David vs. Goliath: An Analysis of Asymmetric Mixed-Strategy Games and Experimental Evidence

Wilfred Amaldoss • Sanjay Jain
Fuqua Business School, Duke University, Durham, North Carolina 27708
R. H. Smith School of Business, University of Maryland, College Park, Maryland 20742
wilfred.amaldoss@duke.edu • sjain@rhsmith.umd.edu

Mixed strategies are widely used to model strategic situations in diverse fields such as economics, marketing, political science, and biology. However, some of the implications of asymmetric mixed-strategy solutions are counterintuitive. We develop a stylized model of patent race to examine some of these implications. In our model two firms compete to develop a product and obtain a patent. However, one firm values the patent more because of its market advantages, such as brand reputation and distribution network. Contrary to some intuition, we find that the firm that values the patent less is likely to invest more aggressively in developing the product and will also win the patent more often. We argue that the reason for these counterintuitive results is inherent in the very concept of mixed strategy solution. In a laboratory test, we examine whether subjects' behavior conforms to the equilibrium predictions. We find that the aggregate behavior of our subjects is consistent with the game-theoretic predictions. With the help of the experience-weighted attraction (EWA) learning model proposed by Camerer and Ho (1999), we show that adaptive learning can account for the investment behavior of our subjects. We find that the EWA learning model tracks the investment decisions of our subjects well, whether we hold out trials or an entire group of subjects.
(*Competitive Strategy; R&D; Marketing; Decision Analysis; Learning; Experimental Economics*)

1. Introduction

Noncooperative game-theoretic models are widely applied in various fields, including economics, marketing, political science, and biology. The most fundamental concept of noncooperative game theory is the Nash equilibrium, which allows players to use pure strategies and/or mixed strategies. Mixed strategies have been applied in various contexts such as promotion, capacity choice, pricing, bargaining, R&D, and product standardization.[1] This popularity of mixed-strategy solutions has an important basis: Mixed-strategy equilibrium often exists even when no equilibrium in pure strategy exists.

There is also some empirical support to justify the use of symmetric mixed strategies. In laboratory experiments, researchers have found that at the aggregate level mixed strategies can adequately account for the actual behavior of subjects when players are symmetric (e.g., O'Neill 1987, Rapoport and Boebel 1992, Rapoport and Budescu 1992, Rapoport and Amaldoss 2000; see also Chapter 2 in Camerer 2001 for a survey of experimental research on mixed-strategy games).[2]

[1] For instance, we see the use of a mixed-strategy equilibrium in the price promotion models of Varian (1980), in the capacity choice model of Deneckere and Peck (1995), in the pricing model of Lan and Kanafani (1993), in the R&D model of Fudenberg and Tirole (1985), and in the standard setting model of Farrell and Saloner (1988).

[2] There is also some evidence which suggests that mixed strategies do not survive a laboratory test (see, for example, Brown and Rosenthal 1990, Ochs 1995).

Management Science © 2002 INFORMS
Vol. 48, No. 8, August 2002 pp. 972–991

0025-1909/02/4808/0972$5.00
1526-5501 electronic ISSN

However, the use of mixed strategies when players are asymmetric can lead to theoretical results that are not consistent with intuition. For example, Farrell and Saloner (1988, p. 250) consciously avoid using asymmetric mixed strategies in the context of a model of standards, given the counterintuitive nature of its implications. Cheng and Zhu (1995) recommend the use of quadratic utility rather than expected utility to circumvent some of these issues. Furthermore, there is little empirical evidence to suggest that asymmetric mixed strategies can account for the actual behavior of economic agents. In this paper, we illustrate some of the counterintuitive results of asymmetric mixed-strategy games using a model of patent races and then examine whether subjects' behavior conforms to the game-theoretic predictions. We also study whether adaptive mechanisms can track how subjects learn to play the game.

We model a patent race as an asymmetric mixed-strategy game. In our game, two firms compete to develop a product and obtain a patent. However, one firm values the patent more because of its market advantages, such as brand reputation and distribution network. A firm, for example, can extend the protective umbrella of its preexisting brand-name reputation to the innovation and exploit consumers' willingness to pay a brand premium (Bresnahan et al. 1997). Alternatively, a firm can potentially leverage its preexisting distribution network to distribute the innovation at a low incremental cost. Intuition would suggest that the firm that values the patent more will invest more and, consequently, win the patent more often. However, our theoretical results show that, in fact, in equilibrium the opposite is true: The player who has less to benefit from winning the patent will invest more aggressively and win the patent more often.

We later argue that the reason for this result is intimately connected to the very concept of mixed strategies and is therefore likely to appear in a variety of models. Consequently, it is important to assess the descriptive validity of such a theoretical result. If, as intuition would suggest, asymmetric mixed-strategy solution fails to account for actual behavior, then we should attempt to identify the cause of its failure, and also be cautious in using the concept to build models. On the other hand, if asymmetric mixed strategy solution can describe actual behavior, then we can more confidently apply the solution concept in a variety of situations where players are asymmetric.

We submit our model to a controlled laboratory test that permits us to examine to what extent the actual behavior of economic agents conforms to the game-theoretic predictions. In general, the experimental results suggest that equilibrium predictions closely track the aggregate behavior of financially motivated subjects. The low-reward players invested more than the high-reward players. They also won the competition more often. Furthermore, the aggregate distributions of investments and payoffs follow the model predictions. However, we notice some significant departures from the theoretical norm in one of the replications of the experiment, namely Group 1. While aggregate behavior conforms to theory, the behavior of individual subjects varied widely. Hence, the model survives at the aggregate rather than the individual level.

We then examine whether adaptive learning mechanisms can account for the investment behavior of our subjects. For this investigation we use the EWA model proposed by Camerer and Ho (1999), which has been shown to have good predictive accuracy in a variety of contexts (see Camerer et al., forthcoming). The EWA model tracks the investment decisions of our subjects well, both in the calibration and validation samples.

Our research adds to the growing body of literature on strategic thinking in games (e.g., Costa-Gomes et al., forthcoming, Johnson et al., forthcoming; see also Camerer 2001 for a recent review). We show that the aggregate behavior of our subjects is consistent with the asymmetric mixed-strategy solution. Although our primary focus is on assessing the descriptive validity of the mixed-strategy solution concept, our results provide yet another plausible explanation for why a firm with market advantage might fail to innovate: Market advantage can provoke aggressive investment by competitors and potentially impede innovation by the firm that enjoys market advantage. We show that such a failure to innovate may be a consequence of strategic behavior rather than any inefficiency in management. The implications of our theoretical model and experimental evidence can potentially be useful in understanding a broader set of phenomena that involves asymmetric mixed-strategy games.

The rest of the paper is organized as follows. In §2, we develop a parsimonious model of a patent race where the players are asymmetric. In §3, we discuss a laboratory test of the model. Later, in §4 we discuss the results of the EWA learning model. Finally, we conclude the paper in §5.

2. An Illustrative Model of a Patent Race Between Asymmetric Firms

In this section we develop a parsimonious model of a patent race to illustrate the behavioral implications of asymmetric mixed strategies.[3] Consider a situation in which two firms in the market are competitively investing in R&D to win a patent. We assume that both firms can invest an amount $x \in [0, c]$, where c represents the financial constraints imposed on the firm by the capital market. These financial constraints may be a consequence of information asymmetries that exist between the innovating firm and the capital market (Lerner and Merges 1998). We assume that firms invest in discrete units of size c/k. In other words, a firm can invest $(0, c/k, 2c/k, 3c/k, \ldots, c)$ units. Thus, there are $(k+1)$ levels of feasible investment. We also examine the case in which the strategy space is continuous by taking the limit as $k \to \infty$.[4]

[3] Although our model is developed to highlight some theoretical implications of asymmetric mixed strategies, it addresses an important issue in the context of patent races. Prior theoretical research in the R&D literature has primarily concentrated on the issue of timing of entry (e.g., Gilbert and Newberry 1982, Reinganum 1983). This research has examined whether, and if so when, an incumbent firm would continue to dominate the market. There is also a large body of empirical work that has examined whether market pioneers enjoy enduring market leadership (e.g., Robinson and Fornell 1985, Golder and Tellis 1993). In contrast to prior theoretical research, we model a situation in which both firms are either incumbents or entrants. Our model allows us to study how firm-level asymmetry in market advantages can impact their incentives to innovate. For a review of the literature on R&D models see, for example, Reinganum (1989).

[4] In Technical Appendix B we present the analysis for the continuous case. The technical appendices are available from the authors upon request and also on the INFORMS website at http://mansci.pubs.informs.org/ecompanion.html.

Following Gilbert and Newberry (1982), we assume that the firm that invests more wins the patent (see also Fudenberg et al. 1983). If firm i wins the patent, it makes profit r_i, $i = \{L, H\}$ for the patent. We assume that the firms are asymmetric in terms of the profits they make if they win the patent (i.e., $r_L \neq r_H$). This may be due to the differences in the ability of these firms to exploit the new technology. For example, a firm can extend the protective umbrella of its preexisting brand-name reputation to its new innovation and exploit consumers' willingness to pay a brand premium (Bresnahan et al. 1997). Furthermore, a firm can leverage its preexisting distribution network to efficiently reach diverse market segments at a low incremental cost. Thus, factors such as marketing ability and synergy with current product lines can influence a firm's ability to profit from an innovation. Without any loss in generality we assume $r_H > r_L$. If both firms invest the same amount in R&D, then we assume that firms engage in Bertrand competition and both firms make zero profits from the patent. Thus, firm i's profits if it invests x_i are:

$$\Pi_i(x_i) = \begin{cases} r_i - x_i & \text{if } x_i > x_j;\ i \neq j, \\ -x_i & \text{otherwise.} \end{cases} \tag{1}$$

The payoffs associated with various levels of investment are presented in Table 1. We assume that $c < r_L$. This ensures that the capital constraint is binding. The assumption can be relaxed easily. Finally, in our model we assume that r_i and c are common knowledge.

Now we proceed to characterize the equilibrium solution. We present the formal analysis in Technical Appendix A. Note that there can be no pure strategy equilibrium in the game when we have at least three levels (i.e., $k > 1$). To see this, first consider the case in which both firms invest zero. In this case one firm can do better by investing c/k and winning the reward. Hence, investing zero is not a pure strategy equilibrium. Next, consider the case in which both players are investing positive amounts. In this case the losing firm can do better by deviating to zero investment. So, both firms investing positive amounts is not a pure strategy equilibrium. Finally, consider the case in which one firm is investing zero. In this case

Table 1 Payoff Matrix for the Asymmetric Mixed-Strategy Game

			x_L			
x_H	0	c/k	$2c/k$	$3c/k$	$\dots$	c
0	$0, 0$	$0, r_L - c/k$	$0, r_L - 2c/k$	$0, r_L - 3c/k$	$\dots$	$0, r_L - c$
c/k	$r_H - c/k, 0$	$-c/k, -c/k$	$-c/k, r_L - 2c/k$	$-c/k, r_L - 3c/k$	$\dots$	$-c/k, r_L - c$
$2c/k$	$r_H - 2c/k, 0$	$r_H - 2c/k, -c/k$	$-2c/k, -2c/k$	$-2c/k, r_L - 3c/k$	$\dots$	$-2c/k, -c$
$3c/k$	$r_H - 3c/k, 0$	$r_H - 3c/k, -c/k$	$r_H - 3c/k, -2c/k$	$-3c/k, -3c/k$	$\dots$	$-3c/k, r_L - c$
.	.	.	.	.	$\dots$	.
.	.	.	.	.	$\dots$	.
.	.	.	.	.	$\dots$	.
c	$r_H - c, 0$	$r_H - c, -c/k$	$r_H - c, -2c/k$	$r_H - c, -3c/k$	$\dots$	$-c, -c$

the pure strategy equilibrium must involve the other firm investing c/k. However, this cannot be an equilibrium because the losing firm can then win the game by investing $2c/k$.[5] Thus, there is no pure-strategy equilibrium. Hence, the equilibrium must necessarily involve mixed strategies. The equilibrium solution is specified in the following proposition.

Proposition 1. *If $k > 1$, the unique equilibrium of this discrete game is for Firm H to invest ic/k discrete units of capital (where $i = 0, 1, \dots, k$) with probability*:

$$p_H\left(i\frac{c}{k}\right) = \begin{cases} \frac{c}{k}\frac{1}{r_L} & \text{if } i = 0, 1, \dots, k-1, \\ 1 - c/r_L & \text{if } i = k. \end{cases} \tag{2}$$

Similarly, for Firm L we have:

$$p_L\left(i\frac{c}{k}\right) = \begin{cases} \frac{c}{k}\frac{1}{r_H} & \text{if } i = 0, 1, \dots, k-1, \\ 1 - c/r_H & \text{if } i = k. \end{cases} \tag{3}$$

The corresponding c.d.f. for Firms H and L converges to (4) *and* (5), *respectively, as $k \to \infty$ where*:

$$F_H(x) = \begin{cases} \frac{x}{r_L} & 0 \le x < c, \\ 1 & x \ge c, \end{cases} \tag{4}$$

$$F_L(x) = \begin{cases} \frac{x}{r_H} & 0 \le x < c, \\ 1 & x \ge c. \end{cases} \tag{5}$$

[5] If $k = 1$, then the strategy space reduces to two levels, namely $x \in \{0, c\}$. Thus, if one firm invests c, then the other firm cannot defect and invest $2c$. This game has no symmetric pure-strategy equilibrium, although it has a symmetric mixed-strategy equilibrium. In addition, there are also two asymmetric pure strategy solutions, in which one firm invests nothing and the other firm invests c. However, if $k > 1$, then these asymmetric pure-strategy solutions do not exist.

Proposition 1 leads to the following results:

Result 1. *Firm L's equilibrium strategy does not depend on r_L, but rather depends on r_H.*

To see this, note that the cumulative distribution function F_L and F_H in the discrete case can be given by:

$$F_L\left(i\frac{c}{k}\right) = \begin{cases} \frac{ic}{kr_H} & i = 0, 1, \dots, k-1, \\ 1 & i = k, \end{cases} \tag{6}$$

$$F_H\left(i\frac{c}{k}\right) = \begin{cases} \frac{ic}{kr_L} & i = 0, 1, \dots, k-1, \\ 1 & i = k. \end{cases} \tag{7}$$

Result 2. *On average, Firm L invests more than Firm H.*

To see this, note that from (6) and (7) it follows that F_L dominates F_H in the first-order stochastic sense since $r_H > r_L$. This implies that the mean investment of Firm L must be higher than that of Firm H. Therefore, the firm which values the patent less will invest more aggressively.

Result 3. *Firm L is more likely to win the patent.*

To see this, note that the probability that Firm L wins is given by:

$$\text{Prob}(L \text{ wins}) = \left(\frac{r_H - c}{r_H}\right)\left(\frac{c}{r_L}\right) + \sum_{i=1}^{k-1}\left(\frac{c}{kr_H}\right)\left(\frac{ic}{kr_L}\right). \tag{8}$$

The first term is the probability that Firm L will invest c and Firm H will invest less than c, and the second term is the probability that Firm L invests an amount

ic/k and Firm H will invest less than ic/k. This simplifies to:

$$\text{Prob}(L \text{ wins}) = \frac{c\big((r_H - c/2)k - c/2\big)}{k r_L r_H}. \tag{9}$$

Similarly,

$$\text{Prob}(H \text{ wins}) = \frac{c\big((r_L - c/2)k - c/2\big)}{k r_L r_H}. \tag{10}$$

As $r_H > r_L$, it follows from (9) and (10) that Firm L is more likely to win the patent.

Results 1–3 are not consistent with intuition, but are the only results consistent with strategic thinking about the game. To better appreciate the sharp contrast between strategic and nonstrategic thinking in the context of this game, let us now examine the game from the perspective of naïve, nonstrategic players. Assume that Firm L believes that Firm H will invest x with probability $f_H(x)$, and let $F_H(\cdot)$ be the associated cumulative distribution function. Then Firm L will choose x_L^* such that:

$$x_L^* \in \arg\max_{x\in[0,c]} \mathbb{E}\Pi_L(x_L) = \arg\max_{x\in[0,c]} (r_L - x_L)F_H(x_L) + (1 - F_H(x_L))(-x_L). \tag{11}$$

For simplicity, assume that there is an interior maximum. Clearly, x_L^* will be independent of r_H. Furthermore, using the implicit function theorem, we have:

$$\text{sign}\left(\frac{dx_L^*}{dr_L}\right) = \text{sign}\left(\frac{\partial^2 \mathbb{E}\Pi_L(x_L)}{\partial x_L \partial r_L}\right) = \text{sign}(f_H(x_L^*)) > 0. \tag{12}$$

This implies that a firm will invest more as its reward for winning the game increases. Thus, a nonstrategic analysis of this game provides results that might be consistent with some of our intuition. The equilibrium solution is exactly the opposite, and strategic behavior on the part of firms is essential for the game-theoretic results. As we have seen earlier, it cannot be an equilibrium for both firms to play pure strategies. Consequently, the equilibrium has to be such that both firms randomize among some of their pure strategies. However, for a firm to randomly choose between any two pure strategies, the payoffs associated with these two strategies must be the same, thereby making the firm indifferent between these strategies. Otherwise, the firm would choose the action that gives it the higher payoff, and would not randomize. Thus, if a firm gets zero from not investing, then its expected payoffs from investing a positive amount must also be zero. If firm i invests a positive amount x, then its expected payoff is

$$\mathbb{E}(\text{Payoff of firm } i) = r_i \text{Prob}(i \text{ wins}) - x. \tag{13}$$

However, the expected payoff of investing x has to be zero, as investing nothing is part of the mixed-strategy equilibrium. This implies that Firm L, which gains less from winning the patent, should win the game more often when it invests a positive amount. Such a high winning probability ensures that Firm L remains indifferent between not investing and investing a positive amount. Thus, we see that the mixed strategy equilibrium has to be such that the firm with less to gain from winning the patent plays the game more aggressively and wins the patent more often. Finally, note that the probability that a firm wins the patent, if it invests a positive amount x, depends only on the mixed-strategy distribution of the *other* firm. Thus, Firm L's equilibrium mixed strategy must be such that it makes Firm H indifferent between its various pure strategies. Consequently, Firm L's mixed strategy has to depend on the payoffs that Firm H receives by either investing or not investing. The maximum payoff that Firm H receives by investing is r_H and is independent of r_L. Therefore, Firm L's strategy will only depend on r_H and not r_L.

Discussion

In sum, the mixed-strategy solution in an asymmetric setting requires players to focus on their competitor's reward instead of their rewards. Although such a behavior is counterintuitive, it is the only behavior that is consistent with strategic thinking about this game. It is useful to note that, although such results are only observed in the context of asymmetric mixed-strategy games, the fundamental reason for these results is inherent in all mixed-strategy solutions (including the symmetric case). At the heart of all mixed-strategy solution lies an important issue: Mixed-strategy solution demands that a firm (Firm L) randomizes its strategies such that the other firm

(Firm H) is indifferent to all the strategies in its (Firm H's) support. It is important to note why this issue is not apparent in symmetric mixed-strategy games. In a symmetric mixed-strategy game, when the payoff-related parameter of one firm is changed, the payoff-related parameter of the other firm is concomitantly changed. Thus, looking at the symmetric case it is difficult to say whether each firm is mixing its strategy to make itself indifferent to all the strategies in its support *or* to make the other firm indifferent to all the strategies in the other firm's support. Therefore, in order to examine the independent effects of a firm's parameters on the equilibrium strategies it is necessary to move to asymmetric games.

In our model, we assumed that if firms invest the same amount in R&D, then they compete away all potential profits. It is conceivable that in some instances of ties, the competing firms get a share of the market instead of dissipating all potential profits. For instance, if both firms invest an amount greater than zero and there is a tie, then player i receives reward tr_i where $0 \leq t \leq \frac{1}{2}$. As long as there are no dominated strategies, the qualitative implications of our theoretical results still hold.

Another conceivable asymmetry in the context of our patent race model is asymmetry in the amount of capital available to each firm. That is, $c_H \neq c_L$. Firms may have access to different amounts of capital because of the financial constraints imposed by the capital market. Informational asymmetries surrounding the research projects (Lerner and Merges 1998) and prior reputation of the firms could potentially lead to such asymmetries in capital constraints. We now briefly consider the implications of such asymmetries. Consider the case where two firms with different capital constraints, but with the same valuation for the patent, compete to develop a product. Without loss of generality, assume that $c_H > c_L$ and $r_H = r_L = r$. Details of the equilibrium solution for this extension are presented in Technical Appendix C.[6] We find that the firm with a financial advantage invests more in innovation and is more likely to win the patent races. Thus, our results show that while market advantage can impede innovation, financial advantage can stimulate firms to invest more in innovation.

3. Empirical Investigation

Our goal is to examine to what extent the actual behavior of financially motivated agents conforms to the game-theoretic predictions. Toward this goal, we test our model in a controlled laboratory environment, so that we can observe the actual investment decisions of economic agents and make causal attributions on the competitive outcomes (e.g., Bolton and Zwick 1995, Smith 1982, Smith 1989, and Roth 1995, p. 302). We derive the Nash equilibrium solution assuming complete information. It is unlikely that our subjects will solve for equilibrium behavior and accordingly make their investments. It is possible that their investment decisions are guided by some heuristics that have limited normative basis. For example, subjects might think that high-reward subjects should invest more than the low-reward subjects. After all, the high-reward subjects benefit more from winning the competition. If high-reward subjects actually invest more, then we might observe a positive relationship between reward size and investments—a relationship that is counter to the theoretical prediction. Consequently, the high-reward subjects would win the competition more frequently. Thus, it is not clear that subjects would invest as predicted by theory. Hence, we subject the model to a laboratory test.

[6] If the strategy space is continuous, then in this case the equilibrium (mixed) strategy solutions are:

$$F_L(x) = \begin{cases} (r - c_L)/r & \text{if } x = 0, \\ (r - c_L + x)/r & x \in (0, c_L), \\ 1 & \text{otherwise,} \end{cases}$$

and

$$F_H(x) = \begin{cases} x/r & x \in [0, c_L), \\ 1 & x \geq c_L. \end{cases}$$

F_H stochastically dominates F_L, and consequently Firm H invests more than Firm L and wins the patent more often. The reason why the results are different in this case is that Firm H is assured of positive profits $(r - c_L)$, and consequently the profits are not dissipated away. Therefore, the support of the mixed-strategy distribution for *both* firms depends on c_L. In a discrete version of this game, players would have to iteratively delete the dominated strategies and then use mixed strategies over the reduced strategy space. For an empirical test of such a game see Rapoport and Amaldoss (2000).

Parameters. For the experimental study, we use the following parameter values: $c = 2$, $r_H = 7$, $r_L = 4$, and $k = 2$. Then, the equilibrium solution for the high-reward player is:

$$p_H = (p_H(0), p_H(1), p_H(2)) = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{2}\right). \tag{14}$$

Similarly, the equilibrium behavior for the low-reward player in this example is:

$$p_L = (p_L(0), p_L(1), p_L(2)) = \left(\frac{1}{7}, \frac{1}{7}, \frac{5}{7}\right). \tag{15}$$

The theoretical mean investments of the low-reward and high-reward players are 1.57 and 1.25, respectively. Hence, the low-reward player invests more than the high-reward player, as implied by our Theoretical Result 2. Also, the low-reward player should win the competition 39.29% of the times, while the high-reward player should win the competition 17.86% of the times. We limited the strategy space to three levels so that the decision problem is cognitively simpler. Even with three levels of investment strategy, the decision problem is nontrivial in this asymmetric game.[7]

Subjects. Thirty-six undergraduate and graduate students participated in the study. Subjects were recruited through advertisements placed on bulletin boards on campus and by class announcements. They were promised a monetary reward contingent on their performance in a decision making experiment. In addition to their earnings, the subjects were promised a show-up fee of \$5.00.

Procedure. We conducted the study in two groups so that we obtained two replications of the same experiment. Each group comprising 18 subjects participated in a single session that lasted about 90 minutes. We held the experiments in a laboratory with a computer facility for multiplayer interactive decision making. Upon arrival at the laboratory, the subjects were randomly seated in separate computer booths. Before the commencement of the experiment the subjects were asked to read the instructions for the experiment (see Technical Appendix D), which included examples. After reading the instructions, subjects participated in five practice trials to familiarize themselves with the decision task. The researcher answered any questions about the experimental procedure. Any form of communication between subjects was not possible. In this experiment all investments were made in terms of a laboratory currency called "franc." At the end of the experiment, the individual payoffs were totaled and converted into U.S. dollars.

The subjects were randomly divided into nine low-reward players and nine high-reward players, and their reward condition remained the same throughout the experiment. On each trial each low-reward subject was randomly assigned to compete with a high-reward subject. Subjects had no way of knowing the identity of their competitors on any given trial. The assignment schedule ensured that each high-reward subject would be matched with each of the low-reward subjects approximately the same number of times and that a subject would not be paired with the same competitor twice in a row.

At the beginning of each trial the subjects were provided with a capital of two francs. The prize for winning the competition was four francs for the low-reward players ($r_L = 4$) and seven francs for the high-reward players ($r_H = 7$). Subjects were informed that whoever invested more would win the competition and secure the prize. Ties were counted as losses. Furthermore, the investments were nonrecoverable, irrespective of the outcome of the competition.

Each subject decided how much to invest based on the capital at hand and the prizes (r_L and r_H) for winning the competition. After the subjects made their investments privately, the computer compared each player's investment with that of the competitor and determined the winner. As discussed earlier, ties were counted as losses. At the end of each trial, each subject was informed of the investments made by the two competing players, the winning player, his or her payoff for the trial, and the cumulative payoff.

[7] Experimental results in the context of symmetric strategies have shown that the results are robust to increases in the levels of strategy space. For example, Rapoport and Amaldoss (2000) find experimental support for a symmetric mixed strategy game that allows for six levels of strategy.

This stage game was played repeatedly for 80 trials. During these 80 trials the reward and capital remained the same, though the competitor changed randomly from trial to trial. At the end of the experiment, the subjects were paid according to their cumulative earnings, debriefed, and dismissed.

Results

In this section we study the ability of the model to predict the investment decisions, winning frequencies, and payoffs of our subjects. We first test the model at the aggregate level, and then at the individual level. Each trial of our experiment is a replication of the game, and it provides a data point on the aggregate behavior of low-reward subjects and another data point on high-reward subjects. We use these data points to assess the ability of the model to predict the aggregate behavior of subjects. Later, we use the average behavior of each subject over the several iterations of the game to perform individual-level analysis.

Investment Strategy Profiles. The model makes point predictions about the probabilities of investing zero, one, and two francs. We find support for these predictions at the aggregate level, with a few exceptions in the case of high-reward players. We first examine the actual investments made by the low-reward and high-reward subjects. Then we contrast the investments under the two reward conditions.

Table 2 presents the empirical distribution of the investment choices made by low-reward players (upper panel) and high-reward players (lower panel). Columns 2, 3, and 4 present the observed relative frequencies for the first 40 trials, the next 40 trials, and the total 80 trials of Group 1. Similarly, Columns 5, 6, and 7 show the corresponding frequencies for Group 2. Column 8 presents the empirical distribution across both groups, and Column 9 presents the theoretical distribution. We discuss below in detail the investment strategy profiles.

Low-Reward Subjects ($r_L = 4$). In the aggregate, the low-reward subjects conform to the model predictions. As per theory, subjects should invest zero, one, and two francs with probabilities 0.143, 0.143, and 0.714, respectively. Across the two groups, subjects played these strategies on 22.85, 7.08, and

Table 2 Aggregate Distribution of Investments

For Low-reward players, $r_L = 4$.

	Empirical Distribution							
	Group 1			Group 2				
Investment	First 40 Trials	Second 40 Trials	80 Trials	First 40 Trials	Second 40 Trials	80 Trials	Both Groups (160 Trials)	Theoretical Prediction
0	0.2600	0.2389	0.2500	0.2917	0.1222	0.2069	0.2285	0.1429
1	0.0750	0.0167	0.0458	0.1194	0.0722	0.0958	0.0708	0.1429
2	0.6639	0.7444	0.7042	0.5889	0.8056	0.6972	0.7007	0.7143
Winning Frequency	0.1917	0.3333	0.2625	0.3250	0.6222	0.4736	0.3681	0.3929

For High-reward players, $r_H = 7$.

	Empirical Distribution							
	Group 1			Group 2				
Investment	First 40 Trials	Second 40 Trials	80 Trials	First 40 Trials	Second 40 Trials	80 Trials	Both Groups (160 Trials)	Theoretical Prediction
0	0.0417	0.0750	0.0583	0.1278	0.2278	0.1778	0.1181	0.2500
1	0.2611	0.3972	0.3292	0.4000	0.5500	0.4750	0.4021	0.2500
2	0.6972	0.5278	0.6125	0.4722	0.2222	0.3472	0.4799	0.5000
Winning Frequency	0.2833	0.2361	0.2597	0.3083	0.1111	0.2097	0.2347	0.1786

70.07% of the trials. The observed and predicted distributions are not statistically different ($D_{160} = 0.0856, p > 0.1$). We obtain similar results for Group 1 ($D_{80} = 0.107, p > 0.2$) and Group 2 ($D_{80} = 0.064, p > 0.2$). In Table 2 we also present the empirical distribution of strategies for the first 40 and the next 40 trials of each group. We cannot discern significant differences between the theoretical and empirical distribution of investments in either halves of the 80 trials of both groups ($p > 0.2$).

In equilibrium play, the low-reward subjects should make a mean investment of 1.57 francs. Using the aggregate behavior in each trial as the unit of analysis, we find that across both groups subjects invested 1.47 francs on average ($z = 1.06, p > 0.3$). Even at the level of individual groups, there is no statistical difference between the predicted and actual investments: Group 1 (mean $= 1.454, z = 1.21, p > 0.1$), and Group 2 (mean $= 1.4902, z = 0.89, p > 0.1$).

High-Reward Subjects ($r_H = 7$). In equilibrium, the high-reward players should invest zero, one, and two francs in proportions 0.25, 0.25, and 0.5, respectively. In actuality, the high-reward subjects across the two groups invested zero, one, and two francs in proportions 0.12, 0.40, and 0.48. This tendency to invest zero francs less frequently than predicted is seen in both groups. In Group 1 and Group 2 subjects invested zero francs in 5.83% and 17.78% of the trials, respectively, instead of 25% of the trials as predicted by theory. The discrepancy between the observed and predicted distribution of investment across the two groups is significant ($D_{160} = 0.1319, p < 0.01$). Analyzing the two groups separately, we find that for Group 1 the departure from equilibrium prediction is significant ($D_{80} = 0.1917, p < 0.1$). However, the deviation from equilibrium prediction is significant only in the first 40 trials ($D_{40} = 0.2083, p < 0.1$), but not in the later 40 trials ($D_{40} = 0.1171, p > 0.2$). For Group 2 we find that the departure from equilibrium prediction is not significant ($D_{80} = 0.1528, p > 0.1$).

We also examined the mean investment made by the two groups. Again using the aggregate behavior in each trial as the unit of analysis, we find that the mean investment of high-reward players in Group 1 is significantly different from the predicted mean investment of 1.25 francs (mean $= 1.554, z = 4.5, p < 0.001$). The results are similar, even if we look at the first 40 trials ($p < 0.001$) and the last 40 trials ($p < 0.021$). However, the mean investment of high-reward players in Group 2 (1.169) is not different from the model prediction ($z = 1.02, p > 0.16$). Thus, we find support for the model in Group 2, but the departures from the model predictions are significant in Group 1.[8]

Comparison of the Low-Reward and High-Reward Subjects. Across the two groups, we find that the low-reward subjects invested zero, one, and two francs on 22.85, 7.08 and 70.07% of the trials, respectively. On the other hand, the high-reward subjects made these investments on 11.81, 40.21, and 47.99% of the trials, respectively. The Kolmogorov-Smirov two-sample test rejects the null hypothesis that these two distributions are similar ($D_{80,80} = 0.2209$, $p < 0.01$). Furthermore, the low-reward subjects (mean = 1.47) invested more than the high-reward subjects (mean = 1.36). While the overall mean investments are directionally consistent with theory, the trend in mean investment over the first and last 40 trials is interesting. In the first 40 trials the high-reward subjects, rather than the low-reward subjects, invested more. The difference in investment is significant in Group 1 (low-reward subjects' mean = 1.40, high-reward subjects' mean = 1.66, $z = 1.545$, $p < 0.07$) but not in Group 2 (high-reward subjects' mean = 1.344, low-reward subjects mean = 1.29, $z = 0.265$, $p > 0.2$). In the later 40 trials, the low-reward subjects invested more than the high-reward subjects. The difference in investment is significant in Group 2 (low-reward subjects' mean = 1.68, high-reward subjects' mean = 0.99, $z = 4.57$, $p < 0.001$) but not in Group 1 (low-reward subjects' mean = 1.506, high-reward subjects' mean = 1.453, $z = 0.31$, $p > 0.2$).[9]

[8] We obtain similar results if we perform the analysis on the mean investments of individual subjects averaged over the 80 trials of the game (overall mean for low-reward subjects = 1.47, $t = 0.50$, $p > 0.2$; overall mean for high-reward subjects = 1.554, $t = 2.14$, $p < 0.02$).

[9] We can reject the null hypothesis that high- and low-reward players in Group 2 made similar investments, when this analysis is performed on the mean investments of individual subjects ($p < 0.06$). We cannot reject such a null hypothesis for Group 1 ($p > 0.2$).

Winning Frequency. Theory suggests that low-reward players should win more often. More precisely, the low-reward subjects are expected to win on 39.29% of the occasions, while the high-reward subjects are expected to only win on 17.86% of the occasions. In general the empirical results are directionally consistent with the model predictions, except for a major exception in Group 1. In Table 2, we report the winning frequencies in the last row of each of the panels.

In each trial, there were nine competitions and we computed the aggregate winning frequency in each trial for low- and high-reward players. We used the average of these trialwise aggregate winning frequencies to assess the predictive accuracy of the model. Across both groups, the low-reward subjects won on 36.81% of the occasions. We cannot reject the null hypothesis that the actual and predicted frequency of winning are equal ($z = 0.650$, $p > 0.2$). The high-reward subjects won on 23.47% of the occasions, and that is more than the theoretical prediction ($z = 1.675$, $p < 0.05$). However, there is some variation between groups.

In Group 1, the low-and high-reward subjects won on 26.25 and 25.97% of the occasions, respectively. We can reject the null hypothesis that these proportions are not different from the theoretical prediction. On further examination, we find that the departure from theoretical norm is significant in the first 40 trials but not in the last 40 trials. In the last 40 trials of Group 1, the low-reward and the high-reward subjects won on 33.33 ($p > 0.2$) and 23.61 ($p > 0.2$) % of the trials, which is not different from the theoretical norm.

In Group 2, theory closely tracks the actual winning frequencies. The low-reward subjects won 47.36% of the time while the high-reward subjects won 20.97% of the time. We cannot reject the null hypothesis that the observed frequencies are not different from the theoretical norm ($p > 0.2$).

Payoffs. Payoff depends on the investments of the two competing players. It is possible that the payoffs earned may not conform to the theoretical prediction, even if the aggregate distribution of investments follows the theoretical norm. An important implication of the equilibrium solution is that the mean payoff per trial is independent of the reward ($r_L = 4$, $r_H = 7$) for winning the competition.

The predicted payoff of low-reward players is two francs per trial. The actual mean payoff of the low-reward players across the two groups is not different from the theoretical prediction (combined mean = 1.999, $z = 0.002$, $p > 0.2$). The expected payoff of high-reward players is the same as that for low-reward subjects. The actual mean payoff of both groups of high-reward players is 2.281. This mean does not differ significantly from the theoretical value ($z = 0.763$, $p > 0.2$).[10]

Individual Differences. So far we have examined the aggregate behavior of subjects, but such an analysis does not reveal individual-level variations in behavior. For instance, it is conceivable that individual subjects playing the same pure strategies from trial to trial could well have produced an aggregate distribution that conforms to the mixed strategy solution. Across the two groups, there are 18 low-reward subjects and 18 high-reward subjects. In Figure 1 we report the frequency distribution of the number of times low- and high-reward subjects invested two francs.

We observe considerable individual differences, with the frequency of investing the entire budget ranging all the way from 11 to 80. Six low-reward subjects invested two francs 71–80 times. However, only one of them played the pure strategy of investing two francs in all 80 trials. Of the three high-reward subjects in the 71–80 frequency class, only two played the pure strategy of investing two francs in all 80 trials. Although the behavior of individual subjects shows substantial heterogeneity, we still observe in Figure 1 that low-reward subjects tend to invest more than the high-reward subjects. We also examined whether the distribution of the investments of each individual subject conforms to the theoretical prediction. We can reject the null hypothesis that the predicted and actual

[10] We also find that the empirical and the theoretical distributions of payoffs are not significantly different in the case of low-reward subjects ($D_{160} = 0.06$, $p > 0.2$) and high-reward subjects ($D_{160} = 0.075$, $p > 0.2$).

Figure 1 Distribution of Subjects by How Frequently They Invested the Entire Capital

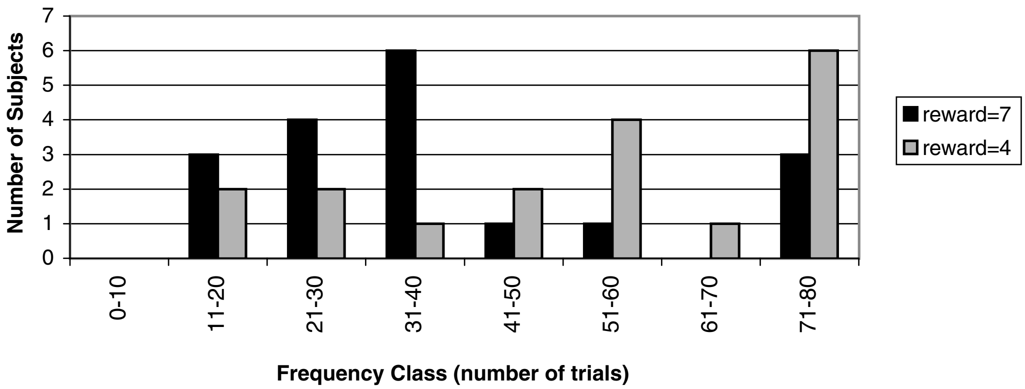

distributions are similar for 10 of the low-reward subjects and 12 of the high-reward subjects ($p < 0.1$).[11]

When we examine the winning frequency, we again see variations at the level of individual subjects. In Figure 2, we present the distribution of subjects by the proportion of times they won the competitions. The proportion of times the low-reward subjects won the competitions ranges all the way from 0 to 65%, whereas the corresponding range for high-reward subjects is 13.75% to 30%. We can reject the null hypothesis that the predicted and actual frequency of winning is similar for 11 of the low-reward subjects and seven of the high-reward subjects ($p < 0.1$).

We also see much variation in the individual payoffs of subjects. The mean payoff of low-reward players ranges from 1.32 to 2.84, whereas that of high-reward subjects ranges from 1.43 to 2.75. In Figure 3 we report the distribution of the mean payoffs of the 36 subjects. The differences in the mean payoffs between subjects is significant ($p < 0.01$), and paired comparisons reveal more than two distinct clusters ($p < 0.05$).

[11] In general, a mixed-strategy solution can be interpreted in three different ways (see Rosenthal 1979, Harsanyi 1973, Aumann et al. 1983, and Osborne and Rubinstein 1994). None of these interpretations demand that each individual subject perfectly randomizes strategies. Indeed, in the experimental study we observe that the equilibrium solution survives at the aggregate rather than the individual level. Furthermore, we also observe first-order sequential dependencies in the choice of the subjects ($p < 0.001$). This is despite the fact that in our experiment, the subject pairings were deliberately changed from trial to trial to eliminate sequential dependencies.

The individual-level analysis of investment decisions, payoffs, and winning frequencies suggests that although the model seems to survive at the aggregate level, it may fail at the individual level. Similar observation has been made in experimental tests of symmetric mixed-strategy games (e.g., O'Neill 1987, Camerer 1990, Rapoport and Boebel 1992, Rapoport and Budescu 1992, Rapoport and Amaldoss 2000, Amaldoss et al. 2000).

Discussion. The key inferences from the laboratory test are summarized below:

(1) *Do low-reward subjects invest more than the high-reward subjects on the average?* At the aggregate level, the low-reward subjects invested more than the high-reward subjects, as predicted by theory. Interestingly, consistent with the naive analysis, in the first 40 trials the high-reward subjects tended to invest more, although the difference is significant only in Group 1.

(2) *Do low-reward subjects win the competition more often than the high-reward subjects?* Theory predicts that low-reward subjects should win more often. Our experimental results provide support for this hypothesis, with a major exception in Group 1. In the first 40 trials of Group 1, the high-reward subjects won more often. However, in the later 40 trials the behavior of subjects in Group 1 was consistent with the theoretical predictions.

(3) *Does the aggregate distribution of investment strategies conform to the point predictions of the model?* The theoretical predictions track the aggregate distribution of investments made by subjects. However, we

Figure 2 **Distribution of Subjects by Frequency of Wins**

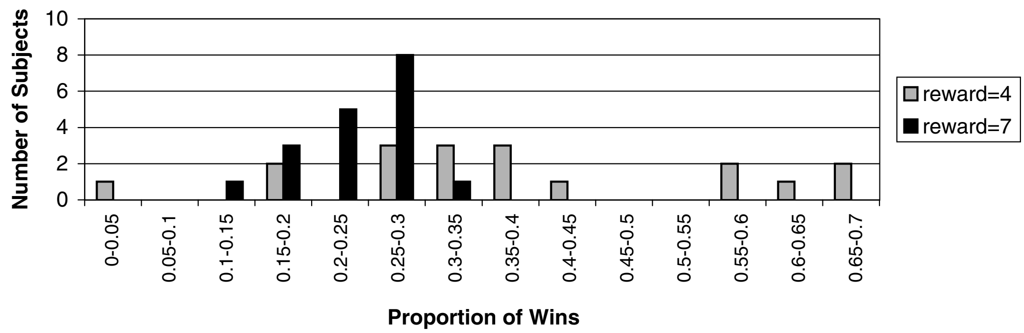

find some exceptions among high-reward players of Group 1, especially in the first 40 trials.

(4) *Does the distribution of payoffs conform to the predicted payoffs?* The payoff is a function of the joint distribution of the investments of competing players. We find that the actual payoffs received by subjects conform to the normative benchmark.

Thus, we find support for most of the major implications of the model at the aggregate level. However, in the individual-level analysis of investments, payoffs, and winning frequencies we see significant departures from the theoretical predictions. Hence, the model seems to survive at the aggregate rather than the individual level.

There is yet another implication of the equilibrium solution that we have not examined: As the reward for low-reward players increases, the high-reward players should invest more while low-reward players' investment should remain unaffected. To see this, let us increase the reward for low-reward players from four to six, such that $c = 2$, $r_L = 6$, and $r_H = 7$. The resulting equilibrium behavior for high- and low-reward players is as follows:

$$p_H = (p_H(0), p_H(1), p_H(2)) = \left(\frac{1}{6}, \frac{1}{6}, \frac{4}{6}\right);$$

$$p_L = (p_L(0), p_L(1), p_L(2)) = \left(\frac{1}{7}, \frac{1}{7}, \frac{5}{7}\right).$$

To assess the descriptive validity of this equilibrium prediction, we conducted another study with a group of 18 subjects for 80 trials, with 9 of the subjects assigned to the low-reward condition and the other 9 of them assigned to high-reward condition. To save space, we will briefly highlight the key findings. The high-reward players invested zero, one, and two francs on 5.28, 29.72, and 65% of the trials, respectively. We cannot reject the null hypothesis that the actual and predicted investments are similar ($D_{80} = 0.11$, $p > 0.2$). Furthermore, this investment pattern is different from that observed in high-reward subjects when $r_L = 4$ ($D_{160,80} = 0.17$, $p < 0.1$). On examining the mean investments of individual subjects,

Figure 3 **Distribution of Subjects by Mean Payoff**

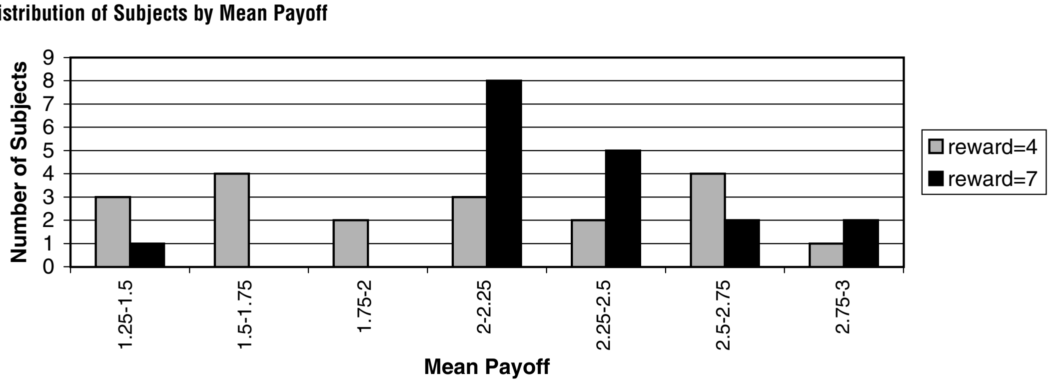

we observe that the average of the mean investments of high-reward players increased from 1.362 to 1.597 as r_L grew from four to six francs ($p < 0.1$). Thus, as predicted by theory, high-reward players invest more aggressively as r_L increases. Also, theory predicts that we should not see significant change in the behavior of low-reward players since r_H remains unchanged. The low-reward players in this game invested zero, one, and two francs on 25.14, 12.36, and 62.5% of the trials. On comparing this empirical distribution of investments with that of low-reward players when $r_L = 4$, we cannot reject the null hypothesis that these distributions are similar ($D_{160,80} = 0.075$, $p > 0.2$). Thus, overall the results of this experiment are directionally consistent with the comparative statics predictions of the theoretical model. Furthermore, we cannot reject the null hypothesis that the predicted and actual investments of low-reward players in this game are similar ($D_{80} = 0.11$, $p > 0.2$).

Earlier we discussed that in some situations it might be appropriate to allow competing firms to get a share of the market instead of dissipating all potential profits. To explore the behavioral implications of such a game, we considered the case where $c = 2$, $r_L = 2.2$, $r_H = 2.9$, and $t = \frac{1}{2}$. The equilibrium solutions for this case are:

$$p_H = (p_H(0), p_H(1), p_H(2)) = (0.0909, 0.7272, 0.1818);$$

$$p_L = (p_L(0), p_L(1), p_L(2)) = (0.3103, 0.0689, 0.6207).$$

Note that Firm L should invest, on the average, 1.3103, while Firm H should invest 1.0909. We conducted an experiment with a group of 18 subjects for 80 trials to test whether actual behavior conforms to these theoretical predictions. Nine subjects played the role of low-reward players and another nine subjects played the role of high-reward players. The findings of this additional study are qualitatively consistent with our earlier experimental results. The low-reward players invested zero, one, and two francs on 37.22, 10.69, and 52.08% of the trials. We cannot reject the null hypothesis that the actual and predicted investments are the same ($D_{80} = 0.10$, $p > 0.2$). The high-reward subjects invested zero, one, and two francs on 37.22, 24.58, and 38.19% of the trials, and their behavior departs from the point predictions of our model ($D_{80} = 0.10$, $p < 0.01$). We note that the low-reward players invested, on the average, more than the high-reward players (low-reward players' mean = 1.1485, high-reward players' mean = 1.0908, $z = 1.4307$, $p < 0.1$). We see evidence of first-order sequential dependency (low-reward players: $\chi^2 = 25.75$, $p < 0.001$; high-reward players: $\chi^2 = 8.73$, $p < 0.07$), and also shifts in the investment decisions over the trials, implying that our subjects might be learning over the several replications of the game.

4. Adaptive Learning

Recall that a nonstrategic analysis of the game suggests that low-reward subjects should invest less than the high-reward players and should also win less often. A strategic analysis, however, shows that the opposite should hold in equilibrium. Our subjects seem to broadly conform to the predictions of the strategic analysis. Probably it is too much to expect our naive subjects to compute the equilibrium solution. Then, how did they come to play strategies that are even close to the equilibrium solution? Our analysis of investment trends suggests that our subjects moved toward equilibrium behavior over the several iterations of the game.[12] In other words, our subjects might have *learned* to play the game over the course of the experiment. We explore this possibility in this section.

Our goal is to understand whether adaptive learning models can account for the investment trends observed in our experiment. Toward this goal, we estimated the parameters of the experience-weighted attraction (EWA) learning model (Camerer and Ho 1999).[13] We chose the EWA for two

[12] In each group 18 subjects participated in 80 trials. To test for learning effects, we divided the 80 trials into eight blocks of 10 trials. Then we conducted an analysis of variance with one between-subject factor (reward condition) and one within-subject factor (block) for each of the groups. The dependent variable was investment level. We find the interaction effect of block and reward is significant in both groups (Group 1: $F_{(7,112)} = 2.15$, $p = 0.0436$; Group 2: $F_{(7,112)} = 10.70$, $p = 0.0001$). Yet the main effect of block is not significant ($p > 0.3$). These results suggest that the high- and low-reward subjects learned to invest, but in different ways, over the eight blocks of trials.

[13] Learning in games is a rapidly growing field of research and several different classes of learning models have been proposed

principal reasons:

(1) The EWA model allows for three important learning effects: actual, simulated, and declining effects. The first effect, common in all reinforcement-based models, implies that past successes with a strategy increase its chance of being chosen in the future. The second effect, common in belief-based learning models, implies that unchosen strategies, which might have yielded high payoffs, are more likely to be chosen in the future. The third effect implies that, with experience, players move to reduce discrepancies between actual and foregone payoffs. The EWA model hybridizes the intuition of both reinforcement- and belief-based learning models and includes these as special cases.[14]

(2) The EWA model has been applied in various contexts and has been shown to have good predictive accuracy (see Camerer et al., forthcoming for the predictive accuracy of the model in 31 different data sets spanning over a dozen different types of games).

We discuss below the model estimates and their implications for adaptive learning.

The EWA Learning Model. The EWA model uses the history of the game to estimate the probability that a player i will invest an amount x_i in the next period. The probability of player i investing $x_i = m$ on trial $t+1$ can be given by the logit function:

$$p_i^m(t+1) = \frac{e^{\lambda A_i^{x_i^m}(t)}}{\sum_{j=0}^{c} e^{\lambda A_i^{x_i^j}(t)}}, \tag{16}$$

where $A_i^{x_i^m}(t)$ is the attraction of investing $x_i = m$ at time t for player i. The parameter λ measures sensitivity of the players to attractions. The parameter λ can also be interpreted as a measure of noise in the strategy choice process. At the end of every trial, a player updates the attractiveness of a strategy based on the actual payoff and also the expected payoffs corresponding to strategies that were not chosen. While updating the attraction of a strategy, payoffs corresponding to chosen strategies are given a weight equal to one, while expected payoffs corresponding to unchosen strategies are given a weight of δ $(0 \le \delta \le 1)$. The parameter δ measures the relative weight given to foregone payoffs, compared to actual payoffs, in updating attractions. It can be interpreted as a kind of "imagination" of foregone payoffs. Previous attractions are depreciated by another parameter ϕ $(0 \le \phi \le 1)$. The decay rate ϕ reflects a combination of forgetting and the degree to which players recognize that other players are adapting and thereby place lower weight on the history of the game. If ϕ is lower, players discard old observations more quickly and become more responsive to the most recent observations.

In updating the attractiveness of a strategy the model also uses another parameter denoted by $N(t)$. This parameter is interpreted as the number of observations-equivalents of past experience. In other words, this is a weighted combination of the number of times the game has been played. The updating rule employed is:

$$N(t) = \rho N(t-1) + 1, \qquad t \ge 1, \tag{17}$$

where the parameter ρ $(0 \le \rho \le 1)$ is the rate of depreciation. This parameter $N(t)$ is used to combine the prior attractiveness with the actual/expected payoffs from the current period. The attraction of investing x_i^m, namely $A_i^{x_i^m}(t)$, is a weighted average of the payoff for period t and the previous attraction $A_i^{x_i^m}(t-1)$:

$$A_i^{x_i^m}(t) = \frac{\phi N(t-1) A_i^{x_i^m}(t-1) + [\delta + (1-\delta) I(x_i^m, x_i(t))]\pi_i(x_i^m(t), \mathbf{x}_{-i}(t))}{N(t)}, \tag{18}$$

where $\pi_i(x_i^m(t), \mathbf{x}_{-i}(t))$ is the payoff received by firm i by investing $x_i^m(t)$ in period t given that the other players invested $\mathbf{x}_{-i}(t)$ in time period t. The

(e.g., Camerer, et al. forthcoming; Fudenberg and Levine 1998; Ho and Weigelt 1996; Roth and Erev 1995; Stahl 1996, 1999, 2000).

[14] The key point of EWA is that the seemingly different belief-based and reinforcement-based learning are indeed very closely related: Reinforcement models are based on some weighted average of previously received payoffs and exclude foregone payoffs, whereas in belief models the expected payoffs are exactly equal to a weighted average of previously received payoffs and foregone payoffs of strategies that were not chosen. Thus, the fundamental difference between reinforcement and belief models is the extent to which *foregone* payoffs are assumed to influence the choice of strategies. In the EWA model, the effect of foregone payoffs on choice of strategies is captured in a single parameter, $\delta \in [0, 1]$.

$I(x_i^m, x_i(t))$ function is an indicator variable, which is defined as:

$$I(x_i^m, x_i(t)) = \begin{cases} 1 & \text{if } x_i(t) = x_i^m, \\ 0 & \text{otherwise.} \end{cases} \tag{19}$$

Thus, if player i invests x_i^m on trial t, then the payoff is added to the attraction of the strategy $x_i = m$, $A_i^{x_i^m}(t)$. However, if player i plays strategy $x_i \neq m$ on trial t, then only δ fraction of the payoff is added to the attraction of strategy $x_i = m$, $A_i^{x_i^m}(t)$. The initial values of the two parameters of the EWA model are denoted by $N(0)$ and $A_i^{x_i^m}(0)$.

Note that if $\delta = 0$, $\rho = 0$, and $N(0) = 1$, then the attractions are mathematically equivalent to the reinforcement of classical reinforcement-based learning models. If $\delta = 1$ and $\rho = \phi$, then the attractions are exactly the same as updated expected payoffs as in belief-based learning models (see Camerer and Ho 1999 for a proof of this claim).

Results

We calibrated the EWA model using the data from the first 50 trials of the experiment. We estimated the model parameters using the maximum likelihood method. Table 3 reports the parameter estimates and

Table 3 Parameter Estimates for the Learning Model

		Group 1			Group 2		
Reward	Parameter	EWA	Reinforcement-based	Belief-based	EWA	Reinforcement-based	Belief-based
	ρ	0.999*	0.000	0.990*	0.935*	0.000	0.990*
	ϕ	0.979*	0.922*	0.990*	0.900*	0.838*	0.990
	δ	0.000	0.000	1.000*	0.000	0.000	1.000*
	λ	9.302*	0.225*	0.038*	1.934*	0.237*	0.375*
	$N(0)$	65.276*	1.000*	100.000*	1.967*	1.000*	20.000*
	$A0(0)$	3.375*	0.001	67.780*	3.490*	1.113*	4.837*
	$A1(0)$	3.289*	0.001	21.209*	3.224*	0.001	0.212*
$r = 4$	$A2(0)$	3.553*	5.000*	99.571*	4.242*	5.000*	8.350*
	Log-likelihood	−282*	−299*	−358*	−286*	−294*	−419*
	AIC	−290*	−304*	−364*	−294*	−299*	−425*
	BIC	−306*	−314*	−377*	−310*	−310*	−437*
	pseudo-ρ^2	0.430*	0.396*	0.275*	0.422*	0.404*	0.152*
	χ^2		33.960*	153.264*		17.386*	266.760*
	(p-value, d.o.f)		(0.000, 3)	(0.000, 2)		(0.001, 3)	(0.000, 2)
	ρ	0.885*	0.000	0.990*	0.860*	0.000	0.990*
	ϕ	0.889*	0.878*	0.990*	0.858*	0.870*	0.990*
	δ	0.000	0.000	1.000*	0.000	0.000	1.000*
	λ	1.311*	0.175*	1.182*	0.963*	0.125*	1.634*
	$N(0)$	14.711*	1.000*	20.000*	21.147*	1.000*	25.000*
	$A0(0)$	2.331*	0.001	2.762*	2.678*	0.001	3.097*
	$A1(0)$	3.552*	0.001	5.563*	3.694*	2.611*	3.975*
$r = 7$	A2(0)	5.077*	8.000*	6.939*	4.588*	8.000*	3.890*
	Log-likelihood	−260*	−269*	−345*	−379*	−382*	−436*
	AIC	−268*	−274*	−351*	−387*	−387*	−442*
	BIC	−284*	−284*	−363*	−403*	−397*	−454*
	pseudo-ρ^2	0.475*	0.456*	0.303*	0.234*	0.228*	0.119*
	χ^2		18.480*	169.673*		6.612*	113.925*
	(p-value, d.o.f)		(0.000, 3)	(0.000, 2)		(0.085, 3)	(0.000, 2)

Notes. In reinforcement-based learning models the following three parameters are fixed: $\delta = 0$, $\rho = 0$, and $N(0) = 1$. In belief-based models, the fixed parameters are as follows: $\delta = 1$ and $\rho = \phi$.
*Significant with $\alpha = 0.01$.

goodness-of-fit statistics for the EWA model and its special cases—reinforcement-based and belief-based learning. The upper panel reports the results for the low-reward players, whereas the lower panel reports the results for the high-reward players. We conducted separate analyses for Group 1 and Group 2 in each of the reward conditions. We briefly summarize our results below.

Overall Model Fit. Our results suggest that the EWA model tracks the investments quite well, as evidenced in the fit statistics. Table 3 presents the log-likelihood (*LL*), Akaike Information Criterion (*AIC*), Bayesian Information Criterion (*BIC*), pseudo-$R^2(\rho^2)$, and χ^2 statistic for model comparisons. The pseudo-ρ^2 of the EWA model indicates the extent to which the EWA model can perform better than the null model.[15] We observe that the pseudo-ρ^2 for Group 1 and Group 2 of low-reward players are 0.430 and 0.422, respectively. The corresponding statistics for high-reward players in Group 1 and Group 2 are 0.475 and 0.234. These pseudo-ρ^2 values of the EWA models compare favorably with those reported in Camerer and Ho (1999), where the values range from 0.063 to 0.4858. The fit statistics also suggest that the EWA model, which hybridizes the basic intuitions of both reinforcement and belief learning, outperforms those special cases. Later, while discussing the predictive accuracy of the model, we will summarize the performance of the model within the calibration sample as well.

Interpretation of the Parameter Values. We observe that $\delta = 0$ in all four groups. This implies that the choice of strategies in these games were not influenced by expected payoffs. Rather, the investment decisions were based on the payoffs earned in previous trials. Such an inference is also validated by the Akaike Information Criterion (*AIC*), Bayesian Information Criterion (*BIC*), pseudo-R^2 (ρ^2), and χ^2 statistic reported for model comparisons.

[15] The pseudo-ρ^2 is the difference between the *AIC* measure and the log-likelihood (*LL*) of a model of random choices, normalized by the random model log-likelihood. $AIC = LL - k$, and $BIC = LL - (k/2)\log(M)$, where k is the number of degrees of freedom and M is the sample size.

We observe that in all four cases ϕ is approximately equal to ρ. For instance, for low-reward players in Group 1, $\phi = 0.979$ and $\rho = 0.999$. This implies that the strategy choices of our subjects were guided by the average performance of past strategies. The parameter λ measures the sensitivity of subjects to payoffs. Our estimates for λ range from 0.963 to 9.302. These values fall within the range of values reported in Camerer and Ho (1999): λ is 0.182 in 4×4 constant games, and 17.987 in median effort games. Our estimates suggest that the subjects were sensitive to the changes in payoffs.

Pregame Disposition. The estimates of the parameters $A0(0)$, $A1(0)$, and $A2(0)$ indicate the predisposition of the subjects to invest zero, one, and two francs, respectively, before the first trial of the experiment. These pregame disposition parameters are reported in Table 3. Low-reward subjects were predisposed to invest two units more often than either 0 or 1 units in a fashion that is qualitatively consistent with equilibrium prediction. The high-reward subjects also were predisposed to invest two units more often. It seems that sheer introspection might not lead to the type of counterintuitive behavior implied by an asymmetric mixed-strategy solution. Furthermore, the strength of pregame disposition as evidenced in the empirical estimates of $N(0)$ is more than 14 observations-equivalent, except in Group 2 of the low-reward subjects. This suggests that our subjects gradually learned to play the game based on past experience.

Model Validation. We validated the model in two types of holdout samples. First, we assessed the predictive accuracy of the model in the last 30 trials. Second, we validated the model on the group of subjects who were not used for calibration. That is, we predicted the behavior of subjects in Group 1 using the estimates of subjects in Group 2, and vice versa.

To assess the predictive accuracy of the model, we use the log-likelihood, average predicted probability of investment choices actually made by subjects, percentage of choices correctly predicted by the model (hit ratio), and the mean predicted investment. In addition, we report the performance of the model in

Table 4 Validation of the EWA Model

	Reward = 4		Reward = 7	
	Group 1	Group 2	Group 1	Group 2
Calibration sample of 50 trials				
Log-likelihood	−282	−286	−260	−379
Average probability	0.645	0.649	0.682	0.503
Hit ratio	0.733	0.756	0.778	0.629
Actual investment	1.420	1.313	1.613	1.296
Predicted investment	1.537	1.429	1.724	1.573
p-value	0.304	0.289	0.342	0.008
Holdout sample of 30 trials				
Log-likelihood	−185	−118	−226	−236
Average probability	0.624	0.788	0.579	0.465
Hit ratio	0.819	0.881	0.744	0.607
Actual investment	1.511	1.785	1.456	0.959
Predicted investment	1.880	1.904	1.622	0.893
p-value	0.018	0.474	0.241	0.504
Holdout sample of 18 subjects (80 trials)				
Log-likelihood	−515	−484	−601	−464
Average probability	0.734	0.570	0.558	0.596
Hit ratio	0.789	0.749	0.651	0.772
Actual investment	1.490	1.454	1.169	1.554
Predicted investment	1.744	1.476	1.288	1.660
p-value	0.034	0.845	0.953	0.346

the calibration sample, for it could be a useful benchmark in evaluating the performance in the holdout samples. In Table 4 we report these summary statistics for the calibration sample of 50 trials, holdout sample of 30 trials, and holdout sample of 18 subjects (members of the other group). The 18 subjects in the holdout sample played 80 trials. Overall, the model tracks the investment decisions in the validation sample quite well, as implied by these statistics.

The mean of the average predicted probability reported for the four calibration samples is 62%, and it is comparable to that observed in the holdout sample of trials (61.4%) and holdout sample of subjects (61.4%). Although there is some variation from group to group, we do not see a substantial deterioration in the average predicted probability in any of the groups. The hit ratios of the validation samples show marginal improvement over that of the calibration sample. For instance, the average hit ratio in the four calibration samples is 72.4%, whereas the corresponding average hit ratios for the holdout sample of trials and holdout sample of subjects are 76.3% and 74%, respectively. Next we examine how the actual investments of our subjects compare with EWA predictions. The results are mixed for the low-reward condition. For low-reward subjects, we can reject the null hypothesis that the predicted and actual investments are similar in each of the validation samples of Group 1 ($p < 0.04$), but not for Group 2 ($p > 0.4$). For the high-reward condition, we cannot reject such a null hypothesis ($p > 0.2$) in any of the validation samples.

As the EWA model performs well in the calibration and validation samples, it raises an interesting question: Is the underlying learning mechanism guiding the choices of low- and high-reward subjects similar? We see some decline in the predictive accuracy of the model if we use the model calibrated on low-reward subjects to predict the behavior of high-reward players, and vice versa. For instance, when we use the model calibrated on Group 1 of low-reward subjects to predict the choices of high-reward players, the average predicted probabilities and hit ratios decline. Specifically, the average predicted probability of the choices made in the holdout sample of 30 trials is 62.4% for Group 1 of low-reward subjects (see Table 4, Column 1). However, when we use the same model to predict the choices of high-reward subjects, the average predicted probability declines to 58.8% in Group 1 and 51.5% in Group 2. Similarly, the hit ratio drops from 81.9% to 75% for high-reward subjects. The differences between predicted and actual investments are significant in both the groups of high-reward players (Group 1: $p < 0.03$, Group 2: $p < 0.07$). We also see similar decline in predictive accuracy when we use the model calibrated on Group 2 of low-reward players to predict the choices of high-reward players. The average predicted probability is 78.8% in its own validation sample of 30 additional trials, but the average predicted probability declines to 68.3% and 61.6% when used to predict the choices of high-reward subjects in Group 1 and Group 2. The hit ratio moves downward from 88.1% to 77.8% and 78.8%, respectively. However, we cannot reject the null hypothesis that the predicted and actual investments are the same (Group 1: $p > 0.11$, Group 2: $p > 0.13$). We obtain similar results when we predict the behavior of low-reward subjects based on models calibrated on high-reward subjects. These results imply that the underlying adaptive learning mechanisms in the two reward conditions might not be very radically different, or else the predictive accuracy would have deteriorated substantially.[16]

Discussion. The EWA model accounts for the major behavioral regularities observed in our experiments. The model performs well in the hold sample of trials, as well as the holdout sample of subjects. The model parameters imply that the initial disposition of our subjects was not directionally consistent with the predictions of the equilibrium solution. We find that both high- and low-reward players were predisposed to invest two francs more often than either zero or one franc. However, over the trials they learned to behave in a fashion that is qualitatively consistent with the prediction. We also note that the choices were mostly guided by a weighted average of previously earned payoffs, rather than a combination of both previously received payoffs and foregone payoffs. We obtain similar results for the two additional studies that we conducted.[17] We noted earlier that a naive analysis of the game leads to behavior that is not consistent with equilibrium behavior. Still, our rather naive subjects were able to move toward equilibrium behavior, because they learned from their past experience.

5. Conclusion

This research was motivated by a desire to examine the descriptive validity of some of the counterintuitive implications of using mixed strategy when players are asymmetric. Toward this goal, we developed an illustrative model of a patent race, where two firms compete to develop and patent a product. However, one firm values the patent more because of its market advantages such as brand-name reputation and distribution network. We show that the firm which values the patent *less* is *more* likely to aggressively invest and also win the patent more often than the firm which values the patent more. The reason for this result is intimately related to the very concept of mixed-strategy solution—mixed-strategy solution demands that a firm randomize its strategy such that the *other* firm is indifferent to all the strategies in its

[16] If we neglect the prior dispositions and focus on the values of the three key parameters of the EWA model, namely ρ, ϕ, and δ, then we see that $\delta = 0$, and $\rho \simeq \phi$ with a value in the neighborhood of 0.9. This implies that the weighting mechanism is not very different in the various groups.

[17] Overall, the EWA model fits our data well. For the game in which $r_L = 6$ and $r_H = 7$, our estimates of the EWA model for high-reward subjects are: $N(0) = 0.24$, $A0(0) = 0.012$, $A1(0) = 6.194$, $A2(0) = 17.556$, $\rho = 0.160$, $\phi = 0.927$, $\delta = 0$, $\lambda = 0.0137$, *log-likelihood* $= -319$, *AIC*$= -327$, *BIC*$= -345$. The corresponding statistics for low-reward players are: $N(0) = 0.0484$, $A0(0) = 0.001$, $A1(0) = 5.938$, $A2(0) = 27.685$, $\rho = 0.001$, $\phi = 0.967$, $\delta = 0$, $\lambda = 0.074$, *log-likelihood* $= -402$, *AIC*$= -410$, *BIC*$= -428$. Similarly, the results for the game, in which players were allowed to share the reward in case of a tie, is as follows. Our estimates of EWA model parameters and fit statistics for high-reward subjects are: $N(0) = 1$, $A0(0) = 0.001$, $A1(0) = 7.864$, $A2(0) = 98.491$, $\rho = 0$, $\phi = 0.99$, $\delta = 0$, $\lambda = 0.013$, *log-likelihood* $= -762$, *AIC*$= -770$, *BIC*$= -786$. For low-reward subjects, we obtained the following EWA model parameter estimates and fit statistics: $N(0) = 0.181$, $A0(0) = 1.313$, $A1(0) = 0.001$, $A2(0) = 38.196$, $\rho = 0.785$, $\phi = 0.99$, $\delta = 0.141$, $\lambda = 0.117$, *log-likelihood* $= -675$, *AIC*$= -683$, *BIC*$= -699$.

support. Consequently, a firm's equilibrium strategy is not directly dependent on its own payoff-related parameters but rather depends on the parameters of the other firm. This feature of mixed-strategy solution, which is difficult to see in symmetric games, becomes very apparent in asymmetric games.

This counterintuitive nature of asymmetric mixed strategy solution, along with the difficulty of providing empirical support using field data, could potentially deter the use of the solution concept (see, for example, Farrell and Saloner 1988, p. 250).[18] However, a modeler would like to examine the effect of various asymmetries on the competitive behavior of firms. Naturally, some of these situations may require the use of asymmetric mixed-strategy games.[19] We find that the Nash equilibrium solution provides a reasonable account of the aggregate behavior of our subjects. With the help of the experience-weighted attraction (EWA) learning model proposed by Camerer and Ho (1999), we show that adaptive learning can account for the investment behavior of our subjects. The EWA model predicts remarkably well the choices of our subjects in the validation samples: holdout sample of trials and holdout sample of subjects.

Although the primary focus of our study was to examine the descriptive validity of asymmetric mixed-strategy solution, our model provides an interesting explanation for why firms with market advantages may fail to innovate—market advantage can provoke aggressive investment by competitors, and such a competitive behavior may impede innovation by the firm that enjoys market advantage. This result is consistent with some field observations. For instance, Lerner (1997) notes that leaders in the disk-drive market innovated less frequently, and often lost their leadership position during the period 1971–1995. A more direct support for the model results is provided by our initial laboratory test of asymmetric mixed-strategy games.

In sum, our theoretical analysis highlighted some of the key implications of asymmetric mixed-strategy solutions. These implications are not in agreement with a naive analysis of asymmetric mixed-strategy games, and with some of our intuition. Yet, as evidenced in the experimental investigation, subjects can learn to play these games in a manner that is consistent with equilibrium behavior. Furthermore, the EWA learning model accounts well for the investment behavior of our subjects.

[18] In the context of price promotions, researchers have searched for empirical regularities in field data that are consonant with the mixed-strategy equilibrium (e.g., Rao et al. 1995, Villas-Boas 1995, Raju et al. 1990). After surveying the empirical literature on price promotions, Rao et al. (1995) conclude that price promotions can be characterized as mixed strategies. Our laboratory investigation is a useful complement to these field studies. However, our focus is on asymmetric mixed-strategy games.

[19] To illustrate that asymmetric mixed strategies can potentially be useful in understanding a broader set of phenomena, in Appendices E and F we discuss stylized models of product standardization and consumer search. Again, the results there show that the analysis of asymmetric mixed-strategy games can lead to very counterintuitive results. Our experimental results at least provide some justification for the use of the mixed-strategy solution concept, even in such cases.

Acknowledgments

Both authors have contributed equally to the paper. They thank the departmental editor, the associate editor, three anonymous reviewers, and Ed Fox, Ambar Rao, and Amnon Rapoport, for their comments and suggestions. They also thank the National Science Foundation for its financial support (NSF Grant SES-089182/SES-0114138).

References

Amaldoss, W., R. Meyer, J. Raju, A. Rapoport. 2000. Collaborating to compete. *Marketing Sci.* **19**(2) 105–126.

Aumann, R., Y. Katznelson, R. Radner, R. Rosenthal, B. Weiss. 1983. Approximate purification of mixed strategies. *Math. Oper. Res.* **8**(3) 327–341.

Bolton, G., R. Zwick. 1995. Anonymity versus punishment in ultimatum games. *Games and Econom. Behavior* **10**(1) 95–121.

Bresnahan, T., S. Stern, M. Trajtenberg. 1997. Market segmentation and the sources of rents from innovation: Personal computers in the late 1980s. *RAND J. Econom.* **28** 517–544.

Brown, J., R. Rosenthal. 1990. Testing the minimax hypothesis: A reexamination of O'Neill's game experiment. *Econometrica* **58** 1065–1081.

Camerer, C. 1990. Behavioral game theory. R. Hogarth, ed. *Insights in Decision Making: Theory and Applications*. University of Chicago Press, Chicago, IL, 311–336.

——. 2001. *Behavioral Game Theory: Experiments in Strategic Interactions*. Princeton University Press, Princeton, NJ.

——, T. Ho. 1999. Experience-weighted attraction learning in games. *Econometrica* **87**(July) 827–874.

——, ——, J. Chong. Sophisticated EWA learning and strategic teaching in repeated games. *J. Econom. Theory* Forthcoming.

Cheng, L., M. Zhu. 1995. Mixed-strategy Nash equilibrium based upon expected utility and quadratic utility. *Games and Econom. Behavior* **9** 139–150.

Costa-Gomes, M., V. Crawford, B. Broseta. 2001. Cognition and behavior in normal-form games: An experimental study. *Econometrica* Forthcoming.

Deneckere, R., J. Peck. 1995. Competition over price and service rate when demand is stochastic: A strategic analysis. *RAND J. Econom.* **26**(1) 148–162.

Farrell, J., G. Saloner. 1988. Coordination through committees and markets. *RAND J. Econom.* **19**(2) 235–252.

Fudenberg, D., D. Levine. 1998. *The Theory of Learning in Games*. MIT Press, Cambridge, MA.

——, J. Tirole. 1985. Preemption and rent equalization in the adoption of new technology. *Rev. Econom. Stud.* **52** 383–401.

——, R. Gilbert, J. Stiglitz, J. Tirole. 1983. Preemption, leapfrogging and competition in patent races. *Eur. Econom. Rev.* **22**(1) 3–31.

Gilbert, R., D. Newberry. 1982. Preemptive patenting and the persistence of monopoly. *Amer. Econom. Rev.* **72** 514–526.

Golder, P., G. Tellis. 1993. Pioneer advantage: Marketing logic or marketing legend? *J. Marketing Res.* **30** 158–170.

Harsanyi, J. 1973. Games with randomly disturbed payoffs: A new rationale for mixed strategy equilibrium points. *Internat. J. Game Theory* **2** 1–23.

Heymann, P. E. 1985. The disk drive industry. Drexel Burnham Lambert, New York.

Ho, T., K. Weigelt. 1996. Task complexity, equilibrium selection, and learning: An experimental study. *Management Sci.* **42**(5) 659–679.

Johnson, E., C. Camerer, S. Sen, T. Rymon. 2001. Detecting backward induction in sequential bargaining with information search. *J. Econom. Theory* Forthcoming.

Lan, L., A. Kanafani. 1993. Economics of park-and-shop discounts: A case of bundled pricing strategy. *J. Transport Econom. Policy* **27** 291–303.

Lerner, J. 1997. An empirical exploration of a technology race. *RAND J. Econom.* **28**(2) 228–247.

——, R. P. Merges. 1998. The control of technology alliances: An empirical analysis of the biotechnology industry. *J. Indust. Econom.* **XLV**(1, 2) 125–156.

Ochs, J. 1995. Games with unique, mixed strategy equilibria: An experimental study. *Games Econom. Behavior* **10** 212–217.

O'Neill, B. 1987. A non-metric test of the minimax theory of two-person zerosum games. *Proc. National Acad. Sci.* **84** 2106–2109.

Osborne, M., A. Rubinstein. 1994. *A Course in Game Theory*. MIT Press, Cambridge, MA.

Raju, J. S., V. Srinivasan, R. Lal. 1990. The effects of brand loyalty on competitive price promotional strategies. *Management Sci.* **36** 276–304.

Rao, R., R. Arjunji, B. Murthy. 1995. Game theory and empirical generalizations concerning competitive promotions. *Marketing Sci.* **14**(3) G89–G100.

Rapoport, A., W. Amaldoss. 2000. Mixed strategies and iterative elimination of strongly dominated strategies: An experimental investigation of states of knowledge. *J. Econom. Behavior Organ.* **42** 483–521.

——, R. B. Boebel. 1992. Mixed strategies in strictly competitive games: A further test of the minimax hypothesis. *Games Econom. Behavior* **4** 261–283.

——, D. V. Budescu. 1992. Generation of random series by humans: A two-person zerosum game paradigm. *J. Experiment. Psych.: General* **121** 352–363.

Reinganum, J. 1983. Uncertain innovation and the persistence of monopoly. *Amer. Econom. Rev.* **73** 741–747.

——. 1989. The timing of innovation: Research, development, and diffusion. R. Schmalensee, R. Willig, eds. *Handbook of Industrial Organization*, Vol. 1. Elsevier Science, New York, 849–908.

Robinson, W., C. Fornell. 1985. Sources of market pioneer advantages in consumer goods industries. *J. Marketing Res.* **22**(3) 305–317.

Rosenthal, R. 1979. Sequences of games with varying opponents. *Econometrica* **47** 1353–1366.

Roth, A. 1995. Bargaining experiments. J. H. Kagel, A. Roth, eds. *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ, 253–348.

——, I. Erev. 1995. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games Econom. Behavior* **8** 164–212.

Smith, V. 1982. Markets as economizers of information: Experimental investigation of Hayek's hypothesis. *Econom. Inquiry* **20** 165–179.

——. 1989. Theory, experiment and economics. *J. Econom. Perspectives* 151–169.

Stahl, D. 1996. Boundedly rational rule learning in a guessing game. *Games Econom. Behavior* **12** 218–254.

——. 1999. Sophisticated learning and learning sophistication. Working paper, University of Texas at Austin, Austin, TX.

——. 2000. Rule learning in symmetric normal-form games: Theory and evidence. *Games Econom. Behavior* **32** 105–138.

Varian, H. 1980. A model of sales. *Amer. Econom. Rev.* **70** 651–659.

Villas-Boas, J. 1995. Models of competitive price promotions: Some empirical evidence from the coffee and saltine crackers markets. *J. Econom. Management Strategy* **4** 85–107.

Accepted by Martin Weber; received September 15, 2000. This paper was with the authors 13 months for 2 revisions.