

On Rationality, Learning and Zero-Sum Betting - An Experimental Study of the No-Betting Conjecture

By: Doron Sonsino, Ido Erev, and Sharon Gilat *

Faculty of Industrial Engineering and Management
Technion, Israel Institute of Technology
Haifa 32000, Israel

June 2001

Sebenius and Geanakoplos (1983) prove that zero-sum betting is an irrational phenomenon. The current paper investigates the descriptive implications of this interesting result. For this sake, we have run six experiments (covering 21 different experimental conditions) in which human subjects played repeatedly the Sebenius and Geanakoplos betting game for 250 rounds. The results (for all 21 conditions) reveal a high initial betting rate that decreases very slowly with experience. Our experimental manipulations (that included changes in payoff information, game complexity, feedback system, incentive scheme, outside options and more) had a significant effect on the initial betting rates in some of the conditions, but did not affect the observed slow learning process. All together, the results demonstrate that human agents' behavior may persistently be very different from the one implied by "iterated removal of dominated strategies" even when it leads the agents into significant losses (up to 39% of the initial endowment). We show, however, that under the assumption that the learning speed decreases with payoff variability, simple 2-parameters adaptive learning models can be used to approximate the observed adaptation processes in all 21 experimental conditions.

* Email correspondence should be sent to sonsino@ie.technion.ac.il. We thank David Budescu, Werner Guth, Dan Levine, Dov Monderer, Rosemarie Nagel, Alvin Roth and Ylva Søvik for comments and suggestions. We also thank the fund for the promotion of research at the Technion for its financial support.

1. Introduction

The experimental literature on *backward induction* and *iterated removal of weakly dominated strategies* is split in its findings concerning the descriptive power of these procedures. Whereas some studies (e.g., Erev & Rapoport, 1990) show that initially most subjects tend to perform at least two levels of backward induction, other studies (e.g., Stahl and Wilson, 1995) exhibit a lower initial level of sophistication. Moreover, while some studies (e.g., Nagel, 1995) demonstrate that the subjects move quickly towards the unique backward induction equilibrium, other studies (e.g., Ball, Bazerman & Carroll, 1991; Cooper et al, 2000) reveal no such convergence.

The current paper deals with a comprehensive experimental investigation of the Sebenius and Geanakoplos (1983) no-betting conjecture. In the Sebenius and Geanakoplos (henceforth: SG) model, agents with asymmetric information are asked to decide simultaneously whether they would like to engage in zero-sum betting. The bet is realized if and only if both individuals voluntarily agree to participate. The actual outcome is determined by the realized state of nature. Sebenius and Geanakoplos (1983) however show that if the agents follow some backward induction type of reasoning then at least one of them must reject the proposed bet in every state of nature and betting will never be realized. In this sense, rationality implies the "impossibility of zero-sum betting" in the Sebenius and Geanakoplos framework.

The SG no-betting result can be easily extended to more general settings. In particular, it is easy to show that the SG reasoning implies the impossibility of speculative trades in efficient economic markets (Milgrom and Stokey, 1982). The no betting conjecture thus seems quite important as it suggests that in economies with rational agents only real, value-adding trades can actually take place.

The main goal of the current paper is to study the potential descriptive implications of the SG conjecture and improve our understanding of the conditions (if any) under which this important conjecture is likely to describe human behavior.

In our experiments, we have let fixed pairs of subjects play repeatedly the two betting games (S and L) sketched in Figure 1 for 250 rounds. We have started with a basic experiment on 48 pairs of subjects. Half of these pairs got precise information about the underlying game. The other half played the game with almost no information.

<Insert Figure 1>

The results for both conditions revealed a persistently strong inclination to bet. The average betting rate at the first 50 rounds of the experiment was 37.5%. The corresponding individual entry rate (at the first 50 rounds) was almost 70%.¹ Convergence to the no-betting equilibrium was extremely slow. The average betting rate at the first 10 rounds of the repeated play of Game S (the simpler game) was 42%; the average betting rate at the last 10 rounds of that repeated game was 31%. The corresponding figures for Game L reveal a decrease in betting proportion from 47% to 39%.

Put differently, the results suggest that our subjects consistently exhibit a low level of sophistication (Stahl, 1993). In particular, the proportion of cases where subjects' behavior was compatible with two (or more) levels of "iterated removal of weakly dominated strategies" has mildly increased from 46.5% at the first 50 rounds of the repeated game to approximately 63% at the last 50 rounds of the game.²

In an attempt to trace the reasons for our subjects' poor performance in the basic experiment, we have manipulated the experimental design in 5 different directions. Some of these manipulations have indeed proved successful in terms of significantly decreasing the subjects' betting rate along the experiment. Yet, in all of these cases the individual subjects kept accepting the bet repeatedly in 45%-50% of the trials. Our subject's low level of sophistication and persistent inclination to bet thus seem quite robust with respect to the experimental manipulations.

A post-hoc quantitative analysis of the experimental data set suggests that simple adaptive learning models might be used to approximate subjects' behavior in all of the experimental designs. In particular, we demonstrate that the adjustment process can either be approximated by a simple reinforcement learning model or by a stochastic fictitious play model. We argue that (in both learning models) under the assumption that the learning speed decreases with payoff variability (Erev, Bereby-Meyer & Roth, 1999), it is

¹ All through the analysis we distinguish between "betting rates" which refer to those cases where both subjects accepted the bet and "entry rates" that address the individual players and refer to the proportion with which the individual players have accepted the bet.

² In particular, these numbers represent the proportion with which "Agent Two" subjects did not enter at the information set {A,B}.

possible to describe the adaptation process in all 6 experiments by a single set of two parameters; i.e., one does not gain much extra fit by approximating a different set of parameters for each experiment. In this sense, the slow learning process seems to be stable across our different experimental manipulations.

The paper proceeds as follows: Section 2 outlines the Sebenius and Geanakoplos framework and the no-betting result. Section 3 describes the basic experiment. Section 4 deals with our five alternative designs. Section 5 presents the quantitative summary of the results. Section 6 concludes.

2. The No-Betting Model

To illustrate the Sebenius and Geanakoplos model we use game S from Figure 1. We now describe the game verbally and sketch the no-betting argument: Assume that two agents, One and Two, are offered a zero-sum bet on a space of states $\Omega = \{A, B, C, D\}$ as follows. If the realized state is A, agent Two pays 32 dollars to agent One; if the realized state is B, One pays 28 dollars to Two; if the realized state is C, Two pays One 20 dollars, and if the true state is D, One pays Two 16 dollars. The agents hold asymmetric information on Ω . Let $\Pi_1 = \{ \{A, B\}, \{C, D\} \}$ denote the information partition of agent One and $\Pi_2 = \{ \{A\}, \{B, C\}, \{D\} \}$ denote the information partition of agent Two. Thus, when the realized state of the world is A or B, agent One only learns that the true state belongs to the subset $\{A, B\}$. Similarly, when the realized state is C or D, One only learns that the true state belongs to $\{C, D\}$. Agent Two, on the other hand, learns the exact state of the world whenever the realized state is A or D, but cannot distinguish between states B and C. The bet is outlined schematically as Game S in Figure 1. Assume also that both agents initially assign equal probability $\frac{1}{4}$ to each possible state and that betting occurs only when both agents accept the proposed bet. If one of the agents rejects the bet, it is not realized and each agent gets a fixed reservation payoff of 1 dollar.³

Note that given his prior uniform probability assessment agent One expects a gain from betting in every state of nature; i.e., if the realized state is A or B, One's conditional expected payoff from betting is $0.5 (+32 - 28) = 2 > 1$ (the reservation payoff); if the

³ SG (1983) assume that the reservation payoff is zero. We choose to modify this framework slightly and assume a fixed positive reservation payoff of 1 dollar. As demonstrated above, this modification does not effect the original no-betting result. Yet, we believe it makes our experimental results more interesting since we show that the subjects keep on betting although they have the option of getting a positive reservation value if they reject the bet.

realized state is C or D, One's expected payoff from betting is $0.5 (+20-16) = 2 > 1$. Similar calculations show that agent Two expects a gain from betting whenever the realized state is B, C or D. When the realized state is A, Two knows the actual realized state and thus expects a loss of 32 dollars.

Assume now that the actual realized state is D and that both agents initially decide to accept the bet at that state. Consider, however, the case where agent One mentally re-evaluates the proposed bet in four steps as follows:

Step 1: If Two is rational, she should clearly reject the bet whenever her realized information signal is $\{A\}$. (This follows from the assumption that Two learns the actual state when the realized state is A and thus expects a certain loss from betting at A).

Step 2: If indeed Two rejects the bet whenever the realized state of nature is A, I (One) can only lose (28 dollars) from betting when my realized information signal is $\{A,B\}$. Anticipating Step 1, I should refuse the bet when my realized information is $\{A,B\}$; i.e., when the realized state is A or B.

Step 3: If indeed I (One) reject the bet when the realized state of nature is B, Two can only lose (20 dollar) from betting when her realized information signal is $\{B,C\}$. Thus, if Two is rational and anticipates Step 2, she should refuse the bet when her realized information is $\{B,C\}$; i.e., when the realized state is B or C.

Step 4: If indeed Two rejects the bet when the realized state of nature is C, I (One) can only lose (16 dollars) from betting when my realized information signal is $\{C,D\}$. Anticipating Step 3, I should refuse the bet when my realized information is $\{C,D\}$; i.e., when the realized state is C or D.

Thus, if agent One indeed applies the above reasoning before making his final betting decision, he will refuse the proposed bet at the outset and betting will not take place.⁴

In general, let Ω denote a finite space of states, P_0 denote the agents' common prior on Ω , Π_1 and Π_2 denote the agents' information partitions, and $x : \Omega \rightarrow \mathfrak{R}$ denote a zero-

⁴ Alternatively, we may assume an explicit prepay communication process where the agents repeatedly announce whether they accept the bet or reject it. Steps 1-4 imply that both agents will accept the bet at the first 3 rounds but agent 1 will reject it at the 4-th round.

sum bet on Ω (where $x(\omega)$ denotes agent Two's payoff to agent One when the realized state is $\omega \in \Omega$, and a negative number represents a payoff of agent One to agent Two at the corresponding state). Let $\omega' \in \Omega$ denote the realized stated of nature and assume that both agents expect a gain from betting given their private information at ω' . SG (1983) prove that if the agents are risk averse then one of them must reject the bet within $\text{Card}(\Pi_1) + \text{Card}(\Pi_2) - 1$ successive re-evaluation steps, as demonstrated above (where $\text{Card}(\Pi_i)$ denotes the number of elements at the information partition Π_i).⁵ In this sense, zero-sum bets will never be realized if the agents have the opportunity to thoroughly re-evaluate their principle willingness to take the bet before it actually takes place.⁶ Moreover, it turns out that no-betting is the only possible equilibrium outcome when we analyze the SG model as a static game of incomplete information; this equilibrium is obtained by iterated removal of weakly dominated strategies. Furthermore, no-betting is the only possible equilibrium result even when we consider an extended game where the agents arrive at the zero-sum bet through a multi-stage negotiation process (Perets and Sonsino, 1999).

3. The Basic Experiment

3.1 Procedure

Forty-eight pairs of Technion students played repeatedly each of the two games outlined in Figure 1 (games S for Short and L for Long). Each of the games was played repeatedly for 250 rounds. Participants arrived to the laboratory in pairs and were seated side by side facing a personal computer screen. A partition was placed between them preventing one player from seeing the other player's screen-display or action.

Two experimental factors were manipulated to test the possible effect of information on subjects' inclination to bet: Half of the subjects (24 pairs) received a detailed complete description of the underlying betting game. To make the story more interesting, however, we have framed it as a "stock-market" game.⁷ We henceforth address these individuals as the "informed subjects". The other 24 "uninformed" pairs were also told that they would

⁵ Note that risk aversion is not required to obtain the no-betting result in game S. This follows from the fact that in this game each stage in the "iterated removal of information elements" process leads one of the agents to the conclusion that he can only lose from betting in a certain information set. Game L however is different and risk aversion is required in order to obtain the no-betting result for the states E,F,G and H in that game.

⁶ For an extension of this result to the case where there exists some small probability ϵ that the agents are irrational in the sense that they accept a bet when they expect a loss from betting, see Sonsino (1998).

⁷ The stock market cover story was removed in experiments 4.3-4.5.

participate in an experiment simulating a stock-market activity. These subjects were told that in every round of the experiment each one of them would receive an information signal describing the possible market conditions and would then have to decide whether to enter the market or stay out. If one of them decides to stay out, both of them would receive a payoff of one point. If both decide to enter, payoffs will be determined by the actual market conditions. Yet, these subjects were not informed of the payoff function of the game and the corresponding information structures. The (translated) instructions for the two types of subjects are enclosed in appendices A.1 and A.2. In the experiments, the instructions were provided in print and explained verbally by the instructors.

The second manipulation involved the order in which the two games (S and L) were played. Twelve of the pairs that have received full information about the underlying game started with a repeated play of the Short game and then played repeatedly the Long game. The other 12 pairs started with the Long game and then switched to the Short game. The 24 uniformed pairs were similarly subdivided to two groups of 12 pairs each. We henceforth address the first order as "order SL" and the second order as "order LS".

The experiment started with 3 fixed practice trials.⁸ These were followed by 250 repetitions of the first game, an intermission, and 250 repetitions of the second game. Each participant started the experiment with 1000 points. The accumulated score was reset to 1000 at the beginning of the repeated play of the second game. In every round of the game, the subjects gained or lost points according to their realized payoffs. In particular, if one agent (or both) have refused the proposed bet at some trial, then each of them got one point for that trial. In the break between the two games, the informed subjects were given additional instructions describing the exact structure of the second game. The uninformed subjects were only told that the second part of the experiment corresponds to a "different market".

Each player's screen-display included 3 windows. The first, presenting the current information signal.⁹ The second, presenting the payoff (in points) at the current round.

⁸ The realized state of nature was revealed to the subjects after each of these practice trials. This data was not provided in the experiment itself. The additional feedback was provided in the practice trials in order to help the informed subjects fully understand the payoff function of the game.

⁹ The subjects in one half of the pairs observed the information signal for 2 seconds and then had to announce their decision. The subjects in the other half could observe the information signal for as long as they wish before making their decision. The differences in results across these two conditions were marginal and statistically insignificant so we have pooled the data together.

The third, presenting the accumulated payoff (in points). In every round of the experiment, each of the participants first observed his realized information signal. He then had to choose one of two keys (the first representing a decision to bet and the other representing a decision not to bet). Payoffs were then realized according to the game's payoff rule and disclosed to the players. The accumulated score was updated accordingly.

Each subject received approximately \$6¹⁰ for participating in the experiment. In addition, each participant could gain \$6 more as a bonus contingent on his performance in the experiment. We have used a binary lottery (independent lottery for each participant) to determine whether the player will receive the additional \$6 at the end of the experiment. The probability of winning the lottery was an increasing function of the subject's accumulated score in the two blocks of the experiment. The lottery was implemented in order to eliminate the possible effects of risk attitude on subjects' behavior (in accordance with Roth and Maluf (1979) binary lottery procedure).

3.2 Results

3.2.1: Initial behavior

The average individual entry rates of the informed subjects at the first round in which each information set was played (for each order and each game) are presented in the upper 4 rows of Table 1. Recall that in two of the information sets, {A} and {D}, the rational behavior is obvious. The number of rounds (levels) required to eliminate each of the other information sets (according to the SG model) is listed at the second row of the Table.¹¹ The figures on the table reveal a surprisingly low initial level of reasoning. Indeed, even in cases where it takes only two levels of reasoning to reach the no-betting conclusion (at the information set {A,B}) the average entry rate was higher than 75%. The coefficient of correlation between the levels required to eliminate each information set according to the SG model (as given at the second row of Table 1) and the average initial entry rate in each treatment was found insignificant. We discuss this observation again in Section 5.

The right hand column of Table 1 presents the average initial entry rate of the uninformed subjects for each condition. Since there is no point in separating the numbers across

¹⁰ 20 Israeli Shekel

¹¹ For example, the information set {A,B} is eliminated at the second round of the SG reasoning process, so the "levels required" are 2.

information sets when subjects are uninformed, we present the average across all information sets. The average initial entry rate across all conditions (.75) is surprisingly high and significantly higher than then the 50% rate implied by random behavior ($t[47]=6.09$ $p < .0001$).

<Insert Table 1>

3.2.2: Very Slow Learning

On average, the experimental results reveal a very slow decrease in the inclination to accept the bet with experience. While the average betting rate at the first 10 rounds of the repeated play of the Short game was 42%, the average betting rate at the last 10 rounds was only 31%. (A significant difference at $p < 0.003$ ($t[47]=3.17$)). The corresponding figures for the Long game show a decrease in betting proportion from 47% to 39% ($t[47]=1.91$, $p < .07$).

The two columns at the left of Figure 2a present the average betting rates in blocks of 50 trials for each one of the 8 conditions.¹² The graphs reveal a very slow negative trend for most of the experimental manipulations. This linear trend (over all experimental manipulations) was significant ($F[1,95]=11.35$, $p < .002$).

<Insert Figure 2a >

In Figure 2b, we have separated the subjects' behavior in the Short game across information sets. The graphs demonstrate that (in game S) Player Two has learned to reject the bet when his realized information signal was {A}. On average, at the last 2 blocks of the experiment, subjects of type {A} have rejected the bet in 92.5% of the trials (uninformed subjects rejected in 92% of the trials, while informed subjects rejected in 93% of the trials). These "nice" results, however, don't carry over to the next information set on the SG chain of reasoning: The average rejection rate of player One at the information set {A,B} was 60.5% in the case of informed subjects and only 47.5% in the case of uninformed subjects. When we move on to the "next" information set, {B,C}, the results are even weaker (40% rejection rate in the case of informed subjects and 48%

¹² The horizontal axis (in figures 2a, 3a and 4a) counts the 50-trials blocks (from block 1 to block 5) in the experiment, the vertical axis measures the average betting rate at the corresponding block. The two columns at the left of the figure refer to the experimental results. The two columns at the right refer to the simulations that will be discussed in Section 5.2.

rejection rate in the case of uninformed subjects). The data for game L reveals similar trends: Nice learning at the information set {A}, which carries over only partially to the information set {A,B} and almost disappears at the information set {B,C}.

<Insert Figure 2b >

Interestingly, it turns out that subjects of type One lost, on average (across both games and all the manipulations) seven points in each of the trials in which they have entered when their realized information signal was {A,B}. The corresponding average loss of these subjects at the last block of the experiment was even higher (7.1 points). In spite of this fact, the subjects kept accepting the bet repeatedly at {A,B}. The average entry rate at {A,B} in the last block of the experiment was 37.5%.¹³

3.2.3: Weak Information Effect

Figure 2 demonstrates that information had a surprisingly weak effect on subjects' behavior. Informed players of type Two were less likely to enter at the information signal {A} and more likely to enter at the information signal {D}. Yet, this reasonable behavior did not translate into an overall significantly lower betting rate of the informed subjects. Even the largest difference in average betting rates between informed subjects and uninformed subjects (game S, block 1, order LS) was not statistically significant ($F[1,22]=4.1, p<0.6$).

3.2.4: Non significant Transfer

The order manipulation was implemented to check whether experience in one game could be carried over to a second different game. We also wanted to check the interaction of information with this possible “transfer” effect. Our results however don't provide clear support to the transfer hypothesis. The initial entry rates (see Table 1) of the informed subjects in game S were indeed lower at the information sets {A}, {A,B}, {C,D} when the game was played second (order LS). Yet, the order effect is negative at the information set {C,D} and it disappears when subjects are uninformed: the average initial entry rate of uninformed subjects in game S is significantly higher in the order LS (then in the order SL). The transfer effects are even weaker for game L.

¹³ Informed subjects of type One lost on average 8.9 points from entering at (A,B). Uninformed subjects lost on average 5.5 points in those cases.

3.2.5: Two Beats One in Game S

Table 2 demonstrates that Player Two was "the winner" in game S in both information conditions in our experiments.¹⁴ The higher gains of that player might be explained by the fact that this player has a "better" information structure than Player One in the two games. In particular, Player Two is able to pre-identify state A in which he might lose a relatively large number of points from betting. He can similarly pre-identify state D in which he has a sure gain. Since Player One, on the other hand, has a coarse information structure that doesn't include any singleton elements, that player plays the game (S) from an inferior position and therefore loses to player Two on average. The relative advantage of player Two, however, gets much weaker in game L. In that game, the information partition of Two contains only one singleton element {A}. Moreover, the probability of the corresponding state A is decreased to 1/8 (compared to 1/4 in game S).

<Insert Table 2>

The fact that Player Two makes nice profits along the experiment brings up an alternative explanation to the consistent inclination to enter at the information set {A}: Player Two might keep entering occasionally at {A} because this behavior pays off. By entering occasionally at {A}, Two confuses player One and breaks down the SG domino-effect. Indeed, if One would have learned to reject the bet at {A,B}, Two could not have ended the experiment with such nice gains. In Section 4.2, however, we will reject this "sophisticated" explanation.

3.3 Discussion

The results described above seem suspiciously discouraging. The most disturbing point is the very slow learning of the informed subjects when playing the Short game. We were especially bothered by the fact that in game S, informed subjects of type {A} have (on average) accepted the bet in more than 10% of the trials. This behavior seemed unreasonable given the simplicity of game S and the fact that these subjects were told at the outset that they could only lose 32 points from betting in this case. We were also surprised to see that player One type {A,B} accepted the Short bet (on average) in 51% of the trials in spite of the fact that he has lost (on average) more than 7 points in each of these entries. We have therefore decided to manipulate the experimental conditions in

¹⁴ The order manipulation didn't have a significant effect in this case as well.

several alternative directions in order to examine possible explanations to the observed behavior.

<Insert Figure 3 >

4. Alternative Experimental Designs

4.1: Stronger Feedback.

In experiment 4.1 we have enhanced the feedback that was provided to the subjects after each trial. In particular, we have disclosed to each subject the payoff that he would have obtained if he chose differently at that round. In game S, for example, when the realized state of nature at the first trial was A, and player One decided to enter while player Two decided not to enter, One got a four lines feedback as follows:

<p>Your response was Y As a result your payoff is 1 Your total score is 1001 If you had chosen N, your payoff would have been 1</p>
--

Similarly, Two got the feedback:

<p>Your response was N As a result your payoff is 1 Your total score is 1001 If you had chosen Y, Your payoff would have been -32</p>
--

Except for the additional feedback the design of this experiment was identical to the one described in Section 3. We had run 8 pairs of subjects in each of the four basic information/ order conditions for this case. Table 1 presents the average initial entry rates for this design. Figure 3a presents (graphically) the average betting rates in blocks of 50 trials for each case. Figure 3b presents the individual entry rates in each information set

for game S.¹⁵ Tables 3 and 4 present the average betting rates by blocks for the Short game and the Long game for this design. The major conclusions are outlined next.

<Insert Tables 3 and 4>

4.1.1. Decrease in Initial Entry Rates in Game S

The data on Table 1 demonstrates that the increased feedback has significantly reduced the initial entry rates of informed subjects at the information signal {A,B} in 3 of 4 cases. The average initial entry rate (across all information sets except for {D}) of informed subjects in Game S has decreased from 63.4% (in design 3) to 51.8% (at the current design). The corresponding proportions for uninformed subjects were 75% (in design 3.1) and 59% (at the current design). In the Long game, the average initial entry rates of informed subjects were not significantly different across the two designs.¹⁶ The average entry rates of uninformed subjects in the Long game has decreased from 74.5% to 61%.

4.1.2 Insignificant Effect on Learning

The figures in Tables 3 and 4 demonstrate that the initial impact of the stronger feedback was carried over to the subsequent rounds. While the average betting rate (across the two games, the 4 information/order manipulations and the 250 repetitions) in the original design was 36.8%, the corresponding figure for the stronger-feedback design (4.1) was only 24.5% (significant difference at $p < .0001$). The linear trend from block 1 to block 5, however, was not significantly different across the two designs ($F[1,158]=1.16$ n.s.).

4.1.3 Positive Information Effect

Information had a positive effect on learning in this design. While the decrease in individual entry rate (from block 1 to block 5) for the informed subjects was 23.3%, the corresponding figure for the uninformed subjects was 6.3%. A linear trend analysis shows that the decrease in betting rates (across the five blocks) was significantly larger for the informed subjects ($F[1,62]=4.68$, $p < .04$). The graphs in Figure 3a demonstrate that the effect was most pronounced in the Long game.

¹⁵ In figures 3b, 4a and 4b we have ignored the insignificant order effect and presented one "curve" for both orders (SL and LS) in each condition.

¹⁶ The initial entry rates (excluding the information set {D}) were 67.25% in design 3 and 67.43% in design 4.1.

<Insert Figure 3>

4.2: Change in payoffs

In experiment 4.2 we have changed the payoff matrix of Game S as follows: The payoff to player Two at the state A was changed from -32 to -128. The payoff to player Two at the state D was decreased from +16 to +4. The main purpose of the change was to provide stronger incentives to player Two to reject the bet at A. We expected that the stronger incentives will speed up learning and significantly decrease the betting rates. We have also decreased player Two's payoffs at D in order to make it more difficult for Two to "beat" One in the repeated experiment. Thus, this design also tests the "sophisticated" explanation discussed at the end of section 3.2.5. If we make it much more expensive for player Two to enter at A (and significantly decrease the compensation he might get at D), then intentional entry at {A} would not payoff as much as it did in the previous setup.

We have run this design only on game S and provided the subjects the same detailed feedback that was used in design 4.1. Table 3 presents the average betting rates of the informed subjects and the uninformed subjects for this case. The results are quite similar to the ones obtained in design 4.1. Statistical analysis reveals that the change in payoffs did not have a significant effect on the average betting rates ($F[1,30]=.79$, n.s.) and the linear trend ($F[1,30]=.83$, n.s.)

Figure 3b, however, demonstrates that the change in payoffs (from design 4.1 to design 4.2) has caused a sharp decrease in the inclination of informed subjects to accept the bet at the information set {A}. The average entry rate at {A} has decreased from 16.6% to 5.8%. The entry rate at {A,B} has also significantly decreased: From 51% in design 4.1 to 17.6% in the current design. The entry rates at {B,C} however were not affected by the change in payoffs (59.8% in design 4.1 and 59.6% in design 4.2).

The effect of the change in payoffs was much weaker in the case of uninformed subjects. The average entry rate of subjects of type {A} has decreased from 16.6% to 11.6%. This however was not strong enough to significantly decrease the acceptance rates at {A,B}.¹⁷

¹⁷ The average entry rate of these subjects in design 4.2 was 41.6%. The corresponding figure for design 4.1 was 42.2%.

Finally note that informed subjects of type Two have gained (on average) 23 points in this design compared to an average gain of 540 points in design 4.1. Recall that the subjects could make a sure gain of 250 points by simply rejecting the bet in each round. This suggests that some of the (significantly larger) entry rates at {A} observed in the previous designs might have been caused by strategic considerations (see Section 3.2.5): Type Two players entered occasionally at {A} since it encouraged subjects of type One to enter at {A,B} and “paid off” on average. The change in payoffs in the current design ruins the profitability of this strategy and significantly cuts the entry rates at the information set {A}. The fact that the subjects kept entering at {A} in more than 5% of the cases even at the current design might be explained by (I) Other reasons for entering at {A} (e.g., plain noise or mistakes) or (II) Claiming that the sophisticated considerations outlined above are partially robust to the failure of that strategy in the current design.

4.3: Change in Cover Story

In experiment 4.3 we have modified the instruction set and the feedback provided to the subjects in attempt to clarify and simplify the experiment. In particular, we have omitted the stock-market cover story and used the original zero-sum bet framing of Sebenius and Geanakoplos. In addition, we have changed the screen design and the feedback system. At the beginning of each round, when the subject received his information signal, he was reminded of his two options through the following screen:

DECIDE:	
Y	Accept the game (see instructions for payoff)
N	Reject the game and your Payoff will be 1 point

In addition, after each round of the experiment, each player observed a feedback window disclosing all the information on the realization of the game at that round. For instance, when the realized state at the first round of the experiment was A and both players accepts the bet, player Two got the feedback:¹⁸

¹⁸ Note however that in this design we have not provided the enhanced feedback scheme used in designs 4.1 and 4.2.

<p>Your response: Y</p> <p>The other player's response: Y</p> <p>The realized state: A</p> <p>Your payoff for this round: +32</p> <p>Your total accumulated score: 2032</p>
--

The subjects could examine the feedback window for as long as they chose. They were instructed to press a certain key to remove the feedback and continue to the next round. The (translated) instructions for his design are enclosed in Appendix A.3. We have run it on 12 pairs of informed subjects. Table 3 shows that the betting rates per block for this design are very similar to (and statistically not different from) the betting rates in the original design (3). Figure 4b reveals that the modified instructions have helped subjects understand that they should not enter in {A}: The average entry rate at {A} has decreased from 19% in the design 3 to 6% in the current design. Similarly, the subjects have learned that they should enter in {D}: The average entry rate at {D} has increased from 86.4% to 94%. The average betting rates in all other information sets were not significantly different across the two designs.

<Insert Figure 4 >

4.4: Interesting Outside Option

To check whether subjects keep entering because "rejecting is boring" we have replaced the fixed 1-point reservation payoff that was paid to each subject when the bet was rejected by a state-dependent lottery with an expected payoff of 1 point. We have run this design on the Short game with 12 pairs of informed subjects. The instruction set was identical to the simplified version used in design 4.3 except for replacing the sentence "*if one of the players rejects the bet, the bet will not take place and each player will get a payoff of +1 point*" with "*If one of the players rejects the bet, the bet will not take place and the payoff will be randomly determined. The distribution of payoffs for this case will appear on your screen.*"

The payoff distribution of the "reservation lottery" was fixed (across trials) for each player and for every possible information signal of that player (but varied across players

and across information signals).¹⁹ Whenever the realized state of nature was D, for example, player One observed the screen:

DECIDE:	
Y	Accept the game (see instructions for payoff)
N	Reject the game and your payoff will be: +20 with prob. 1/3 -16 with prob. 1/3 -1 with probability 1/3

Similarly, player Two observed:

DECIDE:	
Y	Accept the game (see instructions for payoff)
N	Reject the game and your payoff will be: +16 with prob. ½ -14 with prob. ½

The betting rates per block for this design are disclosed in Table 3. The average betting rate across all five blocks was 24%. This rate is 8% lower than the one for design 4.3, but 3% higher than the corresponding rates for designs 4.1 and 4.2. An examination of subjects' behavior at the different information sets (see Figure 4b) gives the impression that the alternative lottery option has confused the subjects and took them even farther from complying with the rational no-betting conjecture. The average entry rate of subjects of type {A} has increased to 24% in this design. The corresponding rate for

¹⁹ In general, the payoff distribution of the lottery was determined as follows: It assigned equal probabilities to each of the payoffs that the player could get given his realized information signal (so that player One could get +20 or -16 with equal probabilities when his realized information was {C,D}) and it assigned the same probability to the payoff that made the expected payoff from the lottery equal to +1 (that is, -1 in the example above).

design 4.3 was 6%. Similarly, the average entry rate at the information set $\{D\}$ has decreased from 94% (in design 4.3) to 78% (in the current design).

4.5: Changing the Incentive Scheme

In this design we have manipulated the payoff scheme in an attempt to simplify the task. In particular, subjects were told that their final actual payoff (in Shekels) would be determined by dividing their accumulated score at the end of the experiment by 100. The subjects have started the experiment with 3000 points. Once again, we have restricted this manipulation to the Short game and to the informed subjects. We have used the simplified instructions introduced in Section 4.3 (modified to the change in payoff scheme).

The average betting rates per block for this design are presented again in Table 3. The results seem close to the ones obtained in design 4.3 (the one with the simplified instructions). The betting rates are somewhat lower but the differences are not statistically significant. The average betting rate at the last block of the experiment was still 23%.²⁰

5. Summary: A Two-Model Framework

To summarize the results quantitatively, it is constructive to distinguish between the initial tendencies to accept the bet and the observed characteristics of the learning process. In this section, we present a post-hoc two-model framework summarizing the experimental data. The specific models we examine are only examples of the large set of models that might be used to summarize our data. Yet, we believe that the analysis provides interesting insight into our subjects' behavior across the experiments.

Before we proceed with the details we need to define the concept of a "strategy" in the Sebenius and Geanakoplos betting game. For this sake, we look at the SG game as a static game of incomplete information. We define the (normal-form) pure strategies of the game as follows: A pure strategy for player i specifies the behavior of that player (accepting or rejecting the bet) for every possible information signal that this player may receive. In game S, for example, a strategy for player One will specify that player's

²⁰ Note however that this incentive scheme ignores the possible effects of individual risk attitudes on the observed behavior.

behavior (accept or reject) when his observed information-signal is $\{A,B\}$ and his behavior when his realized information-type is $\{C,D\}$.

5.1 Modeling the Initial Behavior.

To characterize our subjects' initial inclination to accept or reject the bet (as described on Table 1) we use the concept of level- n strategies as advanced by Stahl (1993, 1996) and Stahl and Wilson (1995). In particular, we assume that subjects choose their behavior from a finite set of Level- n strategies as follows: Level-0 is the strategy in which the agent accepts the bet in every information set.²¹ Level-1 is the strategy in which the agent plays his best response to the assumption that his opponent plays Level-0, and - in general - Level- n (for $n > 1$) is the strategy in which the agent plays his best response to the assumption that his opponent plays the strategy Level- $(n-1)$.

For instance, in game S, if the realized information is $\{A,B\}$ strategy Level-1 implies that Player One will play best replay to the belief that the expected payoff from betting is $0.5(+32) + 0.5(-28) = 2$. Since this is bigger than the reservation payoff 1, One's best response is to accept the bet at that information set. Similarly, it is easy to show that One would also prefer to accept the bet at the information set $\{C,D\}$ when he expects his opponent to enter. Thus, Level-1 for player One is the pure strategy (Bet, Bet). Table 5 presents the strategies Level-1 to Level-7 for each player, in the games S and L.

< Insert Table 5 >

Note that in the Short game all the strategies of level $n \geq 4$ prescribe the SG no-betting behavior (in the sense that each player rejects the bet in those states of nature in which he should have learned to reject the bet according to the SG algorithm). In the Long game, however, we need 7 iterations to reach the no betting result: Strategy Level-7 is the only one in the list that prescribes the SG behavior.

We have added to the list one more strategy, Level-8, that is formally a (weak) best response to Level-7.²² Level-8 is interesting for its own sake since it prescribes betting

²¹ This definition is different from the one used by Stahl and Wilson (1995) that assume random behavior in Level-0. Note however that using Stahl and Wilson's definition for Level-0 does not change the structure of the strategies Level-1-Level 8 as presented on Table 5.

²² That is, the expected payoff from playing Level-7 is the same as the expected payoff from playing Level-8 when the opponent follows Level-7.

only in those information sets where the agent expects a gain from betting independently of his opponent behavior (i.e., it prescribes betting only at the information set D of player II in game S).

To study the effect of our experimental manipulations on the subjects' initial inclination to bet, we study a logistic level-n model, where the probability of accepting the bet at the information set π , $P(\text{accept}|\pi)$, is given by the expression:

$$\text{Ln} [P(\text{accept}|\pi) / (1 - P(\text{accept}|\pi))] = \beta_0 + \beta_1(L_1[\pi]) + \beta_2(L_2[\pi]) + \dots + \beta_8(L_8[\pi]).^{23}$$

In this expression, β_0 - β_8 denote the estimated parameters of the model while $L_n[\pi]$ denote the probability with which the player accepts the bet when his realized information signal is π and he follows the strategy Level-n. In particular, β_0 represents the "independent" initial tendency to bet, while β_1 - β_8 represent the impact of the different Level-n strategies on that tendency.

We have estimated the model's parameters by running logistic regressions in different levels of data-aggregation. First, we have run the model for each one of the 21 experimental conditions separately²⁴ and observed that: (a) In all treatments with informed subjects the only significant coefficients were β_0 , β_1 and β_8 (the coefficients β_2 - β_7 were not significant in all these treatments). (b) In all treatments with uninformed subjects the only significant coefficient was β_0 (the coefficients β_1 - β_8 were not significant in all these treatments).²⁵

²³ When estimating the model for treatments that only involved the Short game we have restricted the equation to the form $\text{Ln}[P(\text{accept}|\pi)/(1 - P(\text{accept}|\pi))]=\beta_0+\beta_1(L_1[\pi])+\beta_2(L_2[\pi])+\beta_3(L_3[\pi])+\beta_8(L_8[\pi])$ (since $L_n[\pi] = L_4[\pi]$ for $n>4$ in this game).

²⁴ See Table 1 for a concise presentation of the 21 experimental conditions (12 with informed subjects and 9 with uninformed subjects; 13 involving the Short game and 8 involving the Long game). Note that since the short game has 5 information sets while the Long game has 8 information sets these conditions translate into a total of $13*5+8*8=129$ curves.

²⁵ The initial betting rates measured in Table 1 ignore the fact that the experience gained while responding to previous information signals, might have affected subjects' behavior in their first encounter with a given information signal. To check the robustness of our results we have also run the regressions for the case where the initial betting rates were measured only from the data for the first round of the repeated game. The results were not significantly different from the originals.

Following observation (a) we adopted the restricted model

$$(5.1.1) \quad \text{Ln} [P(\text{accept}|\pi) / (1 - P(\text{accept}|\pi))] = \beta_0 + \beta_1(L_1[\pi]) + \beta_8(L_8[\pi]),$$

to represent the initial inclination to accept the bet in the 12 treatments with informed subjects. To shorten the exposition and avoid unnecessary details we choose to present the results of this estimation for all 12 treatments together.²⁶ The estimated coefficients and significance levels are: $\beta_0 = -1.66$ ($p < 0.0001$), $\beta_1 = 2.456$ ($p < 0.001$) and $\beta_8 = 0.539$ ($p < 0.05$). These results clearly reflect the low initial level of sophistication of our subjects. The significance of β_8 seems to follow from the high entry rate at the information set {D} in game S (Level-8 is the only strategy that prescribes entry at {D} alone).

Following observation (b) above, we simply let the average entry rate represent the initial inclination to accept the bet in the treatments with uninformed subjects; i.e., we adopt the model

$$(5.1.2) \quad P(\text{accept}|\pi) = \text{Average initial entry rate at } \pi,$$

to represent the initial inclination to bet in the treatments with uninformed subjects.

Since the average entry rate at the 9 treatments with uninformed subjects was 0.68 (see Table 1), we simply let $p=0.68$ represent the initial inclination to accept the bet in the pool of uninformative treatments.

5.2 Modeling the Learning Process

5.2.1: Learning by Reinforcement

The experimental literature on learning in games widely agrees now that human subjects' behavior in repeated interactions might be approximated by quite simple adaptive learning models.²⁷ Erev and Roth (1998), for instance, show that a simple two-parameters learning model can be used to approximate subjects' behavior in 12 different

²⁶ We had $8 \cdot 5 + 4 \cdot 8 = 72$ observations for this estimation.

²⁷ For specific examples, see Roth and Erev (1995), Erev and Roth (1998), Daniel et al. (1998), Fudenberg and Levine (1998), Chung and Friedman (1999), Tang (2001), Camerer and Ho (1999), Mookherjee and Sopher (1997), Sarin and Vahid (1999), Rapoport and Amaldoss (2000).

(independent) experimental data sets. In this section we show that a similar two-parameter model (based on the models used by Erev et al. (1999) and Roth et al.(2000)) can be used to approximate the subjects' behavior in the current task. The main assumptions of our model (henceforth addresses as RL for Reinforcement Learning) are described next:

A1: Initial Propensities

At time $t=1$ (before any experience has been acquired) each player i has an initial propensity to play his k -th pure strategy that is described by a real number $q_{ik}(1)$. Informally, the initial propensity to play strategy k can be thought of as a subjective initial estimate of the expected payoff from playing that strategy.

A2: Updating of Propensities

If player i plays his j -th pure strategy at time t and receives a payoff of x , then his propensity to play strategy j is updated with respect to x according to the following formula:

$$(A2) \quad q_{ij}(t+1) = [q_{ij}(t)[C_{ij}(t)-1+N(1)/m_i] + x] / [C_{ij}(t)+N(1)/m_i],$$

where $C_{ij}(t)$ is the number of times that strategy j has been played by i at the first t rounds of the game, $N(1)$ is a parameter (the first "free" parameter of the model) that determines the strength of the initial propensities, and m_i is the number of pure strategies for player i . Thus, the propensity to play strategy j is increased (when $x > q_{ij}(t)$) or decreased (when $x < q_{ij}(t)$) according to the payoff (x) that it has actually generated at the last round of the game.²⁸

The propensity to play all other strategies ($k \neq j$) that were not actually played at round t , is not changed at that round.

A3: Exponential Derivation of Probabilities

The probability $p_{ik}(t)$ that player i plays his k -th pure strategy at time t is given by:

$$(A3) \quad p_{ik}(t) = \text{EXP}[\lambda q_{ik}(t)/S_i(t)] / \sum \text{EXP}[\lambda q_{ij}(t)/S_i(t)],$$

²⁸ Initial propensities get together a total weight of $N(1)$, so that the initial propensity assigned to each strategy j of player i is $N(1)/m_i$. The division by m_i is required in order to approximate the experimental results for the Short game and the Long game with the same general model (since the number of pure strategies for each player is quite different across the two games).

where the sum is over all of player i 's pure strategies, λ is a "free" parameter of the model that represents the sensitivity of the choice-probabilities with respect to the history of the game, and $S_i(t)$ is a measure of payoff variability that is repeatedly updated as follows: $S_i(1)$ equals the expected absolute value of the difference between the payoff obtained from random choice and the average payoff from random choice²⁹ and for every $t \geq 1$, $S_i(t+1)$ is calculated from $S_i(t)$ by the following formula:

$$S_i(t+1) = [S_i(t)(t-1 + N(1)) + |x - A_i(t)|] / [t + N(1)],$$

where $A_i(1)$ denotes the expected payoff from random choice and for every $t \geq 1$,

$$A_i(t+1) = [A_i(t)(t-1 + N(1)) + x] / [t + N(1)].$$

We derive the choice probabilities using the exponential formula (A3) to avoid the unnecessary complications that may arise when the propensities become negative.

Existing experimental data suggests that the speed of learning decreases with payoff variability (see, for example, Erev et al.(1999)). The normalization of the propensities for time t with respect to $S_i(t)$ is carried out in order to reflect this fact.³⁰

Our model thus has two "free" parameters: λ and $N(1)$. For our specific application, we assume that the strategy space upon which the model is applied is the space of normal-form strategies in the SG betting game. We use the two initial propensities models (5.1.1 for the pool of informed subjects and 5.1.2 for the pool of uninformed subjects) described in Section 5.1 to estimate the initial inclination to accept the bet in each information set and derive a corresponding initial probability distribution over the space of (pure) normal form strategies of the game in the obvious way.

5.2.2: Method and Main Results

A grid search was conducted to find the parameters (λ and $N(1)$) that minimize the Mean Squared Deviation (MSD)³¹ between the observed average betting rates (in blocks of 50

²⁹ I.e., if X denotes the random payoff to each player when both players enter with probability 0.5 in each information set, independently across information sets, then $PV(1)$ is equal to $E(|X - E(X)|)$ where E is the expectation operator.

³⁰ This "normalization" seems especially relevant to the comparison between designs 4. and 4.2.

trials) and the learning model's predictions. The model's predictions were derived by running 500 computer simulations for each experimental condition.³² In each simulation, a pair of virtual agents played a fixed SG betting-game repeatedly. In each round t of the simulation, the following steps were implemented:

- (1) Players' strategies were randomly selected according to Assumption A3.
- (2) The state of nature was randomly selected, and each player's realized information-signal and implied act ("bet" or "don't bet") were determined accordingly.
- (3) Individual payoffs were calculated according to the game's payoff function.
- (4) Individual propensities were updated according to Assumption A2.
- (5) The average payoff measure ($A(t)$) and payoff variability measure ($S(t)$) were updated accordingly.

The best fit in our grid search was obtained for the case where $N(1) = 100$ and $\lambda = 6$. The MSD over the 129 curves (see footnote 29) was 0.018.³³ The simulation results are graphed vis a vis the experimental results in Figures 2-4.

To check the robustness of the learning process with respect to our experimental manipulations, the model's parameters were also estimated separately for each of our 21 tasks. This procedure that increases the number of free parameters from 2 to 42 has decreased the average MSD score from 0.018 to 0.013.

As an additional sensitivity check we have also estimated the initial inclination to bet for each one of the 21 treatments separately using equation 5.1.1 for the treatments with informed subjects and equation 5.1.2 for treatments with uninformed subject and used the estimated equations to generate the initial propensities for the simulation of each treatment. The average MSD for this treatment however was very close to the one obtained in the general estimation: MSD=0.018 in both cases.

³¹ Formally, the MSD score measures the average squared distance between the prediction graph and the experimental data. For a detailed discussion of MSD rules see Selten (1995).

³² Note however that our learning model (with the initial propensities derived from equations 5.1.1 and 5.1.2) does not distinguish between experiments 4.1, 4.3 and 4.5 and the original Experiment 3, the simulation results presented for conditions 4.1, 4.3 and 4.5 are thus identical to those presented for treatment 3.

Finally, we have also run the simulation starting with a uniform initial betting rate of 0.5 in each information sets. The average MSD score of this model was 0.0246. The results of all 4 simulations are presented on Table 6.

To verify that the normalization for payoff variability (as measured by the parameter $S_i(t)$) indeed contributes to the fit of the model, we have run the simulations again restricting $S_i(t)=1$ for each t and i . The MSD of the restricted model was 0.027 (see Table 6 for comparisons with other simulations).

A careful examination of the simulation-results (see, for example, the summarizing graphs in figures 2-4) reveals that our learning model captures the following experimental results:

- (1) The very slow learning in game L.
- (2) The somewhat faster learning in game S.
- (3) The difference in subjects' behavior across information sets.
- (4) The fact that the experimental manipulations outlined in Section 4 did not have a significant effect on the adaptation process (see Figures 2-4).

5.2.3: Stochastic Fictitious Play

Stochastic Fictitious Play models (henceforth: SFP) (see, for examples, Fudenberg and Levine, 1998, Goeree and Holt, 1999) typically assume that the players use all the available information to compute the expected payoff of each strategy. In particular, the SFP models (implicitly) assume that the information concerning the forgone payoff from unchosen strategies is as important as the payoffs actually obtained. Formally, (see Camerer and Ho, 1999) it is possible to show that the same three equations that were used to describe the Learning by Reinforcement model in Section 5.2.1 can be used to present SFP learning when assumption A2 is replaced by the following A2':

A2': Updating of Propensities with known forgone payoffs

If the payoff (that was obtained or could have been obtained) from strategy j at trial t is x_j then the propensity to play strategy j is updated with respect to x_j according to the following formula:

$$(A2') \quad q_{ij}(t+1) = [q_{ij}(t)[t-1+N(1)/m_j] + x_j] / [t+N(1)/m_j].$$

To evaluate the descriptive power of the SFP learning model in the current experiment, we have used again the grid search method + computer simulations outlined in Section 5.2.2 to estimate the parameters $N(1)$ and λ . In particular, we have estimated two versions of the model. At the first, we have normalized for payoff variability exactly as explained in Section 5.2.1.³⁴ At the second, we have constrained the payoff variability measure $S(t)$ to be 1 for all t .

The results of the two estimations are summarized in the last two rows of Table 6. The figures on the Table show that with normalization the fit of the SFP model ($MSD=0.021$) is close to the fit of the RL model ($MSD=0.018$). A closer look at the results suggests that the main reason for the slightly lower fit is that the SFP model predicts a higher learning speed than the RL model. This prediction is inconsistent with the results for some of the treatments.³⁵

As in the case of the RL model the fit is impaired (the MSD increases to 0.031) when the normalization is removed. A closer examination of the data suggests that the normalization is particularly important to capture the small effect of the payoff manipulation described in experiment 4.2. Without this normalization both models (RL and SFP) predict a large difference between the predictions for experiment 4.2 and the experimental results for this game.

5.2.4 Implications

We would like to emphasize again that the current attempt to model the adjustment process was not intended to find the “best” approximation of the observed behavior. It is possible for instance that a generalized model that accommodates the RL model and the SFP model, like the model proposed by Camerer and Ho (1999), would provide closer predictions of the observed data.³⁶ Nevertheless, the results above clearly demonstrate our main point: Simple (two-parameters) adaptive learning models can be used to approximate subjects behavior in all 21 experimental tasks (129 curves) when the learning speed is normalized with respect to payoff variability.

³⁴ Typically, (see for example Fudenberg and Levine 1998, Goeree and Holt, 1999) SFP models do not normalize for payoff variability.

³⁵ The full details on the SFP estimation can be obtained from the authors upon request.

³⁶ Note however that Camerer and Ho use at least 5 parameters to model each Player's role in each task. This translates into 210 parameters in the current study.

Our results also suggest that the adjustment of the learning speed with respect to payoff variability is very important in the current context. This in turn might imply that the main reason for the slow learning observed in our experiment, is the relatively large payoff variance that characterizes our betting games. Under this interpretation the observed slow learning does not follow from the difficulty of backward induction reasoning. Indeed, note that the speed of learning was not significantly affected by the information conditions (while backward induction reasoning can only be applied under full information). Note also that this interpretation suggests that the no-betting result might still be relevant in betting environments with relatively low payoffs. In particular, one may imagine an Internet betting site where betting occurs repeatedly while participants update their (normal-form) betting strategy once in a while after observing the average payoff obtained from their previous strategy. If the frequency of betting is high then the observed payoff-variability will be relatively low. Our results suggest that the no-betting result might apply in this application.

6. Discussion

The rapidly growing literature in experimental game theory demonstrates very clearly that human agents are not always rational but seem to adjust their behavior to their accumulated experience in most circumstances. These observations suggest that theoretical results that assume "common knowledge of rationality" or "equilibrium behavior" might have descriptive implications if adaptive agents learn from experience to follow the corresponding type of behaviors.

The SG model seems like an especially interesting setting for testing human agents' ability to learn from experience to behave in the "rational way". In particular, note that since the SG game is a zero-sum game (with a fixed positive reservation price) the strategic type of explanations that may be used to explain subjects' deviations from unique equilibrium outcomes in games like the Centipede game (e.g., McKelvey and Palfrey, 1992) cannot be invoked in this setting. Our comprehensive experimental investigation also suggests that we cannot explain the subjects' persistent inclination to bet by means of risk-preferences,³⁷ "wrong" instructions, "weak" feedback, boring alternatives or an "unsuitable" payoff scheme.³⁸ The argument: "subjects bet simply

³⁷ Note again that the no-betting result holds independently of the risk-preferences of the agents in game S and in 1/2 of the states in game L

³⁸ Note however that our experimental investigation was restricted to a set-up with low-medium incentives.

because they love betting or enjoy playing games like the SG betting game (e.g., they derive utility from betting)" also seems irrelevant since our subjects kept betting and losing even when the alternative was a degenerated bet (design 4.4). With this regard, note also that the betting rate of our informed subjects was not lower than the betting rate of the uninformed subjects, even though the uninformed types did not know the exact structure of the game.

The current results are consistent with the findings of Ball, Bazerman and Carroll (1991) on the "Acquiring a Company" game. The participants in Ball et al's game were asked to bid a price for a company that is worth X to the current owner and $1.5X$ to the potential buyer. Subjects were told that X is uniformly distributed on $[0, 100]$ and that the current owner would accept bids higher than X . The only rational solution to the game is bidding 0. The average initial bid however was around 60 and almost did not decrease in 20 repetitions of the task. Thus, like the participants in our current study, the bidders in the acquiring game fail to converge to the rational solution.³⁹

It is important, however, to emphasize that our results do not imply that subjects cannot learn to play rational equilibria in games with a unique rational outcome. There is no doubt that very detailed instructions, use of the "strategy method" or role switching across players might help subjects move closer to the rational solutions (Søvik 2000). The main implications of the current study however are that "experience" per se does not guarantee elimination of zero sum betting. The fact that 2-parameters models can be used to approximate the learning trends in 21 different tasks suggests that general approximations of behavior can be obtained. The observed learning process in our experiments, in particular, seems relatively robust with respect to the experimental manipulations.

References

Ball, S.B., Bazerman, M.H., and Carroll, J.S. (1991). "An evaluation of learning in Bilateral winner's curse", *Organizational Behavior and Human Decision Processes*, 1991, 48, 1-22.

Camerer, C. F. and Ho, T. (1999). "Experience-weighted attraction learning in normal form games", *Econometrica* 67, Vol. 4, 827-874.

³⁹Note however that the "Firm Acquiring Game" is quite different from the GS betting game. Moreover, the solution process (leading to the unique rational outcome) is also very different across the two games.

- Camerer, C., Ho, T., Rhee and Weigelt, K. (1997). "The effect of incentives in experimental p-beauty contests", Caltech working paper.
- Chung, Y. W., and Friedman, D. (1999). "A comparison of learning and replicator dynamics using experimental data", forthcoming, *Journal of Economic Behavior and Organization*.
- Daniel, T. E., Rapaport, A. and Seale, D. A., (1998). "Strategic play and adaptive learning in sealed bid mechanism", *Journal of Mathematical Psychology*, 42, 133-166.
- Cooper, D., Feltovich, N., Roth, A. and Zwick, R. (2000). "Relative versus absolute speed of adjustment in strategic environment: responder behavior in ultimatum game. Working Paper. University of California, San Diego.
- Erev, I., Bereby-Meyer, Y. and Roth, A., (1999). "The effect of adding constant to all payoffs: experimental investigation and implications for reinforcement learning models", Working Paper, Technion.
- Erev, I. and Rapoport, A. (1990), "Provision of step-level public goods: The sequential contribution mechanism". *Journal of Conflict Resolution*, 34, 401-425.
- Erev, I., and Roth, A.E., (1998). "Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria", *American Economic Review*, 4, 848-881.
- Fudenberg, D. and Levine, D. K. (1998). *The theory of learning in games*. MIT press.
- Goeree and Holt (1999). "Stochastic game theory: for playing games. Not just for doing theory," *Proceedings of the National Academy of Sciences*, 10564-10567.
- Ho, T., Weigelt, K. and Camerer, C. (1996). "Iterated dominance and iterated best-best response in experimental 'P-beauty contests'", Working Paper Caltech.
- Milgrom, P. and N. Stokey (1982). "Information, trade and common knowledge", *Journal of Economic Theory*, 26, 17-27.
- McKelvey, R. D. and Palfrey, T. R. (1992). "An experimental study of the Centipede game," *Econometrica*, 60, 803-836.
- Mookherjee, D. and Sopher, B. (1997). "Learning and decision costs in experimental constant sum games", *Games and Economic Behavior*, 19, 97-132.
- Nagel, R. (1995). "Unraveling in guessing games: An experimental study", *American Economic Review*, 85, 1313-1326.
- Perets, H. and Sonsino, D. (1999). "On cheap talk and the impossibility of zero sum betting", *International Game Theory Review*, Vol. 1, Nu 2, 193-196.

- Rapoport, A. and Amaldoss, W. (2000). "Mixed strategies and iterative elimination of strongly dominated strategies: An experimental investigation of states of knowledge," *Journal of Economic Behavior and Organization*, 42, 483-521.
- Rapoport, A and Erev, I. (1994). "Provision of step-level public goods: effects of different information structures." In U. Schulz, W. Albers, and U. Mueller (Eds), *Social Dilemmas and Cooperation*. New-York; Springer-Verlag.
- 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*, 3, 3-24.
- Roth, A.E., Erev, I., Slonim, R.L. and Barron, G. (2000), "Learning and equilibrium as useful approximation: Accuracy of prediction on randomly selected constant sum games. Harvard University.
- Roth, A. E., and Malouf, M. W. K. (1979). "Game-theoretic models and the role of information in bargaining", *Psychological Review* , 86, 574-594.
- Sarin, R. and Vahid, F. (1998). "Predicting how people play games: A procedurally rational model of choice", mimeo, Texas A&M University.
- Sebenius, J. and Geanakoplos, J. (1983). "Don't bet on it", *Journal of the American Statistical Association* 78 424-426.
- Selten, R. (1995). "Axiomatic characterization of the quadratic scoring rule", mimeo, University of Bonn.
- Søvik, Y. (1999). "Impossible bets: an experimental study," Mimeo, University of Oslo.
- Søvik, Y. (2000). "Strength of dominance and depth of reasoning," Mimeo, University of Oslo.
- Sonsino, D. (1998). "Sebenius and Geanakoplos model with noise", *International Journal of Game Theory*, 27, 111-130.
- Stahl, D. O.(1993). "The evolution of smart_n players", *Games and Economic Behavior*, 5(4), 604-17.
- Stahl, D. O. (1996). "Boundedly rational rule learning in a Guessing game", *Games and Economic Behavior*, 16, 303-330.
- Stahl, D. O. and Wilson, P. (1995). "On players models of other players: theory and experimental evidence", *Games and Economic Behavior*, 10, 218-254.
- Tang Fang-Fang (2001). "Anticipatory learning in two-person games: some experimental results", *Journal of Economic Behavior and Organization*, 44, 221-232.

Appendix 1: Instructions

A.1: Instruction for Informed Subjects in Experiment 3

In the next hour you will participate in an experiment simulating a simple stock-market activity with two participants. The experiment is subdivided into two parts, with multiple rounds in each part. In each round, you have to decide whether you are interested in entering the market or not. If you decide to enter, you should press the <Y> key. If you decide not to enter, you should press the <N> key.

The screen display has three windows:

The lower window displays, in each round of the experiment, a signal (one letter, two letters or three letters) that gives you partial information regarding the state of the market. The signal should be considered a "hint" concerning the realized state of the market. Along the experiment, you should try and understand the meaning of these signals. The actual state of the market is described by the letters A, B, C and D. The probability of drawing each of these states is equal and independent across rounds.

The middle window displays your profit (or loss) in points, for the current round: If you have decided not to enter, trading will not take place and your profit for the round will be one point. If you decide to enter, but the other player has decided not to enter, again trading will not take place and your profit will be one point. Only if both players decide to enter, trading will take place. Your realized profit (or loss) will then be determined by the following rule:

Realized state of nature	A	B	C	D
Profit for player 1	+32	-28	+20	-16
Profit for player 2	-32	+28	-20	+16

The boxes in the figure represent the information-signals observed by the players. For instance, when the realized state of nature is B, player 1 will receive (in the lower window of the display) the signal AB, while player 2 will get the signal BC. When the realized state of nature is A, player 1 will observe the signal AB while player 2 will observe the signal A.

The upper window displays the accumulated score.

You start each part of the experiment with 1000 points.

You will receive 20 Shekel for participating in the experiment but you may gain a bonus of an additional 20 Shekel contingent on your performance. Your chances to win the additional bonus increase with your final cumulative score at the end of the experiment.

In part B of the experiment you will be asked to perform the same task for a different stock-exchange market. There will be a short break between parts.

You are player 2

Part B

In this part you are asked to repeat the previous task for a different market. The realized state of the market can now be A, B, C, D, E, F, G or H.

You start this part of the experiment with 1000 points; however your chances of winning the 20 I.S. additional bonus will be determined by your total final score for the two parts of experiment.

The payoff rule for this part is as follows:

Realized state of nature	A	B	C	D	E	F	G	H
Profit for player 1	+32	-28	+20	-16	+23	-10	-17	+14
Profit for player 2	-32	+28	-20	+16	-23	+10	+17	-14

The boxes in the figure represent again the information-signals observed by the players. For instance, when the realized state of nature is F, player 1 will observe the signal HF, while player 2 will observe the signal DEF. When the realized state of nature is A, player 1 will observe the signal AB while player 2 will observe A.

You are player 2

A.2: Instruction for Uninformed Subjects in Experiment 3

In the next hour you will participate in an experiment simulating a simple stock-market activity with two participants. The experiment is subdivided into two parts, with multiple rounds in each part. In each round, you have to decide whether you are interested in entering the market or not. If you decide to enter, you should press the <Y> key. If you decide not to enter, you should press the <N> key.

The screen display has three windows:

The lower window displays, in each round of the experiment, a signal (one letter, two letters or three letters) that gives you partial information regarding the state of the market. The signal should be considered a "hint" concerning the realized state of the market. Along the experiment, you should try and understand the meaning of these signals. The actual state of the market is described by the letters A, B, C and D. The probability of drawing each of these states is equal and independent across rounds.

The middle window displays your profit (or loss) in points, for the current round: If you have decided not to enter, trading will not take place and your profit for the round will be one point. If you decide to enter, but the other player has decided not to enter, again trading will not take place and your profit will be one point. Only if both players decide to enter, trading will take place. Your realized profit (or loss) will then be determined by some fixed payoff rule that assigns a given payoff to each state of nature.

The upper window displays the accumulated score.

You start each part of the experiment with 1000 points.

You will receive 20 Shekel for participating in the experiment but you may gain a bonus of an additional 20 Shekel contingent on your performance. Your chances to win the additional bonus increase with your final cumulative score at the end of the experiment.

In part B of the experiment you will be asked to perform the same task for a different stock-exchange market. There will be a short break between parts.

You are player 2

Part B

In this part you are asked to repeat the previous task for a different market. The realized state of the market can now be A, B, C, D, E, F, G or H.

You start each part of the experiment with 1000 points.

You will receive 20 Shekel for participating in the experiment but you may gain a bonus of an additional 20 Shekel contingent on your performance. Your chances to win the additional bonus increase with your final cumulative score at the end of the experiment.

You are player 2

A.3: Instruction for Informed Subjects in Experiment 4.3 (Short Game, Part A)

You are given the possibility to participate in a betting-game as follows:

The game has two players: One and Two

The bet has four possible outcomes represented by letters: A, B, C and D

The probability of each outcome is 1/4

The payoff to both players will be determined by the following rule:

If the realized state is A, Player 2 will pay Player 1 32 points

If the realized state is B, Player 1 will pay Player 2 28 points

If the realized state is C, Player 2 will pay Player 1 20 points

If the realized state is A, Player 1 will pay Player 2 16 points

You have to decide whether you would want to participate in the bet without knowing the actual realized state. Yet, the bet's organizers will give you some partial information (signal) concerning the realized state before your decision making.

Player One will receive the signal **AB** when the realized state is A or B. He will get the signal **CD** when the realized state is C or D.

Player Two will receive the signal **A** when the realized state is A, the signal **BC** when the realized state is B or C, and the signal **D** when the realized state is D.

The structure of the bet is summarized in the next figure:

Realized state of nature	A	B	C	D
Profit for player 1	+32	-28	+20	-16
Profit for player 2	-32	+28	-20	+16

The bets' payoffs will be actually played only if both players (One and Two) will decide to accept the bet. If one of the players will decide to reject, the bet will not take place and each player will get a payoff of +1.

The experiment includes 250 rounds. In each round of the experiment, you have to decide whether you would like to take the bet or not. If you decided to accept the bet, you should press **Y**. Otherwise, you should press **N**.

After each round of the experiment, you will observe a feedback window disclosing the following information:

Your response

The other player's response

The realized state

Your payoff for that round

Your total accumulated score

To exit the feedback window and move to the next trial, you should press **9**.

You start the experiment with 2000 points.

You will receive 20 Shekel for participating in the experiment but you may gain a bonus of an additional 20 Shekel contingent on your performance. Your chances to win the additional bonus increase with your final cumulative score at the end of the experiment.

Figure 1: The Basic Betting Games

Game S

	A	B	C	D
Agent One	+32	-28	+20	-16
Agent Two	-32	+28	-20	+16

Game L

	A	B	C	D	E	F	G	H
Agent One	+32	-28	+20	-16	+23	-10	-17	+14
Agent Two	-32	+28	-20	+16	-23	+10	+17	-14

Figure 2a: Average Betting Rates (Experiment 3)

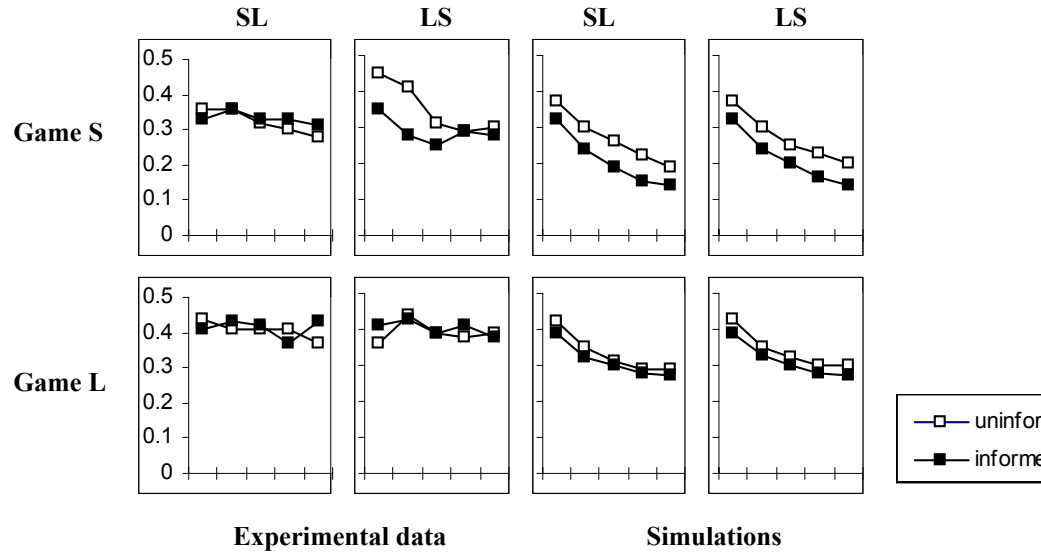


Figure 2b: Average Individual Entry Rates (Experiment 3)

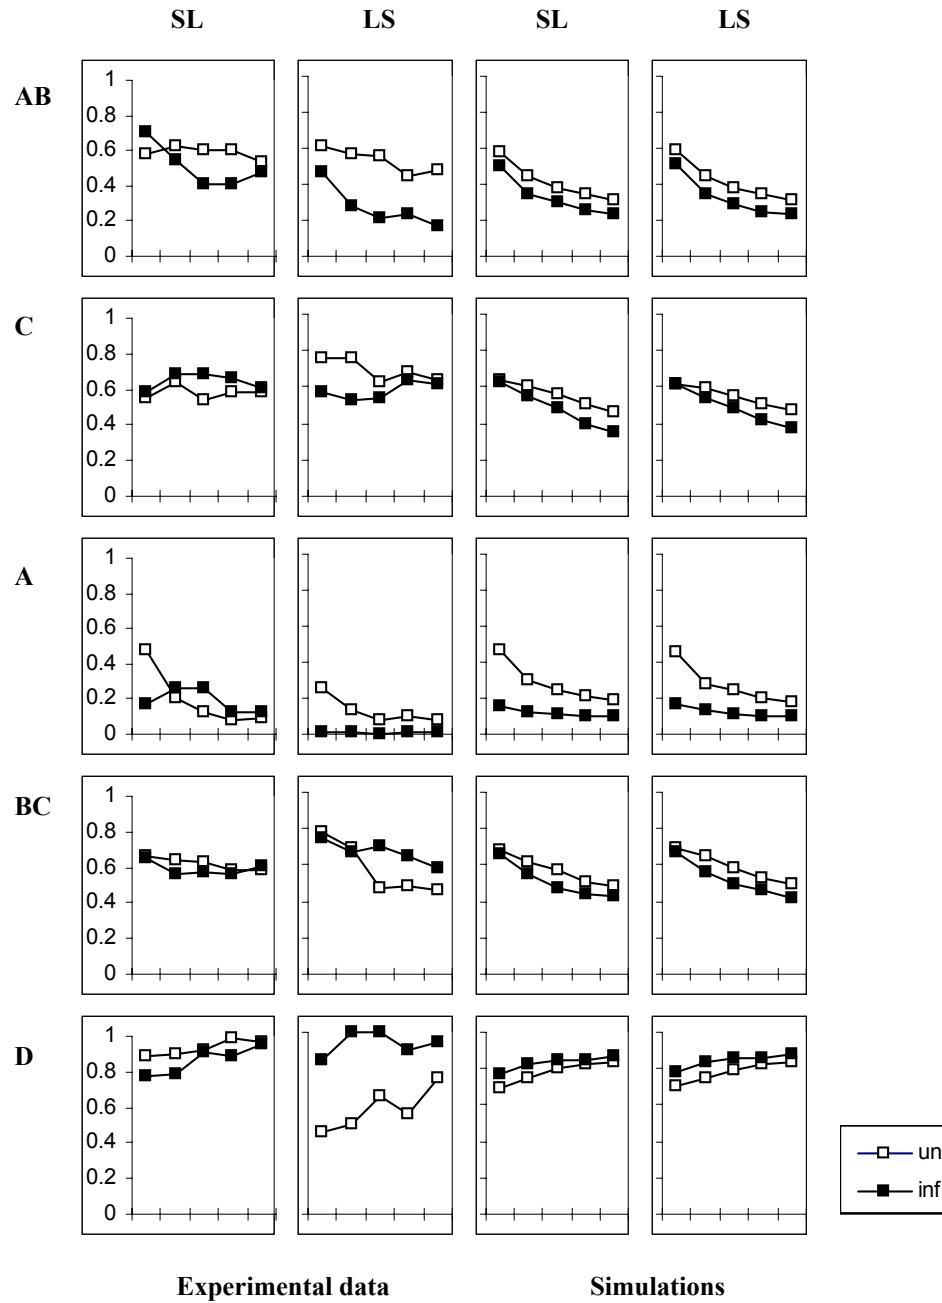


Figure 3a: Average Betting Rates (Experiments 4.1, 4.2)

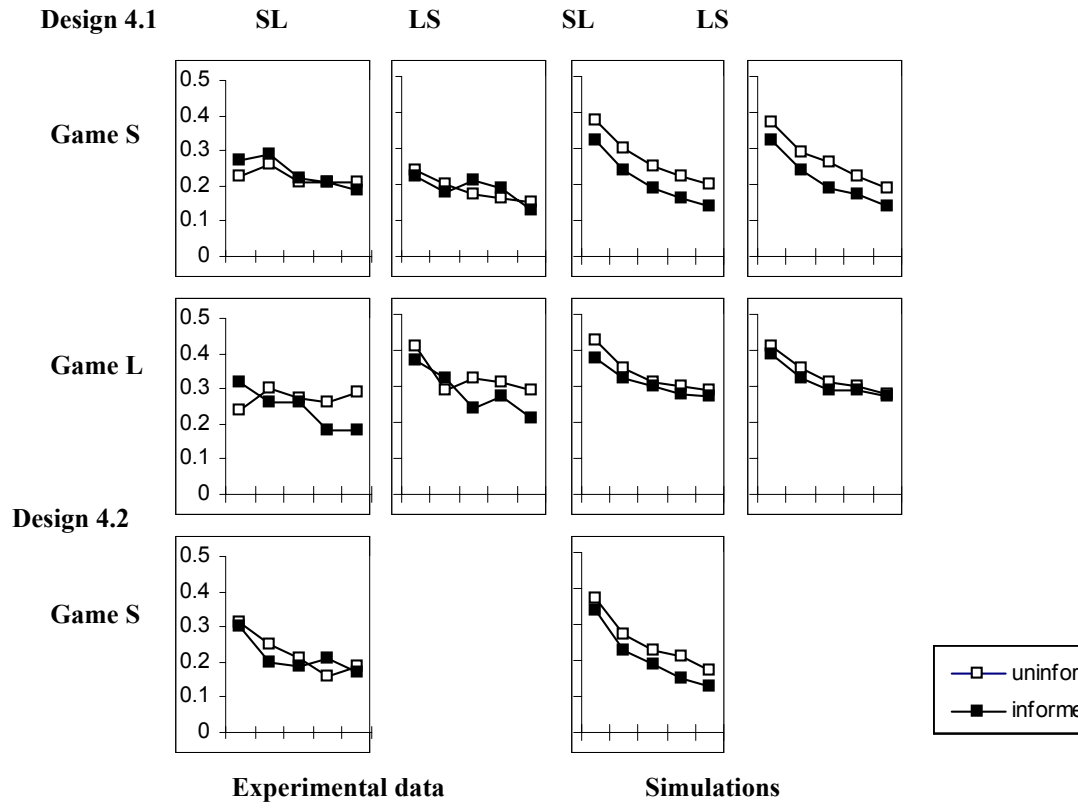


Figure 3b: Average Individual Entry Rates (Experiments 4.1, 4.2)

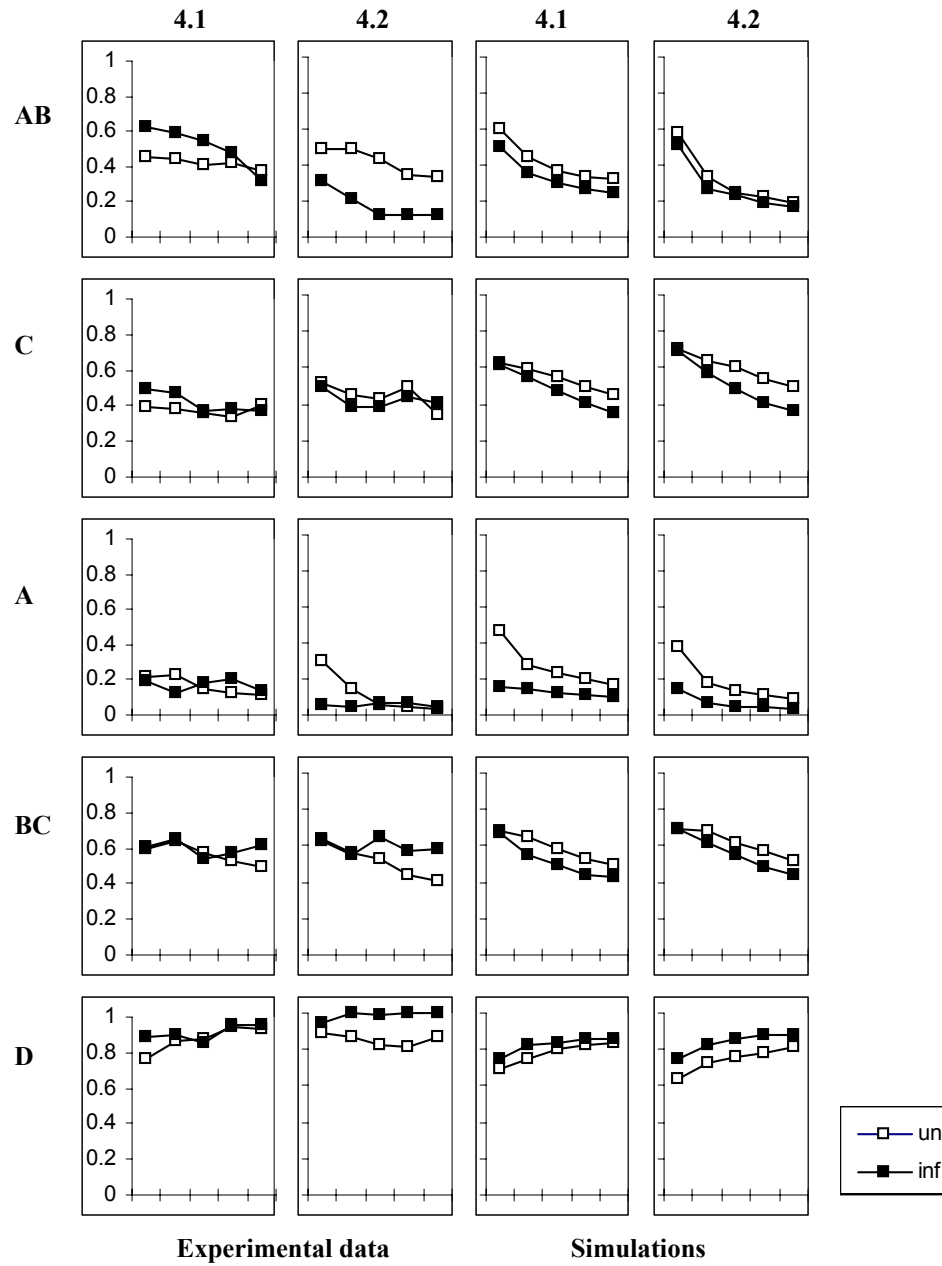


Figure 4a: Average Betting Rate (Experiments 4.3-4.5)




Figure 4b: Average Individual Entry Rates (Experiment 4.3-4.5)

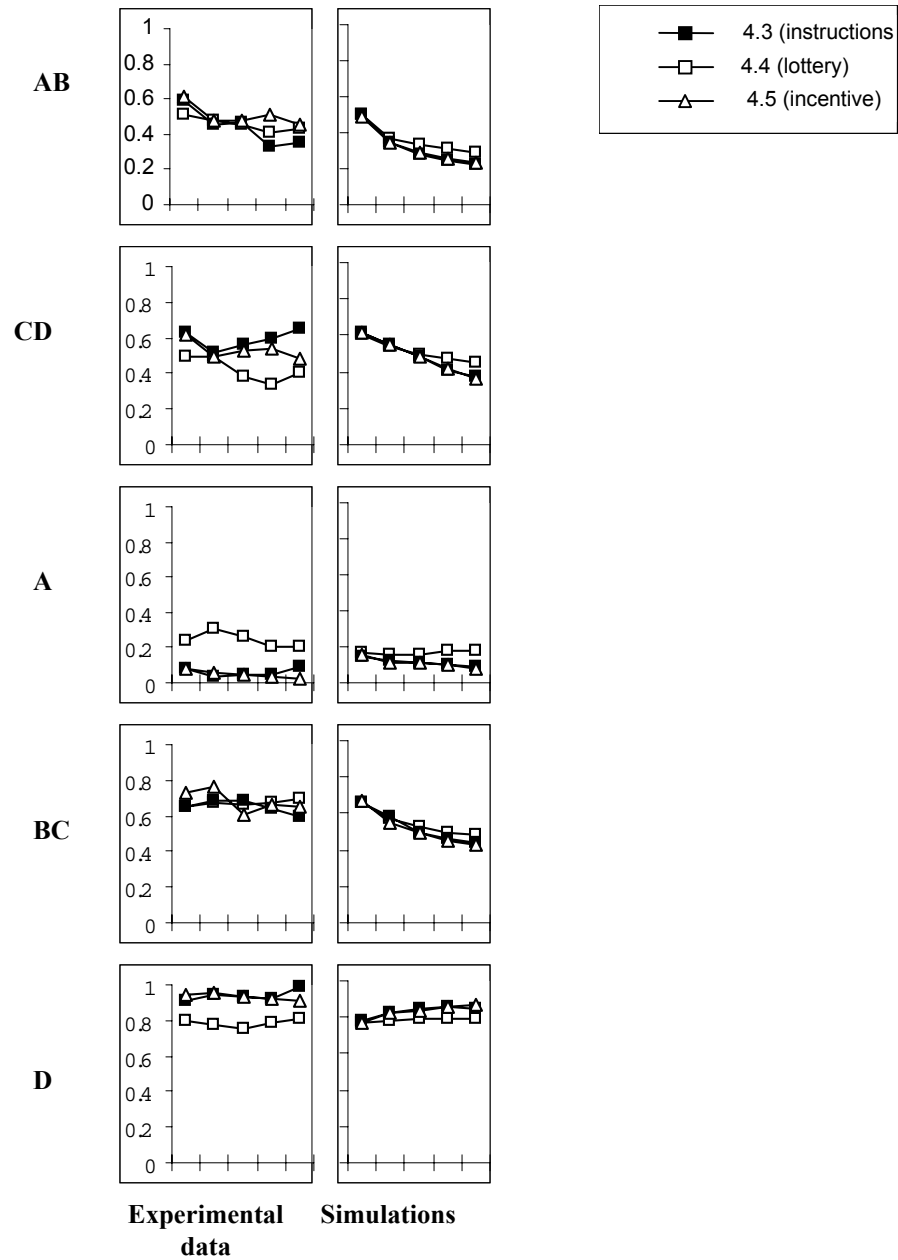


Table 1: Individual Initial Entry Rates in Each Information Set

			Full Information									Low information
Information set Levels required		A 1	D 1	AB 2	BC 3	CD 4	DEF 5	EG 6	FG 7	FH 8		
Design Basic	Game	Order										
	S	SL	0.33	0.50	1.00	0.75	0.83	.	.	.	0.67	
	S	LS	0.08	1.00	0.75	0.75	0.58	.	.	.	0.83	
	L	SL	0.17	.	0.67	0.42	0.92	0.58	0.92	0.75	0.83	
	L	LS	0.25	.	0.67	0.75	0.67	0.83	0.83	0.75	0.75	
4.1	S	SL	0.50	0.75	0.88	0.75	0.38	.	.	.	0.65	
	S	LS	0.00	1.00	0.38	0.63	0.63	.	.	.	0.53	
	L	SL	0.25	.	0.50	0.50	0.75	0.38	0.88	0.75	0.63	
	L	LS	0.25	.	0.88	0.88	0.88	0.75	0.88	0.75	0.88	
4.2	S'		0.13	0.75	0.75	0.88	0.63	.	.	.	0.75	
4.3	S		0.17	1.00	0.83	0.75	0.92	
4.4	S''		0.25	0.92	0.92	0.67	0.92	
4.5	S		0.17	1.00	1.00	0.83	0.67	
Mean			0.21	0.87	0.77	0.71	0.73	0.64	0.88	0.75	0.77	

Table 2: Gain and Loss (in Points) to each Player⁴⁰

	Player One	Player Two
Game S		
Informed	-308	+653
Uninformed	-389	+719
Game L		
Informed	+136	+161
Uninformed	+162	+138

⁴⁰ The figures in the table represent the average of the following measure of gain/loss per player: [Score at the end of the experiment - Score at the beginning] / [Score at the beginning].

Table 3: Average Betting Rates in Blocks of 50 Rounds at the Short Game⁴¹

3a. The Case of Informed Subjects

Design\Block	Block 1	Block 2	Block 3	Block 4	Block 5	Average
Design 3	34%	32%	29%	31%	31%	31%
Design 4.1	25%	24%	22%	20%	16%	21%
Design 4.2	30%	20%	19%	21%	17%	21%
Design 4.3	35%	30%	34%	30%	30%	32%
Design 4.4	27%	31%	21%	18%	24%	24%
Design 4.5	37%	30%	28%	29%	23%	29%

3b. The Case of Uninformed Subjects

Design\ Block	Block 1	Block 2	Block 3	Block 4	Block 5	Average
Design 3	41%	39%	32%	30%	29%	34%
Design 4.1	24%	23%	19%	19%	18%	21%
Design 4.2	31%	25%	21%	16%	19%	22%

⁴¹ We had 12 pairs of subjects in design 3, 4.3, 4.4 and 4.5 and 8 pairs of subjects in designs 4.1 and 4.2.

Table 4: Average Betting Rates in Blocks of 50 Rounds at the Long Game⁴²

4a. The Case of Informed Subjects

Design\ Block	Block 1	Block 2	Block 3	Block 4	Block 5	Average
Design 3	41%	43%	41%	39%	41%	41%
Design 4.1	35%	29%	25%	23%	20%	26%

4b. The Case of Uninformed Subjects

Design\ Block	Block 1	Block 2	Block 3	Block 4	Block 5	Average
Design 3	40%	43%	40%	40%	38%	41%
Design 4.1	33%	30%	30%	29%	29%	30%

⁴² We had 12 pairs of subjects in designs 4 and 8 pairs of subjects in designs 4.1 and 4.2.

Table 5: Level-n Strategies for Games S and L ⁴³

Set: Strategy:	Game S					Game L							
	Player II		Player II			Player I				Player II			
	AB	CD	A	BC	D	AB	CD	EG	FH	A	BC	DEF	GH
Level-0	1	1	1	1	1	1	1	1	1	1	1	1	1
Level-1	1	1	0	1	1	1	1	1	1	0	1	1	1
Level-2	0	1	0	1	1	0	1	1	1	0	1	1	1
Level-3	0	1	0	0	1	0	1	1	1	0	0	1	1
Level-4	0	0	0	0	1	0	0	1	1	0	0	1	1
Level-5	0	0	0	0	1	0	0	1	1	0	0	0	1
Level-6	0	0	0	0	1	0	0	0	1	0	0	0	1
Level-7	0	0	0	0	1	0	0	0	1	0	0	0	0
Level-8	0	0	0	0	1	0	0	0	0	0	0	0	0

Table 6: Comparison of alternative models of the adaptation process.

Model	Initials	Estimated parameters	MSD
RL with normalization	General	$N(1) = 100, \lambda=6$.0181
	General	Task specific (42)	0.0130
	Task Specific	$N(1) = 100, \lambda=6$	0.0184
	Uniform	$N(1) = 100, \lambda=8$	0.0246
RL no normalization	General	$N(1) = 1, \lambda=.25$	0.0270
SFP with normalization	General	$N(1) = 100, \lambda=8$	0.0216
SFP no normalization	General	$N(1) = 1, \lambda=.25$	0.0311

⁴³ The figures at the table specify the probability of entry in every information set.