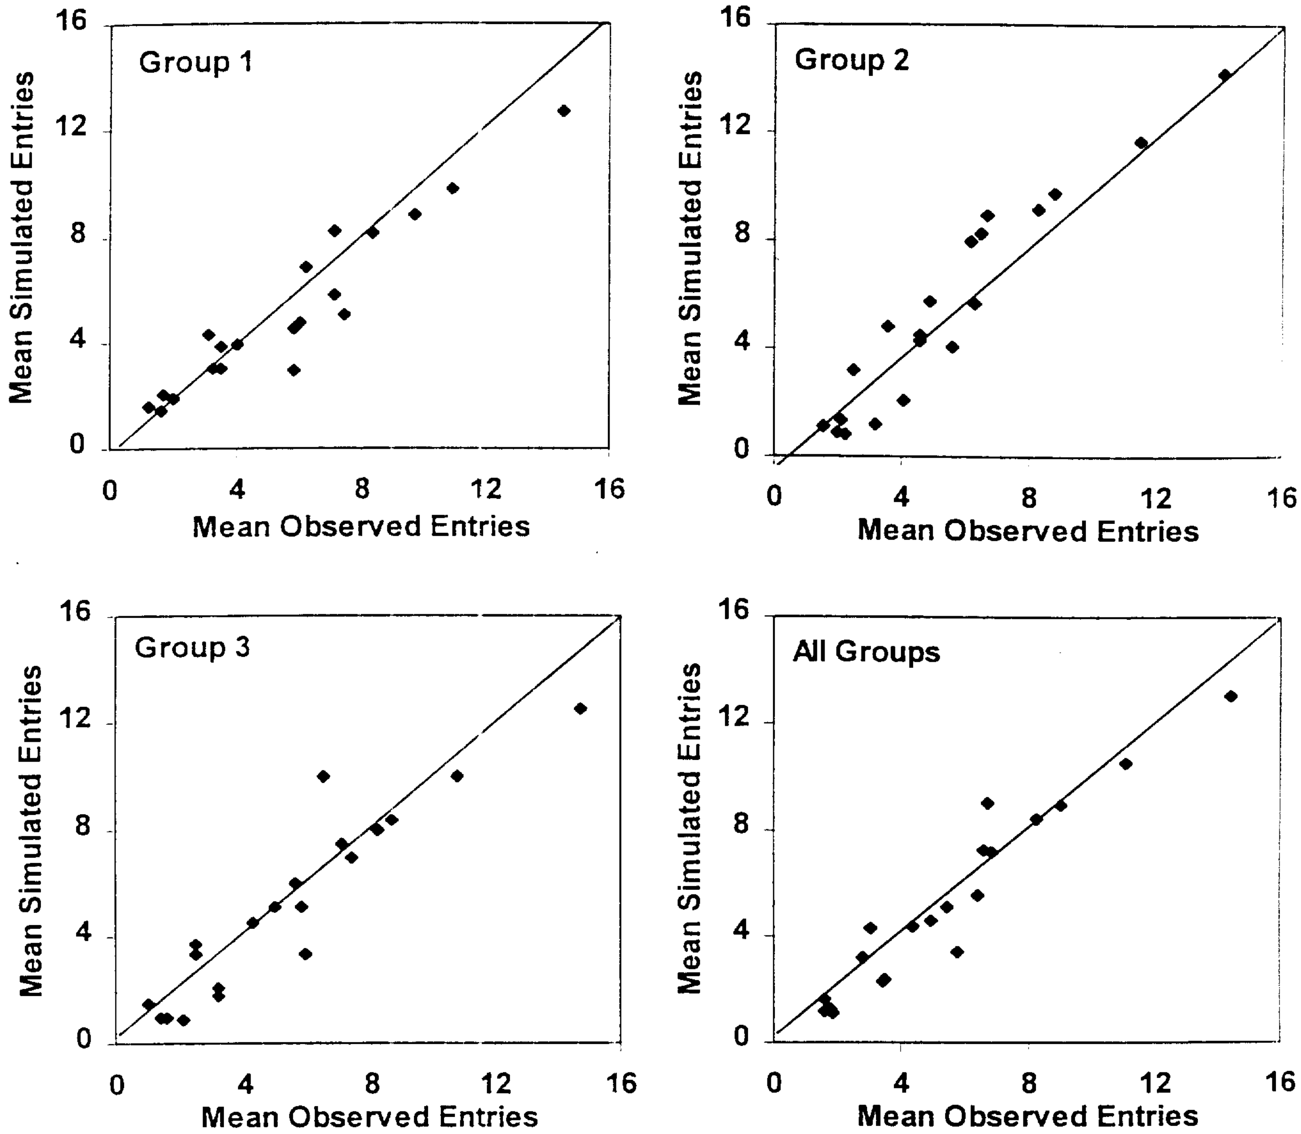

Fig. 3. Scatter plots of simulated against observed mean number of entries across symmetric and asymmetric games.

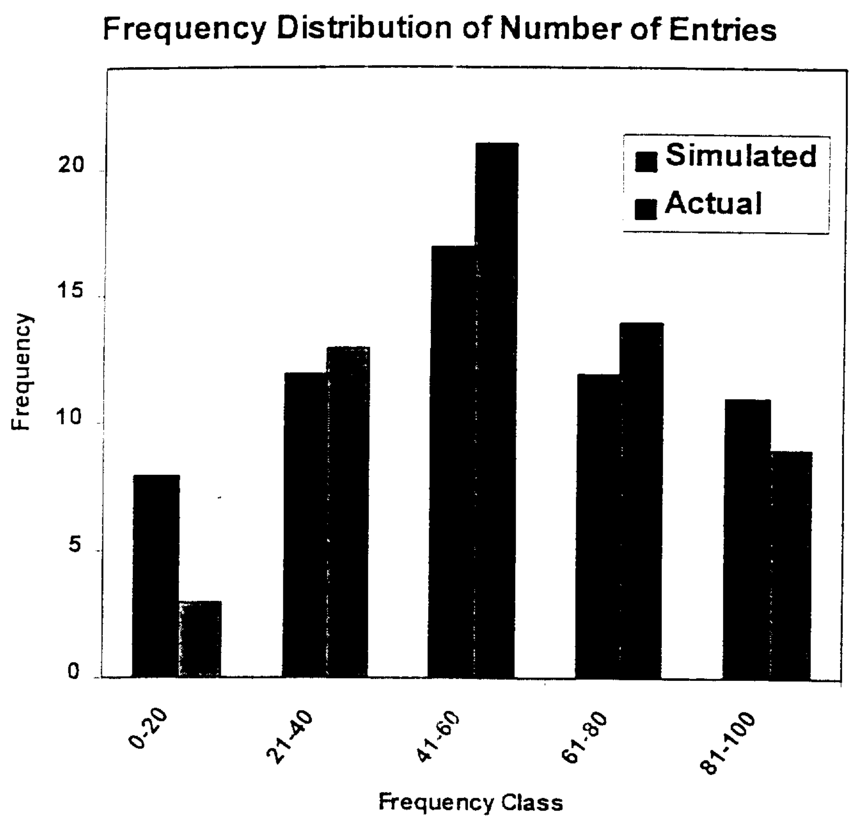

Fig. 4. Frequency distribution of observed and simulated individual total number of entry decisions.

Please forward your requests for working papers to the following address:

Executive Officer
Department of Marketing
School of Business and Management
Hong Kong University of Science & Technology
Clear Water Bay
Kowloon, Hong Kong

www: http://www.bm.ust.hk/~mark/
Enquiries: Tel (852) 2358 7700 Fax (852) 2358 2429