ECONOMIC RESEARCH REPORTS

AN EXPERIMENTAL STUDY OF LEARNING IN ONE AND TWO-PERSON GAMES

BY

Antonio Merlo
and
Andrew Schotter

RR # 94-17 **May 1994**

C. V. STARR CENTER FOR APPLIED ECONOMICS

NEW YORK UNIVERSITY
FACULTY OF ARTS AND SCIENCE
DEPARTMENT OF ECONOMICS
WASHINGTON SQUARE
NEW YORK, N.Y. 10003

AN EXPERIMENTAL STUDY OF LEARNING IN ONE AND TWO–PERSON GAMES

Abstract

In the economic theory of learning, agents approach each learning task identically. If they are perfectly rational Bayesians, they update their priors given revealed information and maximize accordingly, while if they are boundedly rational they mechanically apply an exogenously given and fixed rule of thumb. In both of these approaches agents do not adapt their learning strategy to fit the circumstances they are faced with. In this paper, we view economic agents as boundedly rational. However, we view them as taking a more active and creative part in the learning that they do. This creativity is manifested in their choice of what to learn about depending on the situation they are faced with.

We conduct a series of one and two–person game experiments aimed at investigating how the learning strategy of laboratory subjects changes as they are presented with decision problems of varying complexity across environments with different learning costs. We find that learning is both institution and situation specific. More precisely, our findings show that as the complexity of the decision problem increases, subjects employ simpler learning rules. Such a shift is more pronounced in a payoff environment characterized by low learning costs. Further, holding the complexity of the decision problem constant and looking across different payoff environments, we find that subjects learn less well in a high learning cost environment despite the fact that they appear to sample identically over the domain of their strategy set as their cohorts in low learning cost experiments. In other words, the relatively better performance of subjects in an environment where learning is less costly is a function of how such subjects process the information they generate and not of the type of information they generate. These results are more pronounced in our one–person decision problem experiments than in our two–person game experiments.

Section 1: Introduction

In recent years, both economic theorists and experimentalists have turned their attention to trying to characterize how economic agents learn in markets and in situations of individual optimization or games. In this literature, it is typically assumed that the agents studied are either fully rational Bayesian agents capable of performing all of the calculations required of them or boundedly rational agents pursuing adaptive learning rules.[1] While these two approaches may appear to be opposites in their modelling strategies and their characterizations of human capabilities, they are quite similar methodologically since they both assume that economic agents do not take an active or creative role in the learning that they do. If agents are rational, they employ Bayes rule and calculate Nash equilibria along the way, while if they are boundedly rational they mechanically employ the rule–of–thumb learning rule given to them exogenously by the modeler. The point is that economic agents are not assumed to choose what they want to learn about and their learning strategies do not vary either with the complexity of the problem they face or with other payoff relevant aspects of the environment they are in.

We view economic agents as boundedly rational. However, we view them as taking a more active and creative part in the learning that they do. This creativity is manifested in their choice of what to learn about and how the object of their learning changes to fit the circumstances they are faced with. We study how laboratory subjects go about learning in one–person decision problems and two–person games. Unlike most models of learning in which agents are placed in an environment with incomplete information, our experiments present laboratory subjects with complete information games

[1] For the Bayesian approach see, for example, Jordan (1991), Kalai and Lehrer (1993), and Nyarko (1992) for game–theoretic analyses and Aghion et al. (1991), Easley and Kiefer (1988), and McLennan (1987) for an analysis of individual decision problems. See also El–Gamal, McKelvey, and Palfrey (1992a, 1992b) for experimental analyses. Theoretical examples of the bounded rationality approach can be found in Bray (1982), Brown (1951), Holland (1975), Marcet and Sargent (1989), and Marimon and McGrattan (1994), while Arthur (1990), Crawford (1992), Cheung and Friedman (1994), and Roth and Erev (1993) offer experimental evidence.

with stochastic uncertainty which are easy enough to comprehend yet too complex for them to solve deductively. As a consequence, they must attempt to learn inductively through experience and it is this process that we are interested in studying.

We focus on four learning hypotheses that are intuitively appealing because they are sufficiently simple to actually be employed by boundedly rational agents. These hypotheses are ranked according to the complexity of the computational task they require of the subjects. The definition of complexity we use here refers to the amount of information and the number of computational steps needed to carry out a particular learning strategy.[2]

Our first hypothesis assumes that subjects in one–person decision problems attempt to learn the payoff function they face while in two–person games they try both to learn their best response function and to predict the action of their opponent. This is typically what economists assume will happen. This hypothesis makes the greatest computational demands on our subjects since learning a function requires subjects to store and process all of the data generated over the length of the experiment. However, if the problem facing our subjects is too complicated for them to behave in this manner, we might expect some of them to simplify it. Our second hypothesis postulates that subjects engaged in a one–person maximization problem transform it by selecting a small number of actions and then attempt to learn which one amongst those actions returns the highest payoff. For subjects engaged in a game, this hypothesis predicts that they also transform the game similarly by treating their opponent as a random device and then attempt to learn which action yields them the highest return. Such a learning strategy is simpler than

---

[2] Our complexity measure could be formalized by combining the definitions of Kalai and Stanford (1988), who define the complexity of a strategy by the size (number of states) of the smallest automaton needed to implement it, and Futia (1977), who defines the complexity of a decision rule by the number of computational steps needed to implement it.

our first since it only requires the subjects to compute and remember sufficient statistics (payoff means) associated with a limited number of decision choices. Our third and fourth hypotheses are the simplest since they assume that subjects do not attempt to learn either optimal (equilibrium) or satisfactory actions, but rather attempt to learn a dynamic adaptive rule such as numerical hill climbing or a myopic adjustment rule, respectively. To implement such rules, subjects need only to use data generated either over the last two periods (numerical hill–climbing) or simply the last one (myopic adjustment), and the computations required are trivial. Finally, there is the possibility that when faced with complex decision problems subjects decide not to spend the time and effort actually trying to learn and abdicate learning altogether.

In our discussion below, we formalize these hypotheses and test them with the data generated by our experiments. Our experimental design allows us to easily move between maximization and game–theoretic problems and between payoff environments characterized by different learning costs. What we find is that learning is institution and situation specific. In particular, both the object of learning and the subjects' success in learning (as measured by their payoff) are dramatically affected by the complexity of the decision problem they face and the implicit costs of learning of the payoff environment they are placed in.

Holding the payoff environment constant and looking across decision problems of different levels of complexity (i.e., one–person decision problems versus two–person games) our findings show that as the complexity of the decision problem increases, the distribution of subject behavior across learning hypotheses shifts toward simpler rules. More specifically, the proportion of subjects whose learning can be characterized according to our last two (simplest) hypotheses increases as we move from a one–person decision problem to a two–person game. Such a shift is more pronounced in a payoff environment characterized by low learning costs.

Holding the complexity of the decision problem constant and looking across different payoff environments our findings can be summarized as follows. In our one–person experiments subjects were relatively more successful in learning the relevant payoff function they faced and/or their historically best action in an environment with relatively low learning costs (what we call the Learn–Before–You–Earn environment) than in a more costly (Learn–While–You–Earn) environment, where subjects seemed to focus their learning more on adaptive rules. Such differences in learning led to corresponding payoff differences for the subjects in these two environments with subjects in the Learn–Before–You–Earn experiment earning significantly more.

It is important to note that these differences in learning are not a function of the sampling strategies engaged in across these two environments. From a statistical point of view, over the horizon of the experiment, subjects in these two environments seemed to experiment over the domain of their payoff function in identical ways and generated comparable experimentation data. What differed is how they ended up processing this data in light of what the object of their learning was.

In our two–person game experiments our results are qualitatively similar although less definitive. Subjects in a Learn–Before–You–Earn environment were relatively more successful in learning best responses and/or their historically best action than their counterparts in a Learn–While–You–Earn environment, who seemed to focus their learning more on adaptive rules. Once again, these differences led to subjects in the Learn–Before–You–Earn experiment earning relatively more than subjects in the Learn–While–You–Earn experiment. Such differences can not be attributed to differences in the sampling strategies of subjects during the experiment.

Finally, it is important to point out that the simple hypotheses we consider allow us to characterize the learning of about 90% of the subjects in our experiments. While more complex (less intuitive) hypotheses might certainly be added to our list, the expected benefits are unlikely to compensate the costs from introducing additional hypotheses.[3]

In the remainder of the paper we proceed as follows. In Section 2, we describe our experiments and our experimental design. In Sections 3, we present our results and in Section 4 we offer some conclusions and ideas for future research.

Section 2: Experimental Setting

2.1: The Games Played

All of the experiments performed to investigate learning were of the tournament variety and similar to those of Bull, Schotter and Weigelt (1987) and Schotter and Weigelt (1992).[4] In those experiments, randomly paired subjects must, in each round, each choose a number, e, between 0 and 100 called their decision number. After this number is chosen a random number is independently generated by each subject from a uniform distribution over the interval $[-a,+a]$. These numbers (each player's decision number and random number) are then added together and a "total number" defined for each player. Payoffs are determined by comparing the total numbers of the subjects in each pair and awarding that subject with the largest total number a "big" payment of M and that subject with the smallest total number a "small" payment of m, $M > m$. The cost of the decision number chosen, given by a convex function $c(e) = e^2/k$, is then subtracted from these fixed payments to determine a subject's final payoff. Hence, in these experiments there is a trade-off in the choice of decision numbers; higher numbers generate a higher probability

---

[3] In related work, El-Gamal and Grether (1994) use a likelihood-based procedure to classify the behavior of experimental subjects with respect to the probabilistic updating rules they use. Their approach explicitely considers a penalty for introducing too many rules.

[4] For a description of the theory of tournaments underlying these experiments see, e.g., Lazear and Rosen (1981).

of winning the big prize but also imply a higher decision cost.[5] By letting $k = 500$, $a = 40$, $M = 29$, and $m = 17.2$, the two–person tournament defined has a unique symmetric Nash equilibrium at 37. By replacing one player with a computerized automaton programmed to always choose the Nash equilibrium decision, we transform the problem for the remaining live player into a one–person maximization problem.

We consider these experiments to be good ones for our purposes for a number of reasons. First, although they present subjects with complete information games for which optimal actions or Nash equilibria could be calculated a priori, such games are sufficiently complex so that a deductive solution should be out of the grasp of most experimental subjects. Such complexity forces subjects to learn inductively and it is this process that we are interested in studying. Second, in spite of the complexity of the decision problems they involve, these experiments are simple to describe to subjects and to understand. This feature is appealing since it should reduce the noise in the data. Finally, since one of the objectives of our research is to study the impact of the complexity of decision tasks on the learning rules employed by subjects, our tournament games are appealing in allowing us to easily move between one–person and two–person decision problems.

2.2: Payoff Environments

We distinguish between two payoff environments which we call Learn–While–You–Earn (LWYE) and Learn–Before–You–Earn (LBYE) since they span the spectrum of environments with differing learning costs. A LWYE payoff structure is the typical payoff structure found in laboratory experiments and markets. In this environment time is divided into discrete periods with a known horizon T and in each period subjects or market participants make decisions. These decisions yield them a payoff at the end of the period, and their final payoff from the experiment or market is the sum of their (possibly

[5] In the instructions, a sample of which is contained in Appendix B, we take great care in not using such value ladened terms as "winning" or "losing".

discounted) period payoffs. The LWYE environment is then one in which payoffs occur each period and cumulate throughout the experiment. The cost of learning is the opportunity cost associated with exploring the environment and a trade–off exists between "exploiting" actions already proven to be satisfactory by using them repeatedly and "exploring" to discover new and possibly better actions.

A LBYE environment is a limit case of a payoff environment with low learning costs. It is an artificial environment created by an experimental administrator, although it has some parallels to real world markets. Here time is again divided into T discrete periods but no payoffs are awarded during the first T–1 periods. Rather, subjects make decisions and observe what they would have earned if these were played for real. What does count is their period T decisions and these payoffs are sufficiently large so that their expected payoff from this last round decision is comparable to the expected sum of the payoffs in the LWYE environment. In short, in a LBYE setting there are T–1 practice rounds and one real and lucrative round so that there is no exploration–exploitation trade–off. In terms of markets, we might consider two infinitely lived firms who interact repeatedly in a market in which they both have extremely low (possibly zero) discount rates. In such an environment, firms might treat any finite number of periods as free–learning periods since, with zero discount rates, any finite period would have only a negligible influence on their infinite horizon payoff. Under these circumstances, our LBYE environment is a reasonable approximation to reality.

2.3: Experimental Design

We performed a set of 4 different experiments and 2 computer simulations for purposes of comparison. All the experiments were computerized and conducted at the Experimental Lab of the C.V. Starr Center for Applied Economics at New York University using 117 undergraduate students recruited from economics classes at N.Y.U. Each experiment lasted approximately 45 minutes with average payoffs of about $13.00. No subject engaged in more than one experiment and had previous experience with tournament experiments.

In Experiments 1a and 1b, subjects played the tournament game described above 75 times consecutively against a computer whose strategy was known to be that of choosing 37 in each period. These experiments then presented our subjects with a one–person decision problem under stochastic uncertainty. These experiments were performed under two payoff regimes. In Experiment 1a subjects did the experiment 74 times without receiving a payoff but were paid for the decisions that they made in round 75 (LBYE environment). With the parameters specified above, the (equilibrium) expected amount for this one period choice was $15.27. (Actually, payoffs were denominated in a fictitious currency called francs and converted into dollars at the rate of $0.75 per franc).

In Experiment 1b subjects performed the same experiment 75 times but received a payoff in each of the 75 rounds. Their final payoff was the sum of their 75 round earnings over the course of the experiment (LWYE environment). In order to keep equilibrium payoffs constant across environments, here we converted francs into dollars at the rate of $0.01 per franc. After the 75 rounds of Experiment 1b were over, subjects were then informed that they would perform the experiment one more time with increased payoffs –– actually, with the same payoffs that were used in the LBYE environment. They had not been told about this extra experiment until after they had finished their 75 round experiment. In other words, subjects in Experiment 1b performed the experiment twice. Once for 75 rounds with small payoffs in each round, and once for one extra round with a one round payoff equal to the payoff in the last round of the LBYE experiment. This extra–round with increased payoffs faced LWYE subjects with a decision task identical to the one faced by LBYE subjects in round 75 (their only payoff–relevant round). As such, their extra–round choices should serve as a sufficient statistic for all that these subjects have learned during the course of the experiment as does the 75th round decision of subjects in the LBYE Experiment 1a.

Experiments 2a and 2b, increase the complexity of the learning task of subjects by having them play against a live human subject for 75 rounds instead of a computer. This adds strategic uncertainty to the already existing stochastic uncertainty present in the one–person decision problem of Experiments 1a and 1b. In Experiments 2a and 2b subjects were randomly assigned to an opponent and kept that opponent, whose identity was unknown to them, during the entire course of the experiment. This fact was common knowledge. To make the design symmetric we ran Experiment 2a under the LBYE condition and Experiment 2b under the LWYE condition again with an extra round. In Experiment 2b subjects played the extra round against that person with whom they were matched during the first 75 rounds (and that information was also common knowledge). As was true in Experiments 1a and 1b, after each round of an experiment subjects were informed about their payoff but not about the decision choice or payoff of their opponent.

Finally, in order to facilitate comparisons, we ran two simulations (one for the one–person decision problem, Experiment 1c, and one for the two–person game, Experiment 2c), also with a 75–round horizon. In these simulations, player choices were independently drawn with replacement from a uniform distribution over the set of integers between 0 and 100 representing the feasible decision set. The size of such random samples was chosen to match the size of our live experimental subject pool.

Our experimental design is summarized in Table 2.1.

[Table 2.1 About Here]

Section 3: Results

In this section of the paper we first examine the results of our one–person experiments (Experiments 1a, 1b, and 1c) and then our two–person game experiments (Experiments 2a, 2b, and 2c). In both cases, a description of the experimental data precedes the formulation and test of our hypotheses about the learning of our subjects. A comparison of the results of one–person and two–person experiments concludes the section.

3.1: One–Person Experiments

3.1.1: Overview of the Experimental Results

Our experimental design allows us to compare the learning of subjects in our two payoff environments by comparing their choices in the 75th round of Experiment 1a (the LBYE experiment) and the extra round of Experiment 1b (the LWYE experiment). These two rounds represent situations in which subjects have approximately the same length of experience in the experiment (74 versus 75 rounds) and play an identical game with substantial and identical stakes. The only difference is how they were paid in the rounds preceding these two test rounds. What we find is very different behavior in the last (extra) round of these two experiments despite the fact that subjects appear to have employed indistinguishable sampling strategies during the experiments and hence to have the same information at their disposal to guide their final decision.

Tables 3.1 and 3.2 present the last round and extra round choices of subjects in Experiments 1a and 1b, respectively. Table 3.3 presents the last period choices of 24 random players in our simulated Experiment 1c generated for comparison purposes.

[Tables 3.1, 3.2, and 3.3 About Here]

As we can see, subjects in our LBYE environment made choices in their last round which were substantially closer to the optimal choice of 37 for the one–person decision problem they were engaged in than did either subjects in the LWYE experiment or the random players in Experiment 1c. More precisely, while the mean last period choice for subjects in LBYE Experiment 1a was 42.61 (5.61 units away from the optimal choice of 37), the mean extra round choice for subjects in the LWYE Experiment 1b was 51.33 (14.33 units away from the optimal choice of 37) and the mean last period choice of the simulated random players in Experiment 1c was 51.04 (14.04 units away from 37). Kolmogorov–Smirnov tests of equality of the distributions of last and extra round choices detect significant differences between the LBYE Experiment 1a and both Experiment 1b and Experiment 1c (P–values of 0.025 and 0.048, respectively), while finding no differences between the simulated Experiment 1c and the LWYE Experiment 1b at the 5% significance level (P–value 0.837).

In terms of a money metric, the choices of our subjects in the LBYE environment of Experiment 1a led to a mean expected payoff for the experiment of $14.68 which was $1.11 and $1.16 higher than the mean expected payoff of subjects in the extra round of the LWYE Experiment 1b ($13.57) and the $13.52 mean earnings of the simulated subjects in Experiment 1c, respectively.[6] In addition, while 12 subjects in Experiment 1a had payoff losses (measured by the difference in their expected payoffs when choosing the optimal choice of 37 and their actual last round choice) of less than $0.10, only 2 subjects had such small losses in the LWYE environment of Experiment 1b (5 in Experiment 1c).

One possible explanation for these results would be that these two experiments contain very different learning costs. As a result, the sampling strategies of subjects might differ so dramatically (with subjects in the LWYE environment sampling more conservatively than their counterparts in the LBYE environment) that when they make

[6] Since the actual payoffs in the last (extra) round of the experiment depend on the particular realizations of the additive shocks in that round, our money metric uses the expected payoffs corresponding to the subjects' last (extra) round choices instead.

their choices in the last or extra round they do so using very different information sets. Our data reject this interpretation. To illustrate, we present Tables 3.4–3.6 which contain descriptive statistics of the sampling distributions of our subjects in Experiments 1a, 1b, and 1c over the 74 or 75 rounds of their history prior to making their last choice for big stakes.

[Tables 3.4, 3.5, and 3.6 About Here]

Kolmogorov–Smirnov tests of equality of the distributions of such descriptive statistics between experiments find no statistically significant differences between the distributions of either the mean, median, standard deviation, or interquantile range of these samples in the LBYE Experiment 1a and the LWYE Experiment 1b. However, significant differences are detected between the sampling distributions of subjects in both Experiments 1a and 1b and the randomly generated distributions in Experiment 1c indicating that the sampling strategies of our live subjects were not random.[7]

To summarize, from the data presented in Tables 3.1–3.6 it would appear that subjects in the LBYE environment of Experiment 1a were significantly better at the decision task presented to them than were their counterparts in the LWYE environment of Experiment 1b as measured both by the distance of their last (extra) period choices from

[7] The P–values for these Kolmogorov–Smirnov tests were:

	Mean Choice	Standard Deviation	Median Choice	Interquantile Range
1a vs. 1b	0.098	0.363	0.058	0.197
1a vs. 1c	0.001*	0.000*	0.000*	0.000*
1b vs. 1c	0.040*	0.000*	0.089	0.000*

*indicates that we can reject the null hypothesis of equality of the distributions at the 5% significance level.

the optimal choice and by their payoffs. Furthermore, these differences can not be explained by the different sampling strategies of subjects across these experiments since such differences do not appear to exist. The explanation must lie elsewhere and this is what we turn our attention to next.

3.1.2: Learning Hypotheses for One–Person Decision Problems

As we discussed in the Introduction, we view economic agents as taking an active role in their learning. As a consequence, when placed in different payoff environments and faced with decision tasks of varying levels of complexity, boundedly rational agents will adapt their learning strategies appropriately. Operationally, this means that experimental subjects must decide what it is, if anything, that they will try to learn about and it is our claim that different situations induce subjects to learn about different things. To organize our thinking we propose four simple hypotheses and test them with our experimental data. These hypotheses are ranked in descending order according to the complexity of the computational task required to implement them.

In our one–person experiments subjects face a maximization problem. The objective function in this problem is the conditional expected payoff function obtained from our two–person tournament game after restricting one of the players to be a computer whose choice is constrained to always be 37. To explain, consider Figure 3.1.

[Figure 3.1 About Here]

In Figure 3.1 we see two functions labeled W and L defined as $W(e) \equiv M - c(e)$ and $L(e) \equiv m - c(e)$, respectively. They represent the payoffs to our subjects for each choice of decision number (e) over the interval [0,100] conditional on winning (W) the large fixed payment M or losing (L) the tournament and receiving m. In our experiments subjects choose an integer between 0 and 100 and, depending on their opponent's choice and the realization of two independent random shocks added to their and their opponent's choices, either get W or L. Conditional on a computerized opponent's choice of 37 and our distributional assumptions on the additive shocks, the expected payoff function for a subject, $E\pi$, is defined by the convex combination of these two functions (W and L) with weights given by the probability, p(e), of winning M at each decision number; $E\pi(e) \equiv p(e) \cdot M + (1 - p(e)) \cdot m - c(e)$. Such a (quadratic) function is also depicted in Figure 3.1.[8]

Faced with such a maximization problem, subjects in Experiments 1a and 1b might choose to learn the shape of this expected payoff function. This leads us to formulate our first hypothesis:

H1: *Subjects in the last (extra) round of Experiment 1a (1b) use the data they have generated during the course of the experiment to estimate the expected payoff function they face and choose a decision number consistent with its estimated maximum.*

This hypothesis is our most demanding computationally since it requires subjects to use all of the observations they generated during the course of the experiment to estimate the payoff function they face and then use this piece of information to guide their behavior in selecting their final choice.[9]

We test this hypothesis by asking if subjects in Experiments 1a and 1b acted as if they were econometricians who used the observations generated by their first 74 (75) round decisions to estimate the quadratic function that best fit the data and then chose a decision number in the last (extra) round of the experiment within the 95% confidence interval of their point estimate of the payoff maximizing choice. Tables 3.7 and 3.8

---

[8] The derivation of such function involves straightforward calculations using the formulas reported in Schotter and Weigelt (1992).

[9] This hypothesis was also investigated by Merlo and Schotter (1992) in similar experiments run over a shorter horizon (12 periods) in a LWYE environment.

contain the results of such a test for Experiments 1a and 1b, respectively. For each of the subjects in these experiments, these tables report their estimated payoff maximizing choice together with its standard error, the subject's final choice, and its absolute difference from the estimated payoff maximizing choice.[10]

[Tables 3.7 and 3.8 About Here]

As we can see, the final choices of subjects in the LBYE Experiment 1a were far more consistent with the behavior postulated by hypothesis H1 than were the choices of subjects in the LWYE Experiment 1b. While 18 out of 23 subjects in the LBYE environment made choices in the last round that were within the 95% confidence interval of their estimated payoff maximizing choice, only 9 out of 24 subjects did so in the extra round of the LWYE environment. Also, note that the data generated by subjects in Experiment 1b was sufficient to obtain as good estimates of the true payoff function as was the data generated by subjects in Experiment 1a. The mean of the estimated optimal choices is remarkably close to the true optimal choice of 37 in both Experiment 1a (34.7) and Experiment 1b (40.4) and the means of the estimates' standard errors in the two experiments are similar (9.4 in Experiment 1a versus 10.9 in Experiment 1b).[11] In summation, while a vast majority of subjects in the LBYE Experiment 1a behaved in a manner that was consistent with hypothesis H1, the same can not be said about subjects

[10] Standard errors are computed using the Delta Method from the least–squares estimates reported in Tables A.2 and A.3, as illustrated in Table A.4 in Appendix A. For a general description of this method see, for example, Goldberger (1991).

[11] Also, F–tests of the hypothesis of equality of the estimated coefficients to the coefficients of the theoretical payoff function reported in Tables A.2 and A.3 in Appendix A indicate that the null hypothesis can not be rejected at the 5% significance level for 12 out of the 23 subjects in Experiment 1a (Table A.2) and 13 out of the 24 subjects in Experiment 1b (Table A.3).

in the LWYE Experiment 1b. Again, we emphasize the fact that these discrepancies are not a direct consequence of fundamental differences in the sampling strategies adopted by subjects in the two experiments. Rather, having access to virtually identical data, subjects in Experiment 1b chose not to use it efficiently, unlike their counterparts in Experiment 1a.

A less computationally demanding hypothesis assumes that subjects in our one–person experiments choose to simplify the problem and learn which one of a limited subset of decision numbers returns the highest expected payoff. This implies that during the experiment subjects might choose a few decision numbers relatively often in an effort to obtain estimates of the expected returns they yield. This is, in a static sense, comparable to treating the problem as a multi–armed bandit problem and yields the following hypothesis:

H2: *Subjects in the last (extra) round of Experiment 1a (1b) use the data they have generated during the course of the experiment to estimate the expected payoffs associated with a subset of decision numbers and choose that decision number from amongst this subset which yields the highest expected return.*

Note that this hypothesis is simpler than our first since it only requires subjects to compute a limited number of sufficient statistics (payoff means) instead of estimating a function.

We test this hypothesis by looking at the distribution of choices each of our subjects made in Experiments 1a and 1b. We label a decision number as an "arm" if it was chosen at least 5 times during the first 74 (75) rounds of the experiment and for each arm we compute its historical average payoff. We then ask if a subject chose in the last (extra) round of the experiment that arm whose average return was highest. The results of such a test for Experiments 1a and 1b are reported in Tables 3.9 and 3.10, respectively. For each of the subjects in these experiments, Tables 3.9 and 3.10 report those decision numbers that could be classified as arms along with the number of times they were chosen during the experiment, their realized mean payoff and standard deviation, the subject's final choice, and the number of times that that decision number was chosen during the experiment.

[Tables 3.9 and 3.10 About Here]

As we can see, in the LBYE Experiment 1a 11 of the 23 subjects can be characterized as exhibiting behavior consistent with hypothesis H2, while this is true for only 3 of the 24 subjects in the LWYE Experiment 1b. Once again, such differences in behavior are in spite of the fact that as in Experiment 1a, almost all the subjects in Experiment 1b seemed to generate arms and could have estimated the average payoffs associated with those arms but chose not to.

Hypotheses H1 and H2 concentrate on the actions of subjects in the last (extra) round of Experiment 1a (1b) and assume that they base their choices on the full history of the experiment. Another possibility, however, is that the object of learning for our subjects is an adaptive rule which they use to react to the feedback they receive during the experiment. The two adaptive rules we consider are numerical hill climbing and what we call myopic adjustment. Numerical hill climbing assumes that subjects change their decision choices myopically in the direction of increasing payoffs. Formally, if we let e_t be the choice of a subject in period t ($t = 3,\ldots,75$), and $\pi_t(e_t)$ be the realized payoff associated with such a choice, then numerical hill climbing implies the following set of restrictions on the behavior of that subject:

(i) If $e_{t-1} > e_{t-2}$ and $\pi_{t-1}(e_{t-1}) > \pi_{t-2}(e_{t-2})$, then $e_t > e_{t-2}$
(ii) If $e_{t-1} < e_{t-2}$ and $\pi_{t-1}(e_{t-1}) > \pi_{t-2}(e_{t-2})$, then $e_t < e_{t-2}$
(iii) If $e_{t-1} > e_{t-2}$ and $\pi_{t-1}(e_{t-1}) < \pi_{t-2}(e_{t-2})$, then $e_t < e_{t-1}$
(iv) If $e_{t-1} < e_{t-2}$ and $\pi_{t-1}(e_{t-1}) < \pi_{t-2}(e_{t-2})$, then $e_t > e_{t-1}$
(v) If $e_{t-1} \neq e_{t-2}$ and $\pi_{t-1}(e_{t-1}) = \pi_{t-2}(e_{t-2})$, then e_t is unrestricted
(vi) If $e_{t-1} = e_{t-2}$ and e_{t-1} is consistent with (i)–(v), then $e_t = e_{t-1}$.

This leads us to formulate our third hypothesis:

H3: *Over the course of Experiments 1a and 1b subjects' behavior is consistent with numerical hill climbing as defined by conditions (i)–(vi) above.*

Note that this learning rule uses only information concerning the previous two rounds of the experiment and involves a simple computation of payoff rates of change.

Under the myopic adjustment rule we consider, subjects simply respond to whether or not they won the large fixed payment (M) in the previous period of the experiment. We formulate this rule since previous investigations of laboratory tournament games (Bull, Schotter and Weigelt (1987) and Schotter and Weigelt (1992)) demonstrated that subjects sometimes focus on this factor at the expense of their monetary payoffs. Formally, our myopic adjustment rule imposes the following restrictions on the behavior of subjects:

(i') If $\pi_{t-1}(e_{t-1}) = M - c(e_{t-1})$, then $e_t \leq e_{t-1}$
(ii') If $\pi_{t-1}(e_{t-1}) = m - c(e_{t-1}) > 0$, then $e_t > e_{t-1}$
(iii') If $\pi_{t-1}(e_{t-1}) = m - c(e_{t-1}) \leq 0$, then $e_t < e_{t-1}$.

Under this rule, subjects respond to winning the large fixed payment (M) by either keeping their previous choice and hence the probability of winning M constant, or by lowering their decision number in an effort to save on decision cost. When they lose, subjects increase their decision number in an effort to increase their probability of winning M as long as their payoff is positive. Otherwise, they lower their decision number. These considerations lead us to define our fourth hypothesis:

H4: *Over the course of Experiments 1a and 1b subjects' behavior is consistent with the myopic adjustment rule defined by conditions (i')–(iii') above.*

Note that this learning rule is our simplest since it uses only information concerning the previous round of the experiment and involves no computation whatsoever.

We test hypotheses H3 and H4 by looking at the behavior of our subjects over the last 50 rounds of Experiments 1a and 1b. (We ignore the early rounds of the experiment since subjects may need some experience to learn even a simple adaptive rule). We classify the behavior of a subject as consistent with numerical hill climbing (myopic adjustment) if at least 33 (i.e., two–thirds) of his or her last 50 choices are consistent with this adaptive rule as defined by the conditions described above. The results of our tests of hypotheses H3 and H4 for Experiments 1a and 1b are reported in Tables 3.11 and 3.12, respectively. For each of the subjects in these experiments, these tables list the number of choices that are consistent with numerical hill climbing (top panel) and the number of choices consistent with myopic adjustment (bottom panel).

[Tables 3.11 and 3.12 About Here]

As we can see, significantly more subjects in the LWYE Experiment 1b acted in a manner consistent with hypotheses H3 and H4 than did in Experiment 1a. More precisely, in the LBYE Experiment 1a only 4 (4) subjects had at least two–thirds of their last 50 choices consistent with numerical hill climbing (myopic adjustment), while this was true for 9 (11) subjects in the LWYE Experiment 1b. In light of these findings it appears to us that since many subjects in Experiment 1b seemed to have adaptive rules (rather than payoff functions or best arms) as the objects of their learning, they did not pay sufficient attention to the data they were generating during the experiment. Such myopia could then explain their poor performance relative to their cohorts in Experiment 1a.

Finally, we have to account for the possibility that subjects in our one–person experiments abdicated learning altogether by choosing the same decision number throughout the entire experiment (following Merlo and Schotter (1992) we call these people "theorists"). While none of the 23 subjects in Experiment 1a could be classified as a theorist, 2 of the 24 subjects in Experiment 1b (Subjects 5 and 22) chose the same decision number (0 and 50, respectively) in all 76 rounds of the experiment.

3.1.3: Summary of the Results for One–Person Experiments

As we indicated before, hypotheses H1 and H2 are behaviorally distinct from hypotheses H3 and H4. While the former assumes that subjects use the data generated during the experiment to solve a static optimization problem in the final round, the latter assumes that they concentrate their attention on dynamic adaptive rules. Therefore, to summarize our results we combine our four hypotheses into two groups (H1 + H2 versus H3 + H4).

Our results indicate that a vast majority of subjects in the LBYE environment attempted to learn the payoff function they faced and/or their historically best arm, while a considerable number of subjects in the LWYE environment seemed to focus their learning on adaptive rules. To illustrate, we present Figure 3.2 which depicts Venn diagrams indicating how our subjects in Experiments 1a and 1b were classified according to our criteria.

[Figure 3.2 About Here]

As we can see from this figure, while hypotheses H1 and H2 jointly explain the behavior of 21 of the 23 subjects in Experiment 1a (91%), those same two hypotheses jointly explain the behavior of only 11 of the 24 subjects in Experiment 1b (46%). By looking at the marginal frequencies for hypotheses H3 and H4 the situation is just the opposite. While these two hypotheses jointly explain the behavior of only 6 of the 23 subjects in the LBYE environment (26%), they characterize the learning of 13 of the 24 subjects in the LWYE environment (54%). Eliminating those subjects whose behavior can be simultaneously explained by both groups of hypotheses, the difference between the two environments remains unchanged. While the behavior of 65% (25%) of the subjects in Experiment 1a (1b) can be characterized by hypotheses H1 and H2 alone, hypotheses H3 and H4 account for the behavior of 0% (33%) of the subjects in this experiment. A chi–square test rejects the null hypothesis that our classifications of subject behavior in Experiments 1a and 1b are independent of the payoff environmnet at the 5% significance level (P–value 0.001).

Finally, note that our analysis allowed us to characterize the learning of about 90% of the subjects in our one–person experiments. In particular, only 2 (3) subjects in Experiment 1a (1b) displayed behavior that was not consistent with any of our hypotheses.[12]

3.2: Two–Person Experiments

3.2.1: Overview of the Experimental Results

As we did for our one–person experiments we begin by comparing the learning of subjects in our two payoff environments by looking at their choices in the 75th round of Experiment 2a (the LBYE experiment) and the extra round of Experiment 2b (the LWYE experiment). While in our one–person experiments the payoff maximizing choice for the decision problem is a natural benchmark against which to compare subject behavior, in a game–theoretic setting no such natural benchmark exists. The problem is that in a game a Nash equilibrium choice is only optimal if one expects one's opponent to play accordingly. Hence, subject behavior can not be judged in reference to some objective standard (i.e., the distance from 37) but rather only in relation to a subject's conjecture about his or her opponent, which is unobservable. It is still possible, however, to ask if

[12] However, 3 of these 5 subjects (Subject 6 in Experiment 1a and Subjects 9 and 18 in Experiment 1b) had a score for our test of hypothesis H4 that was only one unit away from the threshold of 33 that we specified. Allowing a modification of our test to account for the behavior of these subjects would strenghten our conclusions even further.

our subjects, despite these conjectural difficulties, acted in the last (extra) round of Experiments 2a and 2b as if they predicted their opponent's choice correctly and chose a theoretical best response.

Tables 3.13 and 3.14 present the last round and extra round choices of subjects in Experiments 2a and 2b, respectively. Table 3.15 presents the last period choices of 36 random players in our simulated Experiment 2c.

[Tables 3.13, 3.14, and 3.15 About Here]

As we can see from these tables, there appear to be unambiguous differences between the behavior of random players in Experiment 2c and that of subjects in either Experiment 2a or 2b, though no obvious differences, on average, between Experiments 2a and 2b.[13] In terms of a money metric, however, notice that 10 out of 36 subjects in the LBYE Experiment 2a versus only 4 out of 34 subjects in the LWYE Experiment 2b had payoff losses (measured by the difference between the expected payoff corresponding to their last or extra round choice and the expected payoff associated with the theoretical best response to the final choice of their opponent) of less than $0.10.[14] The relatively superior performance of subjects in the LBYE environment is also illustrated by the fact that the choices of our subjects in Experiment 2a led to a mean expected payoff for the experiment of $14.32 which was $0.50 higher than the mean expected payoff of subjects in the extra round of Experiment 1b ($13.57). Finally, as it was true for our one–person experiments, such differences are not justified by differences in the sampling strategies adopted by

[13] These observations are supported by Kolmogorov–Smirnov tests using the data reported in Tables 3.13–3.15.

[14] The justification for the money metric we use here is given in footnote 6.

players in the two payoff environments.[15] An explanation for these differences needs to be found elsewhere.

3.2.2: Learning Hypotheses for Two–Person Games

In our two–person game experiments, the maximization problem faced by subjects in Experiments 1a and 1b is augmented by a conjectural component induced by the presence of strategic uncertainty. As a result, the learning task of subjects is greatly complicated, since any conjecture about one's opponent's choice identifies a different conditional payoff function like the one depicted in Figure 3.1 (where the choice of a computerized opponent was fixed at 37). The locus of maxima of such conditional payoff functions defines the best response function for a subject.

Faced with such a problem, subjects in Experiments 2a and 2b might choose to inductively learn to best respond to their opponent. However, given the 75–round horizon of the experiment, it is not feasible for subjects to estimate their best response function. This is so because estimation of the parameters of such a function would require many observations conditional on each choice of one's opponent. The size of the strategy space of the game in comparison to the number of data points generated during an experiment, makes such estimation impossible. What may be feasible is for subjects to coarsen the strategy space of the game and use the data generated to estimate their best response

[15] P–values for Kolmogorov–Smirnov tests of equality of the distributions of descriptive statistics of the sampling distributions of our subjects in Experiments 2a, 2b, and 2c over the 74 or 75 rounds of their history prior to making their last choice for big stakes (not reported here) are:

	Mean Choice	Standard Deviation	Median Choice	Interquantile Range
2a vs. 2b	0.094	0.001*	0.474	0.509
2a vs. 2c	0.000*	0.000*	0.000*	0.000*
2b vs. 2c	0.000*	0.000*	0.000*	0.000*

*indicates that we can reject the null hypothesis of equality of the distributions at the 5% significance level.

function for the simplified game. More formally, suppose that both players partition their action space into a small number, m, of contiguous intervals. This implicitly defines a m×m bi–matrix game which we call the "Convention Game" since the strategies can be thought of as crude conventions of behavior.[16] The elements of the payoff matrix of such a game can be estimated by taking averages of the payoffs associated with the decision choices defining each cell realized over the history of the play of the game.

Since subjects observe only their own payoffs but not their opponent's payoffs this procedure may not be informative in relation to the conjectural part of their decision problem. Subjects can, however, infer the distribution of their opponent's choices over this simplified strategy space during the course of the experiment and use this piece of information together with the estimated payoff matrix to identify a best response to such a distribution. This is, in a static sense, analogous to treating the problem as fictitious play (Brown (1951) and Robinson (1951)). These considerations yield the first of four hypotheses presented to characterize learning in our two–person game experiments. As in Section 3.1.2 these hypotheses are ranked in descending order of complexity.

H1': *Subjects in the last (extra) round of Experiment 2a (2b) use the data they have generated during the course of the experiment to estimate their payoffs in the Convention Game and choose a decision number which is a best response to the distribution of choices of their opponent.*

To operationalize this hypothesis we partition each subject's strategy space into four intervals ([0,25], [26,50], [51,75], and [76,100]) and use the observations generated by each pair of subjects during the course of the experiment to estimate the payoff matrix defining their Convention Game. After eliminating strategies that were never played during the course of the experiment (i.e., deleting rows and columns of the payoff matrix that

---

[16] The Nash equilibria of the Convention Game could be interpreted as conventions of behavior for our subjects to adhere to during the play of the game (see, e.g., Lewis (1969) and Schotter (1981)).

contain no observations), the estimation procedure we use assigns local payoff averages to all the cells containing a positive number of observations. However, some cells in the Convention Game may be empty. For instance, it is possible that a player chose decision numbers in the [0,25] interval during the course of the experiment and his or her opponent chose decision numbers in the [51,75] interval, but they never did so at the same time. Hence, the cell of the payoff matrix identified by these two strategies (one for the row player and one for the column player) contains no observations although it is part of the Convention Game. To estimate the expected payoff associated with such a cell we exploit two features of the game. First, given M, m, and c(e) the expected payoff in any cell is solely determined by the probability of winning M. Second, given the independence of the uniformly distributed additive shocks, the probability of winning only depends on the differences in subjects' decision numbers. Hence, the probability of winning along any diagonal in the Convention Game payoff matrix is constant since along such diagonals the distance between subjects' decision numbers is also constant. We then estimate the expected payoff in any empty cell by estimating the probability of winning there using observations in other cells along the same diagonal.

We test hypothesis H1' by asking if subjects in Experiments 2a and 2b acted as if they used the observations generated by their and their opponents' first 74 (75) round decisions to estimate the payoff matrix defining their Convention Game and then chose a decision number in the last (extra) round of the experiment within that interval of decision numbers that was a best response against the distribution of choices used by their opponent during the course of the experiment. We also check to see if the choices of pairs of subjects are Nash equilibrium choices for their Convention Games. The results of the test of hypothesis H1' for Experiments 2a and 2b are contained in Tables 3.16 and 3.17, respectively. For each pair of subjects in these experiments, these tables report their estimated payoff matrix along with the average payoff associated with each strategy weighted by the opponent's historical distribution of choices and their choices in the last (extra) round of the experiment.

[Tables 3.16 and 3.17 About Here]

As we can see from these tables, 20 (16) out of 36 (34) subjects in the LBYE (LWYE) Experiment 2a (2b) selected decision numbers in the last (extra) round which were consistent with best responses to the distribution of choices of their opponent during the course of the experiment and 4 (2) of these subjects chose a Nash equilibrium to the Convention Game. Such a difference in the behavior of subjects in the two payoff environments is not as stark as the one observed in our one–person experiments. This may be a consequence of the fact that our partitioning of the strategy space of the game was too coarse.

In the face of the difficulties involved in dealing with strategic uncertainty, human subjects might decide to simplify the decision problem they face even further by ignoring its game–theoretic nature and transforming it into a multi–armed bandit problem. As in our one–person experiment subjects would then attempt to learn which one of a limited subset of decision numbers returns the highest expected payoff. This yields our second hypothesis:

H2': *Subjects in the last (extra) round of Experiment 2a (2b) use the data they have generated during the course of the experiment to estimate the expected payoffs associated with a subset of decision numbers and choose that decision number from amongst this subset which yields the highest expected return.*

Our procedure for testing this hypothesis is identical to the one we used to test hypothesis H2 and the results for Experiments 2a and 2b are reported in Tables 3.18 and 3.19, respectively. For each of the subjects in these experiments, Tables 3.18 and 3.19

report those decision numbers that were chosen enough times (at lest 5) to be classified as arms along with the number of times they were chosen during the experiment, their realized mean payoff and standard deviation, the subject's final choice, and the number of times that that decision number was chosen during the experiment.

[Tables 3.18 and 3.19 About Here]

As we can see, hypothesis H2' accounts for the behavior of 15 of the 36 subjects in the LBYE Experiment 2a but only 2 of the 34 subjects in the LWYE Experiment 2b. This result represents the most noticeable difference between the behavior of subjects in Experiments 2a and 2b.

As in one–person decision problems, it is possible that the object of learning for our subjects is an adaptive rule which they use to react to the feedback they receive during the experiment. The two adaptive rules we consider here are numerical hill climbing, as defined by conditions (i)–(vi) above, and a myopic adjustment rule similar to the one specified for one–person decision problems but appropriately modified to take strategic uncertainty into account. Under what we call a strategic myopic adjustment rule, subjects again respond to whether or not they won the large fixed payment (M) in the previous period, but, because of the game–theoretic nature of the experiment, they also take into account the anticipated action of their opponent. Our strategic myopic adjustment rule posits mutually consistent behavior for both subjects in a game and incorporates the fact that the decision cost schedule is convex. Formally, it can be defined as follows:

(i'') If $\pi_{t-1}(e_{t-1}) = M - c(e_{t-1})$ and $e_{t-1} \leq 50$, then $e_t \geq e_{t-1}$
(ii'') If $\pi_{t-1}(e_{t-1}) = m - c(e_{t-1})$ and $e_{t-1} \leq 50$, then $e_t > e_{t-1}$
(iii'') If $\pi_{t-1}(e_{t-1}) = M - c(e_{t-1})$ and $e_{t-1} > 50$, then $e_t \leq e_{t-1}$
(iv'') If $\pi_{t-1}(e_{t-1}) = m - c(e_{t-1})$ and $e_{t-1} > 50$, then $e_t < e_{t-1}$.

The logic behind this rule is similar to the one we provided to justify hypothesis H3 but it incorporates considerations about the behavior of a live opponent. In particular, our strategic myopic adjustment rule takes into account the fact that a win for one player corresponds to a loss for that player's opponent and anticipation of the opponent's reaction induces subjects to respond differently to winning (losing) the large fixed payment M depending upon the cost of their decision. For example, consider the case where a subject wins at a low (low cost) decision number (case (i")). In this case our strategic myopic adjustment rule prescribes that she either increases her choice next period or keeps it constant for the following reason. Either her opponent lost when choosing a low or a high cost decision number. In the former situation he can be expected to increase his decision number next period since the marginal cost of doing so is small relative to the marginal increase in his probability of winning (see case (ii")). Hence, it is an optimal response for the winner to increase her decision number next period as well. In the latter situation the loser can be expected to lower his decision number next period in an effort to economize on decision costs (see case (iv")). In response the winner either keeps her decision number constant, thereby increasing her probability of winning at no additional cost, or increases her decision number and hence her probability of winning at a minimal cost. No incentives exist for the winner to lower her decision number next period since the cost saving is minimal relative to the impact on her proibability of winning. A similar explanation can be given for cases (ii")–(iv").

These considerations lead us to formulate our third and fourth hypotheses:

H3': *Over the course of Experiments 2a and 2b subjects' behavior is consistent with numerical hill climbing as defined by rules (i)–(vi) above.*

H4': *Over the course of Experiments 2a and 2b subjects' behavior is consistent with the strategic myopic adjustment rule defined by conditions (i")–(iv") above.*

Note that although requiring a conjecture about one's opponent's behavior, the learning rule implied by hypothesis H4' is still our simplest according to our definition of complexity.

We test hypotheses H3' and H4' by looking at the behavior of our subjects over the last 50 rounds of Experiments 2a and 2b. We classify the behavior of a subject as consistent with numerical hill climbing (strategic myopic adjustment) if at least 33 (i.e., two–thirds) of his or her last 50 choices are consistent with this adaptive rule as defined by the conditions described above. The results of our tests of hypotheses H3' and H4' for Experiments 2a and 2b are reported in Tables 3.20 and 3.21, respectively. For each of the subjects in these experiments, these tables list the number of choices that are consistent with numerical hill climbing (top panel) and the number of choices consistent with strategic myopic adjustment (bottom panel).

[Tables 3.20 and 3.21 About Here]

As we can see, in the LBYE Experiment 2a 12 (9) subjects had at least two–thirds of their last 50 choices consistent with numerical hill climbing (strategic myopic adjustment). These figures are comparable to the ones for the LWYE Experiment 2b, where the behavior of 11 (14) subjects seemed to be consistent with numerical hill climbing (strategic myopic behavior).

Finally, 1 subject in Experiment 2a (Subject 36) and 2 subjects in Experiment 2b (Subjects 25 and 33) abdicated learning altogether by choosing the same decision number (0) throughout the entire experiment.

3.2.2: Summary of the Results for Two–Person Game Experiments

To summarize the results for our two–person game experiments, we present Figure 3.3 which depicts Venn diagrams indicating how the learning of our subjects in Experiments 2a and 2b was classified according to our hypotheses. As we did in our one–person experiments, for purposes of comparison we combined hypotheses H1' and H2' as well as H3' and H4'.

[Figure 3.3 About Here]

As we can see from this figure, although the differences are less clear–cut than in our one–person experiments, it is still true that subjects in the LBYE environment tended to focus their learning relatively more on their best response function and/or their historically best arm than on myopic–adaptive rules, while their counterparts in the LWYE environment seemed to do the opposite. While hypotheses H1' and H2' jointly explain the behavior of 24 of the 36 subjects in Experiment 2a (69%), those same two hypotheses jointly explain the behavior of only 14 of the 34 subjects in Experiment 2b (47%). Also, while hypotheses H3' and H4' jointly explain the behavior of 17 of the 36 subjects in the LBYE Experiment 2a (47%), they characterize the learning of 20 of the 34 subjects in the LWYE Experiment 2b (59%). Eliminating those subjects whose behavior can be simultaneously explained by more than one group of hypotheses, the difference between the two environments becomes slightly more clear. While the behavior of 39% (24%) of the subjects in Experiment 2a (2b) can be characterized by hypotheses H1' and H2' in isolation, hypotheses H3' and H4' account for the behavior of 17% (35%) of the subjects, respectively. However, using a chi–square test we can not reject the null hypothesis that our classifications of subject behavior in Experiments 2a and 2b are independent of the payoff environmnet at conventional significance levels (P–value 0.15).

Finally, note that as with our one–person experiments our analysis allowed us to characterize the learning of about 90% of the subjects in our two–person experiments. In particular, only 4 (4) subjects in Experiment 2a (2b) displayed behavior that was not consistent with any of our hypotheses.[17]

3.3: Complexity and Learning: A Comparison of One and Two–Person Experiments

To conclude the presentation of our results, we compare our one–person and two–person experiments by looking at Figures 3.2 and 3.3. As we can see, the most apparent phenomenon we observe is that as we move from a one–person decision problem to a game (i.e., as we increase the complexity of the decision problem our subjects face) while keeping the payoff environment constant, the distribution of subjects across learning hypotheses shifts towards simpler rules. Such a phenomenon is more pronounced in our LBYE environment where hypotheses H3 and H4 jointly explain the behavior of only 26% of the subjects in Experiment 1a while hypotheses H3' and H4' explain the behavior of 47% of the subjects in Experiment 2a. When we restrict our attention to those subjects whose learning can be exclusively characterized by hypotheses H3 and H4 or H3' and H4' these percentages become 0% and 17%, respectively. A chi–square test of independence between our classifications of subject behavior in Experiments 1a and 2a rejects the null at the 5% significance level (P–value 0.05). The extent to which such a shift occurs in the LWYE environment is ambiguous since the proportion of subjects whose learning can be characterized by our two simplest rules (hypotheses H3 and H4 or H3' and H4') goes from 54% in Experiment 1b to only 59% in Experiment 2b (33% to 35% after eliminating those subjects whose behavior can also be explained by our more complex hypotheses). Furthermore, a chi–square test of independence between our classifications of subject

[17] Note that 2 of these 8 subjects (Subject 20 in Experiment 2a and Subject 24 in Experiment 2b) had a score for our test of hypothesis H4' that was only one unit away from the threshold of 33 we specified and a modification of our test to account for the behavior of these subjects would leave our conclusions substantially unchanged.

behavior in Experiments 1b and 2b can not reject the null at conventional significance levels (P–value 0.97).

These results provide some evidence in support of the view that at least in a payoff environment that fosters learning (LBYE), the frequency with which human subjects tend to use relatively simpler learning rules increases with the complexity of the decision task they face.

Section 4: Conclusion

While it is never wise to generalize on the basis of a small number of experimental results, there are a number of lessons that we can learn from our experiments if they hold up to replication elsewhere. To begin, our results add yet another piece of evidence for the growing view that learning is a situation and institution specific phenomenon (see, e.g., Mookerjee and Sopher (1993)). This view has been recently summarized by Milgrom and Roberts (1991) as follows:

> "Taken together, these results [i.e., earlier theoretical results on learning in games, *cfr.*] raise serious doubts about the validity of Nash equilibrium and its refinements as a general model of the likely outcomes of adaptive learning. More fundamentally, they indicate that the 'rationality' of any particular learning algorithm is situation dependent: An algorithm that performs well in some situations may work poorly in others. Apparently, real biological players tailor rules of thumb to their environment and experience: They learn how to learn. Thus, any single, simple specification of a learning algorithm is unlikely to represent well the behavior that actual players would adopt." (p. 84).

Such a view, that learning is situation or institution dependent, creates a problem for economic theory in its quest to present generalizable models of economic behavior. If learning is situation dependent, in the sense that people attempt to learn about different things when placed in different situations, then this raises the possibility that one would have to construct special learning theories for each and every economic institution –– certainly a dismal prospect. This opens the door for experimentalists, however, since if they could classify institutions into equivalence classes across which human learning behavior is similar, then theorists could attempt to characterize these institutions. If successful, a small class of learning theories might be constructed which would explain behavior in a large number of institutions. Our results take a step along this path.

One should not take these remarks as negative or pessimistic since we already consider our results as being supportive of much of the work undertaken by game theorists. For example, in our one–person Learn–Before–You–Earn environments, our subjects' behavior is consistent with the behavior predicted by Aghion et al. (1991) where agents in infinite horizon problems without discounting optimally learn the shape of the payoff function they are facing. Further, in our two–person Learn–While–You–Earn environment, which most closely mimic the non–cooperative games studied, for example, by Jordan (1991), Kalai and Lehrer (1993), and Nyarko (1992), we find that laboratory subjects adhere to adaptive learning strategies with limited memory. While we see no evidence that our subjects perform the type of calculations called for by the modern game–theoretic literature, they do appear to mimic the adaptive pattern described in Milgrom and Roberts (1991), where adaptive learning is shown to be capable of reaching Nash equilibrium outcomes in certain classes of games. Whether our subjects' behavior would eventually converge to Nash equilibrium or first–best optimal behavior can not be answered by our experiments since those convergence results are asymptotic and our experiments were obviously run with only finite (75 round) time horizons.

Our results also have direct bearing on the methodology of experimental economics. This is true because almost all experiments in economics aim to test static theories using a repeated framework. This is typically justified by the claim that doing an experiment once and only once does not allow subjects to learn. Hence, repetition is recommended to foster learning. In most designs, subjects play games repeatedly and earn payoffs each period –– they play in a Learn–While–You–Earn environment. In performing statistical tests on the data generated by these experiments, experimentalists typically use observations collected at the end of the experiment since those supposedly distill all the information learned during the course of the experiment. Our results, however, indicate that it is exactly in these types of Learn—While—You—Earn environments that learning is most problematic and in which subjects seem to pay the least amount of attention to the data generated during their multi—period play. Hence, our results seem to lead towards an advocacy of Learn—Before—You—Earn environments in the laboratory at least for one—person decision problems. For game—theoretic problems the prescription is less clear since the differences in the learning of subjects across our two payoff environments are less clear.

Finally, our results teach us that in payoff environments which facilitate learning, the complexity of the learning rules adopted by subjects is inversely related to the complexity of the decision task they face. This result is consistent with the theoretical work of Heiner (1983), who argues that as the complexity of the decision problems faced by economic agents increases, more complex rules become dysfunctional. Simon's (1976, 1978) view of human problem solving and procedural rationality is also borne out in our experimental data.

References

Aghion, P., P. Bolton, C. Harris, and B. Jullien (1991), "Optimal Learning by Experimentation," Review of Economic Studies, **58**, 621–654.

Arthur, W.B. (1990), "A Learning Algorithm that Mimics Human Learning," Santa Fe Institute, Economic Research Program, Working Paper No. 90–026.

Bull, C., A. Schotter, and K. Weigelt (1987), "Tournaments and Piece Rates: An Experimental Study", Journal of Politcal Economy, **95**, 1–33.

Bray, M. (1982), "Learning, Estimation, and the Stability of Rational Expectations," Journal of Economic Theory, **26**, 318–339.

Brown, G. (1951), "Iterated Solution of Games by Fictitious Play," in Activity Analysis of Production and Allocation, New York: Wiley.

Cheung, Y.W. and D. Friedman (1994), "Learning in Evolutionary Games: Some Laboratory Results," mimeo.

Crawford, V. (1992), "Adaptive Dynamics in Coordination Games," University of California, San Diego, Discussion Paper No. 92–02.

Easley, D. and N. Kiefer (1988), "Controlling a Stochastic Process with Unknown Parameters," Econometrica, **56**, 1045–1064.

El–Gamal, M.A. and D.M. Grether (1994), "Uncovering Behavioral Strategies: Likelihood–Based Experimental Data–Mining," mimeo.

El–Gamal, M.A., R.D. McKelvey and T.R. Palfrey (1992a), "Learning in Experimental Games," mimeo.

El–Gamal, M.A., R.D. McKelvey and T.R. Palfrey (1992b), "A Bayesian Sequential Experimental Study of Learning in Games," Journal of the American Statistical Association, **88**, 428–435.

Futia, C. (1977), "The Complexity of Economic Decision Rules," Journal of Mathematical Economics, **4**, 289–299.

Goldberger, A.S. (1991), A Course in Econometrics, Cambridge: Harvard University Press.

Heiner, R.A. (1983), "The Origin of Predictible Behavior," American Economic Review, **73**, 560–595.

Holland, J. (1975), Adaptation in Natural and Artificial Systems, Ann Arbor: University of Michigan Press.

Jordan, J.S. (1991), "Bayesian Learning in Normal Form Games," Games and Economic Behavior, **3**, 60–81.

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

Kalai, E. and W. Stanford (1988), "Finite Rationality and Interpersonal Complexity in Repeated Games," Econometrica, **56**, 397–410.

Lazear, E.P. and S. Rosen (1981), "Rank Order Tournaments as Optimal Labor Contracts," Journal of Political Economy, **89**, 841–864.

Lewis, D. (1969), Convention: A Philosophical Study, Cambridge: Harvard University Press.

Marcet, A. and T.J. Sargent (1989), "Convergence of Lest Squares Learning Mechanisms in Self Referential Linear Stochastic Models," Journal of Economic Theory, **48**, 337–368.

Marimon, R. and E. McGrattan (1992), "On Adaptive Learning in Strategic Games," mimeo. Forthcoming in A. Kirman and M. Salmon (*eds.*), Learning and Rationality in Economics, Oxford: Basil Blackwell.

McLennan, A. (1987), "Incomplete Learning in a Repeated Statistical Decision Problem," mimeo.

Merlo, A. and A. Schotter (1992), "Theory and Misbehavior of First–Price Auctions: Comment," American Economic Review, **82**, 1413–1425.

Milgrom, P. and J. Roberts (1991), "Adaptive and Sophisticated Learning in Normal Form Games," Games and Economic Behavior, **3**, 82–100.

Mookherjee, D. and B. Sopher (1991), "Learning Behavior in an Experimental Matching Pennies Game," mimeo. Forthcoming in Games and Economic Behavior.

Nyarko, Y. (1992), "Bayesian Learning without Common Priors and Convergence to Nash Equilibrium," New York University, C.V. Starr Center, Research Report No. 92–25.

Robinson, J. (1951), "An Iterative Method of Solving a Game," Annals of Mathematics, **54**, 296–301.

Roth, A.E. and I. Erev (1993), "Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," mimeo.

Schotter, A. (1981), The Economic Theory of Social Institutions, Cambridge: Cambridge University Press.

Schotter, A. and K. Weigelt (1992), "Asymmetric Tournaments, Equal Opportunity Laws and Affirmative Action: Some Experimental Results", Quarterly Journal of Economics, **106**, 513–539.

Simon, H.A. (1976), "From Substantive to Procedural Rationality," in S.J. Latsis (*ed.*), Method and Appraisal in Economics, Cambridge: Cambridge University Press.

Simon, H.A. (1978), "Rationality as Process and as Product of Thought," American Economc Review, **68**, 1–16.

Figure 3.1: Expected Payoff Function

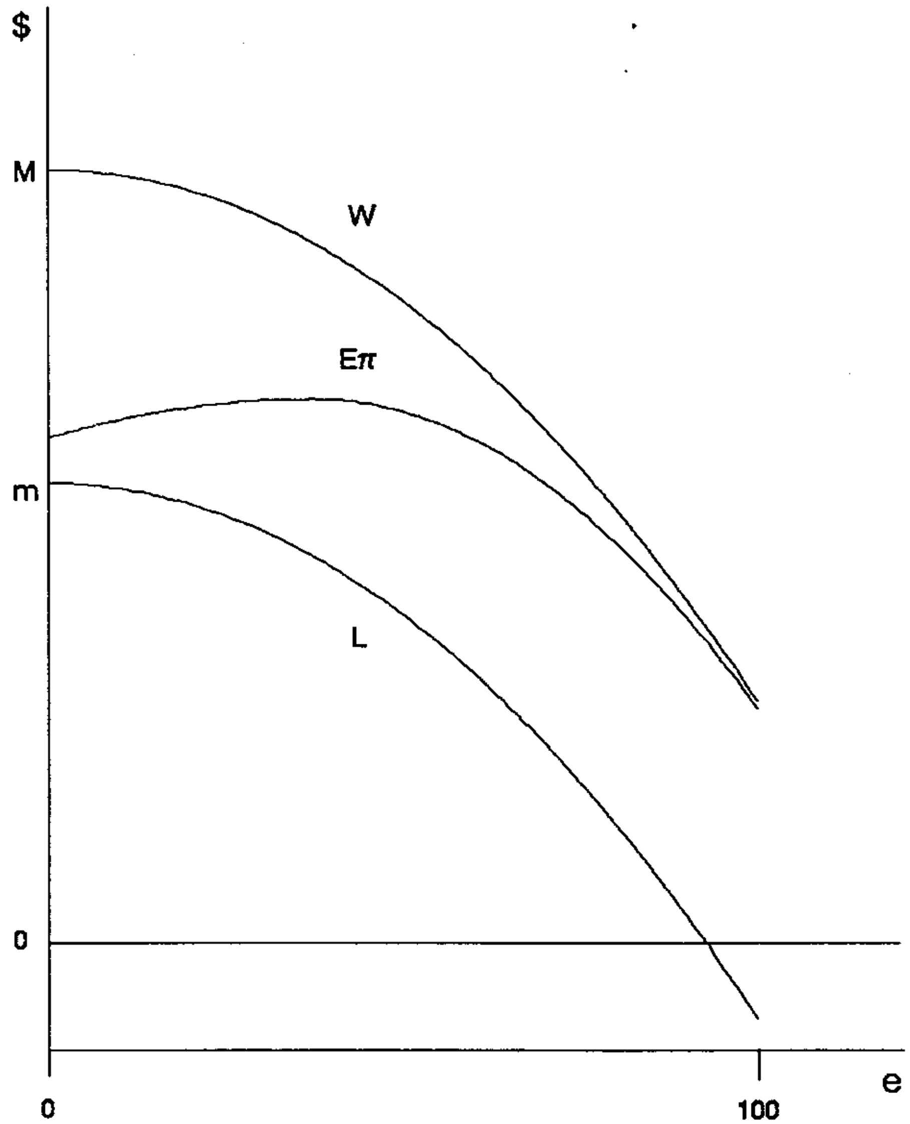

Figure 3.2: Venn Diagrams: Experiments 1a and 1b

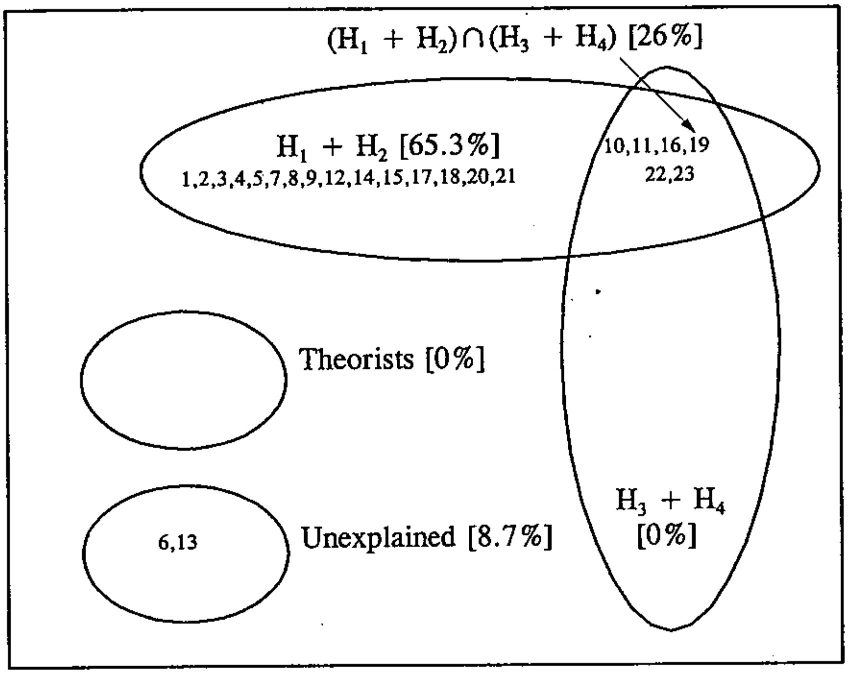

Experiment 1a

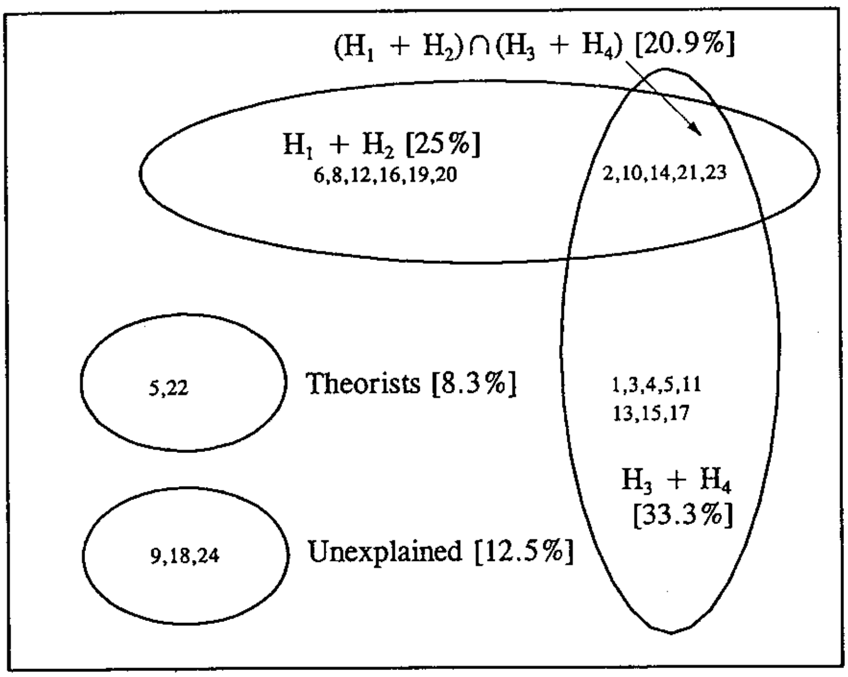

Experiment 1b

Figure 3.3: Venn Diagrams: Experiments 2a and 2b

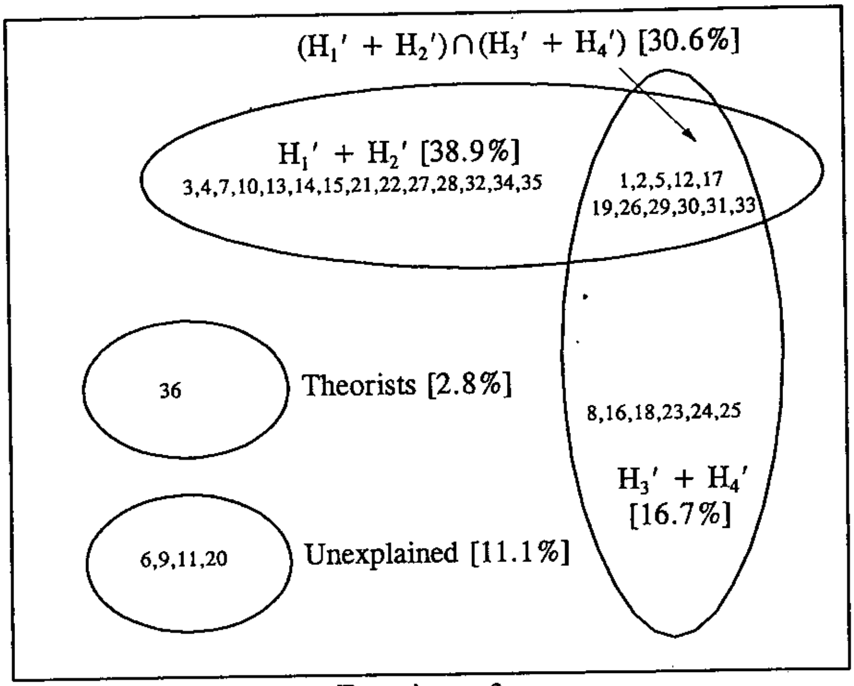

Experiment 2a

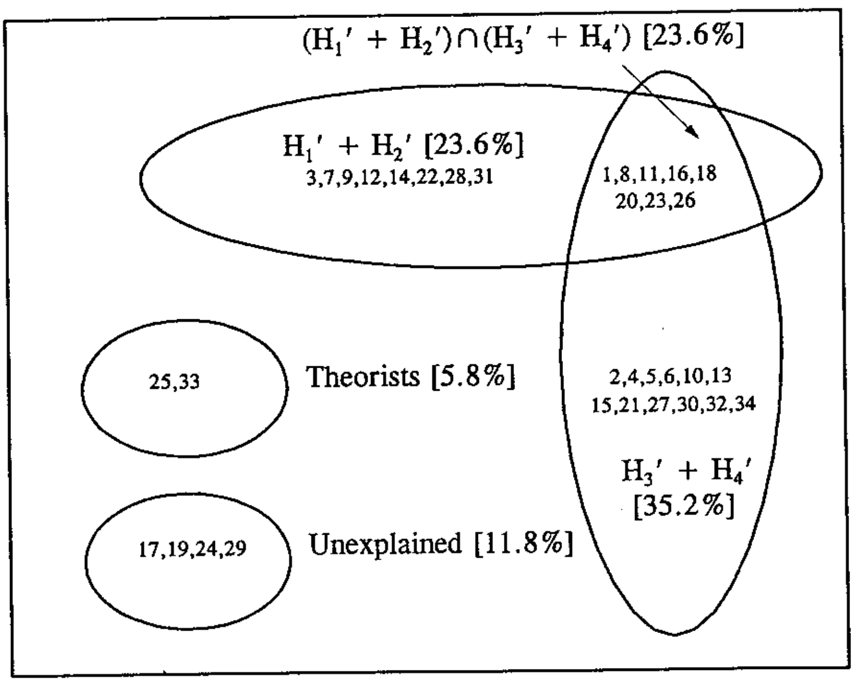

Experiment 2b

Table 2.1: Experimental Design

Experiment	Opponent	Opponent's Strategy	Payoff Environment	Number of Subjects
		One–Person Decision Problems		
1a	Computer	Fixed and Known to be 37	LBYE	23
1b	Computer	Fixed and Known to be 37	LWYE	24
1c	Simulation			24
		Two–Person Games		
2a	Human	N.A.	LBYE	36
2b	Human	N.A.	LWYE	34
2c	Simulation			36

Table 3.1: Last Period Choices — Experiment 1*a*

Subject	Last Period Choice	Expected Payoff	$\|(2)-37\|$	Payoff Loss
1*	40	15.25	3	0.02
2*	38	15.27	1	0.00
3	58	14.30	21	0.97
4	50	14.90	13	0.37
5*	39	15.26	2	0.01
6	96	7.62	59	7.65
7*	30	15.23	7	0.04
8*	33	15.26	4	0.01
9	23	15.12	14	0.15
10	50	14.90	13	0.37
11*	42	15.21	5	0.06
12	15	14.89	22	0.38
13	69	13.02	32	2.25
14*	36	15.27	1	0.00
15*	32	15.25	5	0.02
16	22	15.10	15	0.17
17*	42	15.21	5	0.06
18*	37	15.27	0	0.00
19	50	14.90	13	0.37
20	50	14.90	13	0.37
21*	38	15.27	1	0.00
22*	40	15.25	3	0.02
23	50	14.90	13	0.37
Average	42.61	14.68	11.52	0.59

* indicates that the payoff loss is smaller than $0.10.

Table 3.2: Extra Period Choices — Experiment 1*b*

Subject	Extra Period Choice	Expected Payoff	$\|(2)-37\|$	Payoff Loss
1	65	13.54	28	1.73
2	45	15.13	8	0.14
3	100	6.55	63	8.72
4	77	11.75	40	3.52
5	0	14.18	37	1.09
6	45	15.13	8	0.14
7*	41	15.23	4	0.04
8	68	13.15	31	2.12
9	70	12.87	33	2.40
10	65	13.54	28	1.73
11	45	15.13	8	0.14
12*	35	15.27	2	0.00
13	0	14.18	37	1.09
14	44	15.16	7	0.11
15	68	13.15	31	2.12
16	0	14.18	37	1.09
17	70	12.87	33	2.40
18	69	13.02	32	2.25
19	50	14.90	13	0.37
20	50	14.90	13	0.37
21	45	15.13	8	0.14
22	50	14.90	13	0.37
23*	30	15.23	7	0.04
24	100	6.55	63	8.72
Average	51.33	13.57	24.33	1.70

* indicates that the payoff loss is smaller than $0.10.

Table 3.3: Last Period Choices — Experiment 1*c*

Subject	Last Period Choice	Expected Payoff	$\|(2) - 37\|$	Payoff Loss
1	22	5.10	15	0.17
2	44	15.16	7	0.11
3*	43	15.19	6	0.08
4	1	14.24	36	1.03
5	20	15.04	17	0.23
6	63	13.78	26	1.49
7	90	9.10	53	6.17
8	93	8.38	56	6.89
9	72	12.57	35	2.70
10	51	14.84	14	0.43
11*	38	15.27	1	0.00
12*	41	15.23	4	0.04
13	52	14.77	15	0.50
14	67	13.29	30	1.98
15*	42	15.21	5	0.06
16	73	12.42	36	2.85
17	84	10.41	47	4.86
18	15	14.89	22	0.38
19	58	14.30	21	0.97
20*	29	15.22	8	0.05
21	11	14.73	26	0.54
22	87	9.77	50	5.50
23	45	15.13	8	0.14
24	84	10.41	47	4.86
Average	51.04	13.52	24.38	1.75

* indicates that the payoff loss is smaller than $0.10.

Table 3.4: Descriptive Statistics of Subjects' Choices — Experiment 1*a*

Subject	Mean Choice	Standard Deviation	Median Choice	Interquantile Range
1	45.84	9.93	42	7
2	38.53	23.02	38	29
3	56.52	6.55	57	4
4	45.60	32.92	47	60
5	37.25	11.24	39	7
6	52.40	30.22	52	54
7	35.92	13.76	30	10
8	32.15	17.70	30	17
9	31.40	31.37	29	42
10	54.65	29.25	52	44
11	32.87	18.28	29	24
12	46.44	24.75	40	35
13	64.63	21.00	67	31
14	41.93	28.04	36	38
15	34.51	15.53	32	6
16	28.05	11.52	25	8
17	42.25	17.48	42	17
18	45.24	16.79	40	14
19	48.93	11.24	50	10
20	48.95	17.89	50	23
21	32.03	15.31	38	8
22	41.56	20.42	40	10
23	45.99	10.79	47	16
Average	42.77	18.91	41.39	22.35

Table 3.5: Descriptive Statistics of Subjects' Choices — Experiment 1*b*

Subject	Mean Choice	Standard Deviation	Median Choice	Interquantile Range
1	57.00	11.03	60	15
2	47.73	5.68	50	5
3	73.79	19.37	78	37
4	61.93	27.47	67	44
5	0.00	0.00	0	0
6	44.11	4.74	45	4
7	9.00	19.47	1	4
8	64.72	14.38	65	20
9	48.12	11.33	45	15
10	63.93	6.97	65	5
11	26.99	27.80	9	40
12	53.17	28.22	50	46
13	16.75	34.39	0	0
14	51.86	17.58	45	14
15	59.92	33.05	76	50.5
16	23.47	15.94	35	35
17	61.12	5.32	60	7
18	76.97	17.50	77	22
19	27.49	25.64	35	50
20	39.95	7.57	40	7
21	38.89	11.95	45	2.5
22	50.00	0.00	50	0
23	45.61	21.09	44	21.5
24	69.74	13.05	70	15.5
Average	46.34	15.81	46.33	19.17

Table 3.6: Descriptive Statistics of Subjects' Choices — Experiment 1*c*

Subject	Mean Choice	Standard Deviation	Median Choice	Interquantile Range
1	48.15	29.46	49	50
2	49.99	27.75	53	52
3	54.53	28.94	51	51
4	47.47	32.16	38	64
5	47.32	25.68	44	44
6	44.72	27.82	46	48
7	51.08	29.36	52	54
8	54.61	26.72	54	45
9	45.77	26.89	46	45
10	48.03	27.46	49	43
11	50.81	31.59	52	57
12	50.72	28.37	53	49
13	47.75	28.15	47	50
14	41.47	27.72	37	42
15	49.96	29.67	46	51
16	49.01	27.48	47	43
17	46.79	27.83	45	43
18	52.08	28.29	52	44
19	44.89	26.01	43	39
20	54.88	30.16	61	54
21	46.37	29.04	46	59
22	52.39	29.00	53	48
23	49.71	27.74	45	45
24	53.80	30.68	59	55
Average	49.26	28.50	48.67	48.96

Table 3.9: Test of Hypothesis H2 — Experiment 1a(#)

Subject	Arms	# of Times Chosen	Average Payoff	Standard Deviation	Last Period Choice	# of Times Chosen
1	40	19	16.09	4.39	40	19
	41	5	13.92	4.85		
	42	11	13.47	4.46		
	43	5	17.21	3.96		
	44	5	17.08	3.96		
	other	30	14.42	4.45		
2*	10	10	14.52	3.74	38	16
	38	16	16.27	4.43		
	49	9	16.19	3.90		
	50	7	14.21	4.73		
	other	33	12.68	3.74		
3	55	9	15.25	3.90	58	11
	56	8	13.73	4.58		
	57	11	13.66	4.46		
	58	11	14.29	4.13		
	59	6	12.11	4.85		
	60	8	15.25	3.13		
	65	5	13.64	3.96		
	other	17	14.04	4.59		
4*	50	8	14.69	4.58	50	8
	100	9	6.75	0.00		
	other	58	14.21	3.76		
5	23	10	15.65	4.57	39	37
	32	19	16.49	4.49		
	39	37	15.89	4.40		
	45	6	15.77	4.57		
	other	3	8.67	5.64		
7*	20	6	12.30	0.00	30	29
	30	29	16.43	4.48		
	40	25	15.75	4.43		
	other	15	12.80	3.53		
8	10	6	17.18	4.85	33	1
	23	4	14.32	4.43		
	25	8	16.39	4.73		
	30	7	16.61	4.73		
	40	10	15.81	4.57		
	60	7	15.08	3.35		
	other	33	14.32	4.82		
9	0	10	15.56	4.28	23	1
	10	7	15.28	4.32		
	15	5	17.87	4.85		
	37	7	13.37	4.32		
	100	6	6.75	0.00		
	other	40	14.94	4.56		
10*	50	6	16.52	3.62	50	6
	other	69	12.37	4.31		
11*	20	9	17.22	4.67	42	7
	42	7	17.84	3.35		
	44	6	14.42	4.85		
	other	53	15.03	4.72		
12	37	8	17.48	4.10	15	3
	other	67	13.13	4.93		

Table 3.9 cont'd

Subject	Arms	# of Times Chosen	Average Payoff	Standard Deviation	Last Period Choice	# of Times Chosen
13	78	5	9.08	4.85	69	4
	other	70	12.97	4.58		
14*	36	11	18.20	3.58	36	11
	other	64	12.79	4.07		
15*	0	6	12.90	0.00	32	23
	30	5	15.09	4.85		
	31	10	16.77	4.57		
	32	23	17.52	4.16		
	33	5	14.81	4.85		
	37	16	14.72	4.54		
	50	5	14.46	4.85		
	77	5	12.86	0.00		
16*	20	5	14.07	3.96	22	12
	22	12	16.60	4.62		
	25	27	15.89	4.49		
	30	6	13.03	3.61		
	other	25	14.82	4.50		
17*	25	5	15.50	4.85	42	14
	42	14	17.21	3.77		
	44	5	13.54	4.85		
	other	51	14.78	4.07		
18*	37	17	14.49	4.49	37	17
	other	58	14.33	4.02		
19	40	16	16.04	4.43	50	29
	50	29	14.95	4.28		
	other	30	14.85	3.80		
20	37	14	15.90	4.55	50	16
	45	6	15.77	4.57		
	50	16	14.69	4.43		
	55	9	15.25	3.90		
	60	9	13.40	4.43		
	other	21	13.21	3.54		
21	2	9	13.88	2.95	38	33
	20	6	13.78	3.61		
	37	15	16.16	4.49		
	38	33	14.75	4.48		
	40	5	14.04	4.85		
	other	7	13.34	4.25		
22*	35	5	14.60	4.85	40	21
	40	21	17.24	3.86		
	45	13	15.31	4.49		
	50	10	15.35	4.28		
	other	26	13.33	2.92		
23	37	17	16.05	4.49	50	15
	47	10	15.78	4.28		
	50	15	15.05	4.32		
	53	10	14.88	4.28		
	other	23	15.00	3.59		

(#) subject 6 does not have any choice that can be classified as an arm.
* indicates that the last period choice is the arm associated with the highest average payoff.

Table 3.10: Test of Hypothesis H2 — Experiment 1b(#)

Subject	Arms	# of Times Chosen	Average Payoff	Standard Deviation	Extra Period Choice	# of Times Chosen
1	41	12	14.81	4.62	65	36
	60	12	13.40	4.36		
	65	36	13.20	3.89		
	other	15	13.90	3.83		
2*	45	32	15.40	4.35	45	32
	50	36	14.81	4.31		
	other	7	14.00	3.95		
4	22	5	15.71	4.85	77	5
	33	7	16.32	4.73		
	44	8	15.53	4.58		
	55	6	14.27	4.57		
	77	5	9.32	4.85		
	78	6	12.62	0.00		
	88	5	10.13	0.00		
	other	33	11.28	3.90		
5	0	75	15.50	4.06	0	75
6	39	8	17.26	4.10	45	23
	45	23	16.79	3.74		
	46	6	15.63	4.57		
	48	7	15.77	4.32		
	52	5	14.15	4.85		
	other	26	13.97	4.41		
7	0	6	14.38	3.62	41	0
	1	38	14.06	3.03		
	other	31	14.68	4.06		
8	48	5	14.75	4.85	68	4
	49	6	16.67	3.62		
	65	7	11.62	4.73		
	other	57	12.68	3.86		
9	37	15	13.79	4.32	70	4
	40	16	15.48	4.54		
	45	11	14.69	4.62		
	50	13	13.24	4.59		
	60	6	16.35	0.00		
	other	14	14.62	1.32		
10	55	6	15.74	3.62	65	46
	60	9	13.40	4.43		
	65	46	13.68	3.55		
	other	14	11.60	3.70		
11	2	6	14.37	3.62	45	11
	5	6	12.86	0.00		
	6	7	12.85	0.00		
	45	11	15.50	4.46		
	54	5	17.38	0.00		
	other	40	12.98	4.03		

Table 3.10 cont'd

Subject	Arms	# of Times Chosen	Average Payoff	Standard Deviation	Extra Period Choice	# of Times Chosen
13	0	57	13.83	2.74	0	57
	50	5	14.46	4.85		
	100	9	5.77	2.95		
	other	4	14.51	4.46		
14	41	6	14.81	4.85	44	9
	42	5	15.56	4.85		
	44	9	14.92	4.67		
	45	11	16.30	4.13		
	other	44	12.44	4.84		
15	78	16	12.07	2.21	68	0
	88	7	10.13	0.00		
	other	52	12.62	4.43		
16*	0	21	15.43	4.10	0	21
	35	43	15.18	4.47		
	other	11	12.58	2.10		
17	57	6	15.40	3.62	70	4
	58	8	15.60	3.13		
	59	11	13.31	4.46		
	60	14	13.19	4.40		
	68	6	11.87	4.57		
	other	30	13.18	4.17		
19*	0	26	14.60	3.56	50	29
	50	29	14.64	4.37		
	other	20	14.51	3.21		
20	38	14	16.43	4.40	50	0
	40	21	15.98	4.40		
	45	26	16.33	4.01		
	other	14	13.02	3.79		
21 @	10	8	16.07	4.58	45	53
	20	5	15.84	4.85		
	45	53	16.04	4.10		
	other	9	14.58	4.77		
22	50	75	15.17	4.16	50	75
23	0	5	14.67	3.96	30	3
	40	10	17.58	3.74		
	44	10	16.19	4.28		
	50	10	16.23	3.74		
	70	9	12.44	3.90		
	other	31	13.59	4.08		
24	79	5	12.39	0.00	100	0
	other	70	12.67	3.46		

(#) subjects 3, 12, and 18 do not have any choice that can be classified as an arm.
* indicates that the extra period choice is the arm associated with the highest average payoff.
@ indicates that the extra period choice is an arm and the (dollar) distance from the arm associated with the highest average payoff is smaller than $0.10.

Table 3.11: Test of Hypotheses H3 and H4 — Experiment 1*a*

Number of Choices Consistent with Numerical Hill Climbing (Last 50 Rounds)

Subject	1	2	3	4	5	6	7	8	9	10*	11*	12	13	14	15	16*
# of Choices	25	30	20	10	28	26	26	21	31	37	35	29	31	25	23	33
Subject	17	18	19*	20	21	22	23									**Average**
# of Choices	28	26	42	16	26	20	21									26.5

Number of Choices Consistent with Myopic Adjustment (Last 50 Rounds)

Subject	1	2	3	4	5	6	7	8	9	10*	11	12	13	14	15	16
# of Choices	27	27	31	22	27	32	28	31	30	33	29	26	28	32	27	29
Subject	17	18	19*	20	21	22*	23*									**Average**
# of Choices	28	25	41	29	25	35	33									29.3

* indicates that at least 2/3 of the choices in the last 50 rounds are consistent with the adaptive rule.

Table 3.12: Test of Hypotheses H3 and H4 — Experiment 1*b*

Number of Choices Consistent with Numerical Hill Climbing (Last 50 Rounds)

Subject	1*	2	3*	4*	5	6	7*	8	9	10*	11	12	13*	14*	15	16
# of Choices	33	31	42	37	0	31	36	32	26	36	29	30	49	38	27	30
Subject	17*	18	19	20	21	22	23*	24								**Average**
# of Choices	35	30	23	22	17	0	33	25								28.8

Number of Choices Consistent with Myopic Adjustment (Last 50 Rounds)

Subject	1*	2*	3*	4*	5	6	7	8	9	10*	11*	12	13	14*	15*	16
# of Choices	37	37	45	38	18	32	22	25	31	44	35	26	7	41	36	31
Subject	17*	18	19	20	21*	22	23*	24								**Average**
# of Choices	39	31	29	32	35	16	39	22								32.0

* indicates that at least 2/3 of the choices in the last 50 rounds are consistent with the adaptive rule.

Table 3.13: Last Period Choices — Experiment 2*a*

Subject	Last Period Choice	Opponent's Last Period Choice	Expected Payoff	Best Response	Expected Payoff at BR	$\|(2)-(5)\|$	Payoff Loss
1*	20	54	13.76	22	13.77	2	0.01
2	0	30	14.63	35	16.02	35	1.39
3*	30	15	17.48	30	17.48	0	0.00
4*	30	0	18.68	25	18.72	5	0.04
5	66	54	12.02	22	13.77	44	1.75
6	100	70	5.03	9	13.03	91	8.00
7	0	12	16.10	29	17.75	29	1.65
8	70	100	7.28	0	12.90	70	5.62
9	15	30	15.48	35	16.02	20	0.54
10	54	66	11.72	12	13.16	42	1.44
11	54	20	15.92	32	17.02	22	1.10
12	12	0	18.34	25	18.72	13	0.38
13	0	11	16.19	29	17.83	29	1.64
14	11	0	18.28	25	18.72	14	0.44
15	23	30	15.79	35	16.02	12	0.23
16	69	72	9.86	7	12.98	62	3.12
17*	30	23	16.72	32	16.73	2	0.01
18	72	69	9.88	9	13.06	63	3.18
19	50	0	17.38	25	18.72	25	1.34
20	90	25	9.44	33	16.53	57	7.09
21*	25	45	14.45	30	14.47	5	0.02
22	45	25	16.22	33	16.53	12	0.31
23*	32	25	16.53	33	16.53	1	0.00
24	0	50	13.52	26	14.06	26	0.54
25	96	0	7.93	25	18.72	71	10.79
26	25	32	15.65	35	15.81	10	0.16
27	25	90	12.12	0	12.90	25	0.78
28*	0	96	12.90	0	12.90	0	0.00
29*	37	50	13.95	26	14.05	11	0.10
30	50	31	15.43	35	15.92	15	0.49
31	0	0	17.33	25	18.72	25	1.39
32	0	0	17.33	25	18.72	25	1.39
33*	31	50	14.03	26	14.06	5	0.03
34*	31	0	18.65	25	18.72	6	0.07
35	50	37	14.90	37	15.27	13	0.37
36	0	31	14.56	35	15.92	35	1.36
Average	34.53	34.53	14.32	24.64	15.90	25.61	1.58

* indicates that the payoff loss is smaller than $0.10.

Table 3.14: Extra Period Choices — Experiment 2*b*							
Subject	Extra Period Choice	Opponent's Extra Period Choice	Expected Payoff	Best Response	Expected Payoff at BR	|(2) – (5)|	Payoff Loss
1	40	0	18.25	25	18.72	15	0.47
2	0	40	14.01	34	14.96	34	0.95
3*	30	10	17.91	28	17.92	2	0.01
4	67	99	7.76	0	12.90	67	5.14
5	63	37	13.78	37	15.27	26	1.49
6	10	30	15.24	35	16.02	25	0.78
7	37	63	12.86	15	13.27	22	0.41
8	50	0	17.38	25	18.72	25	1.34
9	11	41	14.45	33	14.85	22	0.40
10	99	67	5.45	11	13.12	88	7.67
11	0	50	13.52	26	14.06	26	0.54
12	41	11	17.50	29	17.83	12	0.33
13*	27	45	14.47	30	14.47	3	0.00
14	23	33	15.50	36	15.71	13	0.21
15	54	50	13.38	26	14.06	28	0.68
16	54	0	16.91	25	18.72	29	1.81
17	45	27	16.06	34	16.34	11	0.28
18	50	56	12.94	21	13.64	29	0.70
19	50	54	13.14	22	13.77	28	0.63
20	55	65	11.75	13	13.19	42	1.44
21	65	55	12.02	21	13.70	44	1.68
22	0	54	13.37	22	13.77	22	0.40
23*	33	23	16.73	32	16.73	1	0.00
24	56	50	13.26	26	14.06	30	0.80
25	0	10	16.29	28	17.92	28	1.63
26	17	98	12.47	0	12.90	17	0.43
27	60	35	14.26	36	15.49	24	1.23
28	10	0	18.21	25	18.72	15	0.51
29	21	32	15.53	35	15.81	14	0.28
30	98	17	7.34	31	17.30	67	9.96
31	35	60	13.16	17	13.41	18	0.25
32	31	100	11.54	0	12.90	31	1.36
33	100	31	6.67	35	15.92	65	9.25
34*	32	21	16.92	32	16.92	0	0.00
Average	40.12	40.12	13.82	24.85	15.39	27.15	1.56

* indicates that the payoff loss is smaller than $0.10.

Table 3.14: Extra Period Choices — Experiment 2b

Subject	Extra Period Choice	Opponent's Extra Period Choice	Expected Payoff	Best Response	Expected Payoff at BR	|(2) − (5)|	Payoff Loss
1	40	0	18.25	25	18.72	15	0.47
2	0	40	14.01	34	14.96	34	0.95
3*	30	10	17.91	28	17.92	2	0.01
4	67	99	7.76	0	12.90	67	5.14
5	63	37	13.78	37	15.27	26	1.49
6	10	30	15.24	35	16.02	25	0.78
7	37	63	12.86	15	13.27	22	0.41
8	50	0	17.38	25	18.72	25	1.34
9	11	41	14.45	33	14.85	22	0.40
10	99	67	5.45	11	13.12	88	7.67
11	0	50	13.52	26	14.06	26	0.54
12	41	11	17.50	29	17.83	12	0.33
13*	27	45	14.47	30	14.47	3	0.00
14	23	33	15.50	36	15.71	13	0.21
15	54	50	13.38	26	14.06	28	0.68
16	54	0	16.91	25	18.72	29	1.81
17	45	27	16.06	34	16.34	11	0.28
18	50	56	12.94	21	13.64	29	0.70
19	50	54	13.14	22	13.77	28	0.63
20	55	65	11.75	13	13.19	42	1.44
21	65	55	12.02	21	13.70	44	1.68
22	0	54	13.37	22	13.77	22	0.40
23*	33	23	16.73	32	16.73	1	0.00
24	56	50	13.26	26	14.06	30	0.80
25	0	10	16.29	28	17.92	28	1.63
26	17	98	12.47	0	12.90	17	0.43
27	60	35	14.26	36	15.49	24	1.23
28	10	0	18.21	25	18.72	15	0.51
29	21	32	15.53	35	15.81	14	0.28
30	98	17	7.34	31	17.30	67	9.96
31	35	60	13.16	17	13.41	18	0.25
32	31	100	11.54	0	12.90	31	1.36
33	100	31	6.67	35	15.92	65	9.25
34*	32	21	16.92	32	16.92	0	0.00
Average	40.12	40.12	13.82	24.85	15.39	27.15	1.56

* indicates that the payoff loss is smaller than $0.10.

Table 3.15: Last Period Choices — Experiment 2*c*							
Subject	Last Period Choice	Opponent's Last Period Choice	Expected Payoff	Best Response	Expected Payoff at BR	|(2) − (5)|	Payoff Loss
1	11	87	12.73	0	12.90	11	0.17
2	87	11	10.39	29	17.83	58	7.44
3	45	84	11.03	0	12.90	45	1.87
4	84	45	10.01	30	14.47	54	4.46
5	99	96	2.95	0	12.90	99	9.95
6	96	99	3.17	0	12.90	96	9.73
7	2	39	14.18	35	15.05	33	0.87
8	39	2	18.19	26	18.57	13	0.38
9	89	38	9.29	36	15.17	53	5.88
10	38	89	11.32	0	12.90	38	1.58
11*	36	13	17.56	29	17.66	7	0.10
12	13	36	14.90	37	15.38	24	0.48
13	59	60	11.99	17	13.41	42	1.42
14	60	59	12.04	18	13.46	42	1.42
15	90	23	9.48	32	16.73	58	7.25
16	23	90	12.23	0	12.90	23	0.67
17	85	38	10.16	36	15.17	49	5.01
18	38	85	11.49	0	12.90	38	1.41
19	48	93	10.29	0	12.90	48	2.61
20	93	48	7.93	27	14.21	66	6.28
21	72	93	7.53	0	12.90	72	5.37
22	93	72	6.37	7	12.98	86	6.61
23	36	7	18.01	27	18.17	9	0.16
24	7	36	14.63	37	15.38	30	0.75
25	46	99	10.23	0	12.90	46	2.67
26	99	46	6.55	29	14.39	70	7.84
27	34	97	11.36	0	12.90	34	1.54
28	97	34	7.44	36	15.60	61	8.16
29	16	5	18.08	27	18.33	11	0.25
30	5	16	16.16	30	17.39	25	1.23
31	99	44	6.62	31	14.57	68	7.95
32	44	99	10.43	0	12.90	44	2.47
33	90	91	5.06	0	12.90	90	7.84
34	91	90	5.01	0	12.90	91	7.89
35	65	10	14.98	28	17.92	37	2.94
36*	10	65	13.19	13	13.19	3	0.00
Average	56.64	56.64	10.92	17.14	14.60	46.50	3.68

* indicates that the payoff loss is smaller than $0.10.

Table 3.16: Test of Hypothesis $H1'$ — Experiment $2a^{(\#)}$

Subjects * 1\11	0–25	26–50	**51–75**	76–100	
0–25	16.64,17.52	15.83,15.79	**16.98,12.70**	12.52,7.92	**15.75**
26–50	17.20,14.71	15.49,15.07	14.05,12.98	9.15,9.93	14.46
51–75	15.18,12.55	7.73,17.09	16.35,8.03	4.46,12.62	11.23
76–100	9.45,12.60	10.40,10.35	8.87,7.64	6.42,6.07	9.08
	14.81	14.21	**10.65**	8.43	

Subjects 2\4 *	0–25	**26–50**	51–75	76–100	
0–25	18.80,15.85	**14.42,17.25**	12.90,13.31	12.90,12.57	**15.48**
26–50	17.97,14.26	18.00,10.25	12.34,14.19	10.73,10.13	17.28
51–75	15.80,12.65	10.08,14.61	12.98,4.46	8.56,8.52	11.70
76–100	10.13,12.65	7.24,9.84	6.90,8.56	7.32,4.10	8.04
	15.53	**16.59**	12.77	11.86	

Subjects * 3\9	**0–25**	26–50	51–75	76–100	
0–25	17.94,16.19	12.30,18.00	.	.	17.64
26–50	**18.45,14.36**	15.98,13.57	.	.	**18.32**
51–75	16.46,12.90	14.41,12.12	.	.	16.35
76–100	.	.	.	.	.
	15.15	15.64	.	.	NE

Subjects 5\10	0–25	26–50	**51–75**	76–100	
0–25	17.85,16.29	14.49,17.74	12.43,17.29	12.52,8.87	13.10
26–50	17.74,14.49	15.94,14.38	14.25,13.80	10.78,10.40	13.97
51–75	13.48,13.93	12.62,11.98	**11.17,11.17**	7.16,10.90	**10.65**
76–100	10.13,12.65	12.05,10.35	8.04,9.52	8.18,1.55	8.67
	14.02	12.47	**11.71**	9.57	

Subjects 6\8	0–25	26–50	**51–75**	76–100	
0–25	12.76,21.15	16.69,14.61	11.96,14.40	13.45,9.33	13.79
26–50	11.06,20.81	9.86,18.00	11.89,14.40	12.84,8.03	10.98
51–75	11.65,15.86	13.12,12.31	11.50,10.98	15.02,4.01	12.13
76–100	9.30,13.04	7.96,11.27	**7.08,10.17**	4.82,6.59	**7.99**
	14.87	12.12	**10.90**	5.94	

Subjects * 7\12	**0–25**	26–50	51–75	76–100	
0–25	**16.77,17.58**	12.58,18.67	.	12.90,6.75	**16.28**
26–50	14.77,17.08	14.68,15.63	.	18.00,−2.10	14.85
51–75	16.35,12.75	14.61,10.50	.	9.28,8.09	15.99
76–100	6.75,12.79	5.71,15.16	.	5.23,6.18	6.61
	17.21	18.20	.	5.71	

Table 3.16 continued

Subjects * 13\14 *	**0–25**	26–50	51–75	76–100	
0–25	**15.60,19.00**	14.41,18.48	12.65,15.80	12.52,8.46	**14.95**
26–50	12.55,18.22	13.77,16.42	13.06,13.47	10.73,10.13	12.88
51–75	15.54,12.55	13.26,12.10	15.37,5.41	9.28,7.81	14.58
76–100	12.15,12.56	10.13,10.73	12.15,5.41	12.15,4.24	11.53
	17.16	16.38	12.65	8.36	NE

Subjects * 15\17 *	0–25	**26–50**	51–75	76–100	
0–25	.	**14.33,18.21**	12.90,15.23	12.65,9.69	**13.43**
26–50	.	10.73,19.15	9.93,17.69	11.17,8.78	10.65
51–75	.	14.50,12.04	9.16,13.58	8.36,7.64	11.19
76–100	.	6.75,10.96	6.85,7.59	2.64,5.77	5.52
	.	**17.33**	14.49	9.12	

Subjects 16\18	0–25	26–50	**51–75**	76–100	
0–25	.	14.48,17.75	12.76,14.11	12.62,11.83	12.76
26–50	.	18.15,9.15	11.60,12.87	10.23,11.38	11.32
51–75	.	17.38,9.15	**10.88,10.13**	9.24,9.41	**10.51**
76–100	.	10.13,10.73	8.93,7.81	4.87,8.30	7.61
	.	11.75	**11.62**	10.43	

Subjects * 19\24	**0–25**	26–50	51–75	76–100	
0–25	20.81,12.90	11.96,18.85	11.96,15.41	12.27,6.88	14.26
26–50	**17.87,14.58**	14.21,15.52	14.34,13.38	10.59,8.14	**14.43**
51–75	15.80,12.65	12.75,13.24	9.43,10.99	8.65,8.43	12.39
76–100	6.75,12.26	7.46,10.41	6.75,7.50	6.75,−2.10	7.07
	13.70	15.69	13.08	8.60	

Subjects 20\27 *	**0–25**	26–50	51–75	76–100	
0–25	16.19,17.94	15.06,16.78	13.90,15.42	12.65,10.13	14.90
26–50	17.90,14.65	14.50,14.47	14.16,13.13	9.86,9.87	15.18
51–75	15.18,12.58	9.94,15.36	11.66,13.42	9.71,7.37	11.88
76–100	**9.32,12.71**	8.81,10.04	7.42,9.82	5.53,5.89	**8.40**
	14.73	14.24	13.04	8.60	

Subjects * 21\22 *	0–25	**26–50**	51–75	76–100	
0–25	19.43,14.67	**14.38,17.40**	12.57,17.09	11.96,6.75	**14.24**
26–50	15.98,17.33	16.42,13.47	10.68,15.11	10.73,10.13	14.15
51–75	16.35,12.60	12.71,13.03	13.70,9.71	8.85,6.88	13.14
76–100	.	.	.	.	.
	14.54	**14.71**	13.73	7.66	

Table 3.16 continued

Subjects 23\26 *	**0–25**	26–50	51–75	76–100	
0–25	16.08,17.64	16.47,14.47	13.99,14.03	12.41,6.75	15.22
26–50	**18.21,13.91**	11.15,18.00	11.48,16.38	11.02,6.75	**14.65**
51–75	17.54,12.72	13.58,12.95	9.67,13.07	9.16,7.92	14.23
76–100	.	.	.	.	.
	16.03	15.91	14.99	6.77	

Subjects 25\28 *	**0–25**	26–50	51–75	76–100	
0–25	.	.	.	.	.
26–50	18.20,14.16	14.27,16.04	12.10,14.43	10.24,6.75	15.34
51–75	15.05,12.90	15.02,9.15	10.49,12.26	6.09,6.75	12.42
76–100	**8.87,12.86**	9.39,9.82	3.79,12.27	3.67,3.21	**7.34**
	13.04	10.59	12.56	4.01	

Subjects * 29\35 *	0–25	**26–50**	51–75	76–100	
0–25	16.66,17.33	15.77,14.94	12.04,14.40	.	16.17
26–50	18.56,13.89	**13.67,15.75**	10.16,15.37	.	**16.53**
51–75	16.35,12.90	14.31,13.00	10.27,12.48	.	15.40
76–100	.	.	.	.	.
	14.33	**15.33**	14.91	.	

Subjects * 30\33 *	0–25	**26–50**	51–75	76–100	
0–25	17.26,16.90	13.58,18.70	13.34,15.20	12.90,11.59	14.13
26–50	18.88,13.35	**18.00,11.43**	11.44,15.09	11.50,9.36	**15.72**
51–75	.	.	.	.	.
76–100	10.13,12.65	3.80,14.25	9.43,7.65	6.87,4.54	6.81
	16.49	**18.04**	14.89	11.16	

Subjects 31\32 *	**0–25**	26–50	51–75	76–100	
0–25	**16.34,18.30**	.	.	12.90,6.75	**16.25**
26–50	18.00,12.90	.	.	10.73,10.13	17.81
51–75	.	.	.	.	.
76–100	.	.	.	.	.
	18.23	.	.	6.80	

Subjects * 34\36	**0–25**	26–50	51–75	76–100	
0–25	17.33,17.33	.	.	.	17.33
26–50	**20.31,12.90**	.	.	.	**20.31**
51–75	.	.	.	.	.
76–100	6.75,12.90	.	.	.	6.75
	17.15	.	.	.	

(#) entries in bold face are the subjects' last period choices and payoffs.
* indicates that the last period choice is a best response to the opponent's historical distribution of choices according to the estimated payoff matrix.
NE indicates that the subjects' choices correspond to a Nash Equilibrium for the Convention Game.

Table 3.17: Test of Hypothesis $H1'$ — Experiment 2b(#)

Subjects * 1\2	**0–25**	26–50	51–75	76–100	
0–25	18.80,15.85	13.63,18.43	12.66,17.32	.	13.75
26–50	**18.39,13.84**	16.16,13.19	9.60,16.64	.	**15.64**
51–75	15.80,12.65	13.44,12.30	6.56,16.35	.	12.89
76–100	.	.	.	.	.
	14.74	15.96	16.97	.	

Subjects * 3\6	**0–25**	26–50	51–75	76–100	
0–25	16.55,17.30	13.94,18.29	12.65,15.80	12.65,10.13	16.32
26–50	**18.50,14.19**	15.18,15.68	11.55,13.31	10.73,10.13	**18.11**
51–75	17.21,11.96	14.50,12.03	11.51,11.24	5.55,12.62	16.75
76–100	10.13,12.65	10.13,10.73	8.84,8.24	12.15,−2.10	10.17
	15.35	16.56	14.18	10.04	

Subjects 4\10	0–25	26–50	51–75	**76–100**	
0–25	19.37,14.20	16.06,15.94	12.56,14.58	12.48,10.27	14.29
26–50	19.03,12.87	13.86,14.80	13.13,12.00	10.03,12.28	13.43
51–75	15.02,12.89	9.23,16.00	6.76,14.94	**14.16,0.84**	**10.03**
76–100	12.51,12.46	9.28,12.81	7.18,10.45	5.73,5.72	8.08
	13.14	14.89	12.87	**8.36**	

Subjects 5\7 *	0–25	**26–50**	51–75	76–100	
0–25	15.56,18.12	12.54,20.24	.	.	13.31
26–50	16.03,15.16	13.96,16.78	.	.	14.48
51–75	12.63,14.61	**12.54,14.81**	.	.	**12.56**
76–100	10.13,12.65	11.96,11.66	.	.	11.50
	15.09	**16.03**	.	.	

Subjects * 8\11 *	**0–25**	26–50	51–75	76–100	
0–25	.	.	.	.	.
26–50	**16.32,14.50**	15.16,15.16	.	9.86,6.75	**16.07**
51–75	15.56,13.03	14.70,11.05	.	6.76,12.86	15.27
76–100	12.74,12.46	12.86,11.06	.	5.71,5.71	12.56
	13.12	11.38	.	12.09	

Subjects * 9\12 *	0–25	**26–50**	51–75	76–100	
0–25	21.46,11.96	**12.62,18.07**	12.58,15.34	12.50,8.83	**12.75**
26–50	20.74,12.86	14.76,15.56	10.73,15.80	10.94,9.02	11.29
51–75	15.80,12.65	15.80,10.73	5.96,15.02	6.73,10.20	7.15
76–100	10.13,12.65	10.13,10.73	7.05,4.46	4.56,4.70	5.06
	12.49	**13.85**	12.31	8.02	

Table 3.17 continued

Subjects 13\17	0–25	**26–50**	51–75	76–100	
0–25	21.41,12.17	16.74,15.29	15.97,12.48	12.65,10.13	17.78
26–50	16.91,14.93	**16.41,13.90**	14.21,10.58	14.05,6.81	**16.24**
51–75	12.94,16.58	14.50,12.32	7.78,14.40	8.53,8.21	13.22
76–100	10.13,12.65	7.45,13.01	7.26,9.82	6.59,4.82	8.09
	14.89	**13.64**	11.61	7.38	

Subjects * 14\23 *	0–25	**26–50**	51–75	76–100	
0–25	21.71,12.17	**14.92,17.47**	12.79,15.77	12.75,10.26	**14.59**
26–50	10.38,21.50	9.15,19.81	10.38,17.38	10.73,10.13	9.52
51–75	.	.	.	.	.
76–100	.	.	.	.	.
	12.54	**17.56**	15.83	10.25	NE

Subjects 15\19	0–25	**26–50**	51–75	76–100	
0–25	20.89,12.90	21.46,11.27	12.40,15.02	.	14.60
26–50	16.93,15.11	19.46,11.27	13.41,11.86	.	14.54
51–75	16.27,12.90	**12.85,13.68**	10.04,11.97	.	**11.30**
76–100	10.13,12.65	10.13,10.73	10.69,7.94	.	10.55
	13.59	**12.49**	11.95	.	

Subjects * 16\22 *	**0–25**	26–50	51–75	76–100	
0–25	.	.	.	.	.
26–50	17.92,14.31	15.16,15.15	9.78,16.35	9.30,6.75	10.89
51–75	**16.88,12.72**	15.20,9.35	11.38,11.71	11.06,7.15	**12.15**
76–100	10.13,12.65	10.13,10.73	9.75,9.09	8.43,7.73	9.69
	12.78	9.67	11.72	7.17	

Subjects * 18\24	0–25	26–50	**51–75**	76–100	
0–25	16.86,17.28	12.52,19.47	12.04,16.88	12.47,10.59	12.39
26–50	16.53,15.70	14.79,14.30	**14.83,11.79**	11.38,7.87	**13.37**
51–75	17.21,12.89	16.81,10.36	11.85,11.22	9.57,6.20	12.01
76–100	10.13,12.65	9.25,11.62	0.75,16.17	−2.10,13.09	1.47
	14.98	13.69	**12.35**	7.85	

Subjects * 20\21	0–25	26–50	**51–75**	76–100	
0–25	18.21,15.93	11.96,18.71	12.65,15.80	.	12.61
26–50	9.73,20.81	16.05,12.52	12.48,14.07	.	14.05
51–75	16.91,11.99	15.39,11.72	**12.95,11.88**	.	**14.33**
76–100	10.13,12.65	12.15,11.27	−1.80,17.21	.	5.53
	15.35	12.09	**12.89**	.	

Table 3.17 continued

Subjects 25\28 *	**0–25**	26–50	51–75	76–100	
0–25	**16.79,17.71**	15.43,16.97	12.90,16.35	.	**16.56**
26–50	.	.	.	.	.
51–75	.	.	.	.	.
76–100	.	.	.	.	.
	17.71	16.97	16.35	.	

Subjects * 26\30	0–25	26–50	51–75	**76–100**	
0–25	19.29,14.55	17.19,16.15	14.68,13.65	**12.61,9.48**	**14.53**
26–50	11.17,21.65	11.27,19.58	13.81,11.86	11.58,9.03	11.71
51–75	15.02,12.90	8.20,20.57	17.05,5.76	9.71,6.71	11.43
76–100	10.13,12.65	9.09,11.78	6.49,10.59	8.52,1.39	8.67
	15.79	17.28	11.88	**8.43**	

Subjects 27\31 *	0–25	**26–50**	51–75	76–100	
0–25	12.90,21.53	16.22,16.01	12.65,15.80	12.65,10.13	13.82
26–50	16.67,14.79	14.63,14.71	13.69,12.23	10.06,11.75	14.32
51–75	16.85,12.34	**11.64,14.54**	11.46,10.79	11.69,5.34	**13.13**
76–100	.	.	.	.	.
	14.19	**14.68**	11.87	9.93	

Subjects 29\34	0–25	**26–50**	51–75	76–100	
0–25	15.47,18.11	**15.43,16.86**	12.56,15.01	12.43,11.00	**14.30**
26–50	20.40,11.96	15.98,15.49	11.55,16.03	18.00,0.75	17.02
51–75	.	.	.	.	.
76–100	.	.	.	.	.
	17.62	**16.75**	15.09	10.18	

Subjects 32\33	0–25	26–50	51–75	**76–100**	
0–25	18.21,16.44	.	.	.	18.21
26–50	.	.	.	.	.
51–75	.	.	.	.	.
76–100	.	.	.	.	.
	16.44	.	.	.	

(#) entries in bold face are the subjects' last period choices and payoffs.
* indicates that the extra period choice is a best response to the opponent's historical distribution of choices according to the estimated payoff matrix.
NE indicates that the subjects' choices correspond to a Nash Equilibrium for the Convention Game.

Table 3.18: Test of Hypothesis $H2'$ — Experiment $2a$

Subject	Arms	# of Times Chosen	Average Payoff	Standard Deviation	Last Period Choice	# of Times Chosen
1*	0	9	13.88	2.95	20	8
	20	8	16.73	4.73		
	35	5	16.37	4.85		
	50	7	12.95	4.73		
	80	7	9.62	4.32		
	100	6	6.75	0.00		
	other	33	12.36	5.19		
2*	0	62	15.76	4.17	0	62
	other	13	11.21	4.87		
3 @	15	6	16.99	4.85	30	12
	20	19	18.82	4.01		
	29	6	19.01	3.62		
	30	12	18.92	3.44		
	other	32	16.75	3.63		
4*	0	21	15.85	4.28	30	7
	30	7	17.87	4.32		
	40	7	16.82	4.32		
	45	14	16.19	4.15		
	50	17	16.44	3.48		
	other	9	10.60	3.59		
5*	65	6	8.04	3.62	66	11
	66	11	11.19	4.62		
	73	6	10.81	4.57		
	other	52	11.33	3.99		
6	88	5	8.36	3.96	100	3
	other	70	10.10	4.71		
7	0	15	14.67	3.65	0	15
	5	40	17.07	4.48		
	25	5	15.50	4.85		
	other	15	13.68	4.99		
8	20	15	15.25	4.32	70	12
	40	7	13.03	4.32		
	50	15	12.10	4.32		
	60	11	10.72	4.46		
	70	12	9.98	4.62		
	other	15	12.86	4.61		
9	0	69	15.21	3.92	15	1
	other	6	14.76	3.78		
10*	53	5	12.23	4.85	54	6
	54	6	15.90	3.62		
	56	7	11.99	4.73		
	65	8	10.99	4.73		
	76	5	13.09	0.00		
	87	5	6.86	4.85		
	other	39	11.05	4.16		

Table 3.18 cont'd

Subject	Arms	# of Times Chosen	Average Payoff	Standard Deviation	Last Period Choice	# of Times Chosen
11	43	6	17.51	3.62	54	4
	57	11	8.83	2.67		
	other	58	12.92	5.12		
12*	0	20	15.56	4.16	12	9
	10	14	16.55	4.55		
	12	9	19.57	3.90		
	15	10	18.76	4.28		
	22	5	19.25	3.96		
	other	17	15.46	2.97		
13	0	32	15.11	3.89	0	32
	10	6	14.23	3.62		
	40	5	14.04	4.85		
	60	7	16.35	0.00		
	70	6	14.40	0.00		
	other	19	13.22	3.61		
14	25	8	17.50	4.58	11	2
	31	10	19.43	2.80		
	50	5	14.46	4.85		
	other	52	15.37	5.61		
15*	0	33	13.70	2.58	23	10
	2	6	12.89	0.00		
	23	10	13.88	3.74		
	100	5	4.98	3.96		
	other	21	10.28	3.80		
16	2	6	12.89	0.00	69	2
	65	5	8.33	3.96		
	67	5	9.71	4.85		
	other	59	11.05	3.95		
17*	30	5	20.40	0.00	30	5
	89	5	8.10	3.96		
	other	65	14.20	4.93		
18	69	6	11.66	4.57	72	5
	70	9	11.45	4.43		
	71	5	14.19	0.00		
	72	5	13.97	0.00		
	75	8	11.10	4.10		
	80	7	8.36	4.73		
	other	35	10.94	3.35		
19	25	16	13.62	3.57	50	5
	30	9	13.52	3.90		
	35	26	15.83	4.50		
	50	5	14.46	4.85		
	100	8	6.75	0.00		
	other	11	11.81	3.40		

Table 3.18 cont'd

Subject	Arms	# of Times Chosen	Average Payoff	Standard Deviation	Last Period Choice	# of Times Chosen
20	50	8	14.69	4.58	90	4
	other	67	12.66	4.82		
21	20	7	17.36	4.73	25	5
	25	5	13.73	3.96		
	30	6	15.98	4.85		
	50	8	11.36	4.10		
	55	7	12.16	4.73		
	60	5	14.58	3.96		
	65	10	11.87	4.57		
	70	6	12.92	3.62		
	other	21	14.68	3.82		
22	0	8	15.11	4.10	45	5
	32	5	13.13	3.96		
	40	5	15.81	4.85		
	42	5	15.56	4.85		
	45	5	13.40	4.85		
	60	6	14.87	3.62		
	other	41	12.88	5.36		
23	23	8	15.43	4.58	32	8
	32	8	14.69	4.58		
	43	6	14.55	4.85		
	other	53	15.12	4.19		
24	0	7	12.90	0.00	0	7
	20	5	15.84	4.85		
	30	6	15.98	4.85		
	40	6	17.87	3.62		
	50	15	14.46	4.49		
	100	13	6.07	2.45		
	other	23	13.04	4.27		
25	90	27	8.29	3.20	96	3
	95	14	6.95	3.22		
	100	8	5.65	3.13		
	other	26	11.57	5.57		
26*	0	10	15.56	4.28	25	22
	25	22	16.39	4.53		
	50	13	15.96	3.88		
	100	14	6.75	0.00		
	other	16	14.75	3.48		
27	2	5	16.43	4.85	25	1
	40	5	14.04	4.85		
	50	8	13.58	4.73		
	51	17	13.16	4.55		
	other	40	14.11	4.55		

Table 3.18 cont'd

Subject	Arms	# of Times Chosen	Average Payoff	Standard Deviation	Last Period Choice	# of Times Chosen
28*	0	40	13.12	1.40	0	40
	50	5	9.15	0.00		
	100	21	4.22	4.10		
	other	9	11.54	2.27		
29	34	7	13.70	4.32	37	10
	36	11	19.01	2.67		
	37	10	18.81	2.80		
	45	5	13.40	4.85		
	other	42	15.99	4.16		
30*	0	52	14.27	3.23	50	5
	50	5	18.00	0.00		
	other	18	12.29	5.11		
31*	0	66	16.12	4.29	0	66
	other	9	17.34	4.43		
32	0	33	17.99	4.44	0	33
	2	29	18.39	4.37		
	5	11	18.50	4.46		
	other	2	6.75	0.00		
33*	31	5	16.77	4.85	31	5
	other	70	15.96	3.70		
34	0	72	17.33	4.46	31	1
	other	3	11.27	7.83		
35*	0	44	14.11	3.08	50	14
	40	9	13.45	4.43		
	50	14	16.10	3.77		
	other	8	16.22	3.56		
36	0	75	17.15	4.45	0	75

* indicates that the last period choice is the arm associated with the highest average payoff.
@ indicates that the extra period choice is an arm and the (dollar) distance from the arm associated with the highest average payoff is smaller than $0.10.

Table 3.19: Test of Hypothesis $H2'$ — Experiment $2b^{(\#)}$

Subject	Arms	# of Times Chosen	Average Payoff	Standard Deviation	Extra Period Choice	# of Times Chosen
1*	0	29	13.51	2.28	40	12
	22	8	15.50	4.58		
	23	5	13.88	3.96		
	40	12	16.40	4.36		
	50	7	12.95	4.73		
	65	11	13.00	4.13		
	other	3	12.75	4.65		
2	30	9	17.45	4.43	0	2
	39	6	15.05	4.85		
	40	14	14.93	4.59		
	45	11	14.69	4.62		
	50	7	15.47	4.32		
	other	28	16.93	3.38		
3	23	9	16.04	4.67	30	9
	24	8	17.57	4.58		
	25	7	18.29	4.32		
	28	6	19.10	3.62		
	30	9	18.44	3.90		
	35	5	16.37	4.85		
	50	7	16.73	3.45		
	other	24	16.77	4.55		
4	22	5	17.48	4.85	67	5
	45	8	14.29	4.73		
	67	5	9.71	4.85		
	other	57	11.32	4.91		
6	0	32	15.39	4.04	10	1
	1	7	17.96	4.73		
	2	8	15.11	4.10		
	5	9	15.81	4.43		
	other	19	13.60	5.27		
7	25	11	15.18	4.46	37	0
	30	20	15.09	4.45		
	31	6	17.36	4.57		
	32	6	18.74	3.62		
	33	8	17.90	4.10		
	35	8	15.49	4.73		
	other	16	14.23	4.01		
8	55	5	17.21	0.00	50	3
	58	5	16.70	0.00		
	65	17	15.41	0.00		
	67	5	15.02	0.00		
	72	5	13.97	0.00		
	other	38	14.94	4.12		
10	46	6	17.10	3.62	99	0
	65	5	11.87	4.85		
	67	6	12.07	4.57		
	other	58	11.95	4.58		

Table 3.19 cont'd

Subject	Arms	# of Times Chosen	Average Payoff	Standard Deviation	Extra Period Choice	# of Times Chosen
11	1	6	12.90	0.00	0	4
	3	5	14.66	3.90		
	5	5	12.86	0.00		
	7	5	12.83	0.00		
	12	5	12.68	0.00		
	15	5	12.56	0.00		
	18	5	12.41	0.00		
	other	39	12.82	2.33		
12	85	9	10.91	0.00	41	0
	95	6	3.79	4.85		
	99	17	6.53	2.15		
	other	43	9.90	4.38		
13	27	12	16.23	4.62	27	12
	45	5	16.94	3.96		
	53	5	13.40	4.85		
	other	53	14.91	5.08		
14	2	11	15.31	4.13	23	0
	6	8	13.95	3.13		
	7	5	14.60	3.96		
	10	32	14.13	3.26		
	other	19	14.18	4.12		
15	45	6	17.24	3.62	54	3
	52	7	12.64	4.73		
	65	6	9.51	4.57		
	68	6	10.39	4.85		
	75	6	11.84	3.62		
	other	44	13.25	4.82		
16	57	21	13.08	4.49	54	0
	65	9	12.47	4.43		
	67	22	11.39	4.46		
	other	23	11.23	4.55		
17	22	14	14.70	4.15	45	22
	32	21	13.47	3.86		
	45	22	14.29	4.53		
	70	8	11.09	4.58		
	other	10	13.25	4.41		
18	30	7	15.35	4.73	50	10
	40	14	13.03	4.15		
	46	5	13.27	4.85		
	50	10	12.69	4.57		
	55	12	12.05	4.57		
	other	27	11.36	4.87		
19	0	12	13.64	2.56	50	0
	33	7	11.27	0.00		
	67	56	11.86	4.28		
	other	1	18.00	0.00		

Table 3.19 cont'd

Subject	Arms	# of Times Chosen	Average Payoff	Standard Deviation	Extra Period Choice	# of Times Chosen
20	46	7	14.78	4.73	55	4
	58	13	15.35	3.32		
	59	5	12.99	4.85		
	60	8	13.04	4.58		
	other	42	13.75	4.79		
21	35	5	12.83	3.96	65	5
	45	13	12.59	4.25		
	55	25	13.32	4.49		
	65	5	11.87	4.85		
	other	27	12.07	3.62		
22	10	5	12.75	0.00	0	0
	50	5	9.15	0.00		
	60	27	11.76	4.51		
	61	19	11.98	4.54		
	70	6	11.45	4.57		
	80	5	8.61	4.85		
	other	8	8.99	4.52		
23	32	7	17.69	4.32	33	8
	33	8	17.90	4.10		
	34	6	18.54	3.62		
	35	9	15.98	4.67		
	45	6	17.24	3.62		
	other	39	15.72	4.13		
24	45	5	13.40	4.85	56	2
	99	6	4.10	4.57		
	100	9	5.77	2.95		
	other	55	12.15	4.49		
25	0	75	16.56	4.39	0	75
26*	17	17	16.63	4.55	17	17
	33	6	12.74	3.62		
	other	52	12.45	4.15		
27	40	16	15.48	4.54	60	6
	45	5	15.17	4.85		
	50	26	13.58	4.52		
	55	8	13.90	4.58		
	60	6	10.45	4.57		
	other	14	14.45	4.10		
28	0	24	16.59	4.46	10	11
	10	11	17.58	4.62		
	20	7	18.62	4.32		
	other	33	18.14	4.22		
29	1	29	14.42	3.41	21	0
	20	31	14.30	3.77		
	25	7	13.23	3.35		
	30	5	15.09	4.85		
	other	3	17.45	4.45		

Table 3.19 cont'd

Subject	Arms	# of Times Chosen	Average Payoff	Standard Deviation	Extra Period Choice	# of Times Chosen
30	18	8	14.63	4.10	98	4
	28	5	17.03	4.85		
	88	7	7.61	4.32		
	other	55	9.77	4.89		
31	12	7	15.21	4.32	35	3
	76	5	9.55	4.85		
	other	63	13.06	4.61		
32	0	57	18.18	4.38	31	0
	1	10	17.33	4.67		
	2	7	19.22	4.32		
	other	1	21.74	0.00		
33	0	75	16.44	4.37	100	0
34	25	10	18.16	4.28	32	1
	34	8	17.81	4.10		
	45	6	14.29	4.85		
	67	5	15.02	0.00		
	88	5	10.13	0.00		
	other	41	14.77	4.76		

(#) subjects 5 and 9 do not have any choice that can be classified as an arm.
* indicates that the extra period choice is the arm associated with the highest average payoff.

Table 3.20: Test of Hypotheses $H3'$ and $H4'$ — Experiment $2a$

Number of Choices Consistent with Numerical Hill Climbing (Last 50 Rounds)																
Subject	1	2*	3	4	5*	6	7	8	9	10	11	12*	13	14	15	16*
# of Choices	27	46	31	17	35	30	28	32	18	31	22	35	28	23	17	33
Subject	17*	18*	19	20	21	22	23*	24	25*	26	27	28	29*	30*	31*	32
# of Choices	35	38	23	26	32	25	34	29	36	25	27	20	34	33	38	8
Subject	33*	34	35	36												Average
# of Choices	41	16	5	0												27.2

Number of Choices Consistent with Strategic Myopic Adjustment (Last 50 Rounds)																
Subject	1*	2	3	4	5	6	7	8*	9	10	11	12	13	14	15	16*
# of Choices	33	21	30	28	31	24	27	44	22	31	30	29	22	18	19	34
Subject	17	18	19*	20	21	22	23	24*	25*	26*	27	28	29*	30	31	32
# of Choices	20	32	37	32	32	32	25	37	40	33	31	21	39	17	18	32
Subject	33*	34	35	36												Average
# of Choices	33	24	13	28												28.3

* indicates that at least 2/3 of the choices in the last 50 rounds are consistent with the adaptive rule.

Table 3.21: Test of Hypotheses $H3'$ and $H4'$ — Experiment $2b$

Number of Choices Consistent with Numerical Hill Climbing (Last 50 Rounds)																
Subject	1*	2	3	4	5*	6*	7	8*	9	10*	11*	12	13	14	15	16*
# of Choices	33	30	30	31	34	33	24	43	25	35	38	32	28	26	31	33
Subject	17	18	19	20*	21	22	23	24	25	26*	27	28	29	30	31	32*
# of Choices	19	28	20	36	26	9	28	29	0	33	25	31	16	28	26	43
Subject	33	34*														Average
# of Choices	0	34														27.6

Number of Choices Consistent with Strategic Myopic Adjustment (Last 50 Rounds)																
Subject	1	2*	3	4*	5*	6	7	8*	9	10	11	12	13*	14	15*	16*
# of Choices	29	33	29	39	43	26	28	47	32	28	18	32	33	28	41	33
Subject	17	18*	19	20	21*	22	23*	24	25	26	27*	28	29	30*	31	32*
# of Choices	26	40	30	30	40	27	38	32	21	32	33	29	18	34	32	34
Subject	33	34*														Average
# of Choices	20	33														31.4

* indicates that at least 2/3 of the choices in the last 50 rounds are consistent with the adaptive rule.

APPENDIX A

Table A.1: Coefficients of the Theoretical Payoff Function

payoff function: $\pi = \alpha + \beta e + \gamma e^2$, e: effort level		
α	β	γ
18.94	0.079	−0.0011

Table A.2: Least–Squares Estimates of the Payoff Function — Experiment 1*a*
(Standard Errors in Parentheses)

Subject	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\gamma}$	F–test*
1	12.56 (14.10)	0.371 (0.509)	−0.0043 (0.0043)	0.78
2	17.37 (1.57)	0.174 (0.728)	−0.0025 (0.0007)	2.06
3	−19.57 (23.97)	1.505 (0.863)	−0.0144 (0.0077)	2.07
4	18.25 (1.20)	0.161 (0.058)	−0.0024 (0.0005)	7.96
5	19.40 (5.67)	0.160 (0.231)	−0.0029 (0.0021)	1.82
6	17.06 (1.68)	0.193 (0.075)	−0.0028 (0.0007)	11.13
7	14.22 (3.41)	0.373 (0.173)	−0.0050 (0.0021)	1.20
8	19.78 (1.99)	0.153 (0.117)	−0.0032 (0.0016)	1.93
9	20.70 (1.08)	0.067 (0.963)	−0.0018 (0.0006)	5.64
10	17.10 (1.63)	0.183 (0.067)	−0.0026 (0.0006)	10.35
11	20.24 (2.23)	0.131 (0.116)	−0.0027 (0.0013)	2.55
12	18.50 (2.19)	0.160 (0.094)	−0.0028 (0.0009)	6.93
13	21.63 (4.29)	0.076 (0.147)	−0.0021 (0.0012)	11.10
14	15.90 (1.55)	0.252 (0.073)	−0.0033 (0.0007)	7.22
15	17.87 (2.24)	0.205 (0.107)	−0.0029 (0.0013)	1.11
16	18.62 (3.95)	0.139 (0.197)	−0.0023 (0.0019)	0.54
17	16.26 (2.54)	0.263 (0.110)	−0.0034 (0.0012)	1.48
18	8.42 (4.18)	0.483 (0.160)	−0.0048 (0.0014)	2.95
19	13.06 (7.12)	0.338 (0.253)	−0.0037 (0.0022)	0.65
20	16.33 (2.90)	0.220 (0.110)	−0.0029 (0.0010)	1.56
21	17.43 (1.76)	0.157 (0.083)	−0.0023 (0.0011)	0.64
22	15.93 (1.61)	0.274 (0.067)	−0.0034 (0.0007)	3.38
23	15.99 (5.15)	0.231 (0.231)	−0.0028 (0.0027)	0.24

* $H_0 : \hat{\alpha} = 18.94$, $\hat{\beta} = 0.079$, $\hat{\gamma} = -0.0011$; $F_{.95}(3,72) = 2.15$.

Table A.3: Least–Squares Estimates of the Payoff Function — Experiment 1*b*
(Standard Errors in Parentheses)

Subject	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\gamma}$	F–test*
1	9.79 (10.08)	0.410 (0.452)	−0.0044 (0.0045)	2.29
2	15.67 (8.36)	0.244 (0.427)	−0.0032 (0.0058)	0.08
3	12.79 (5.24)	0.306 (0.159)	−0.0035 (0.0011)	21.74
4	18.29 (2.20)	0.166 (0.084)	−0.0026 (0.0007)	12.61
5**	–	–	–	–
6	−33.97 (36.93)	2.579 (1.718)	−0.0301 (0.0199)	0.83
7	18.55 (0.78)	0.344 (0.247)	−0.0055 (0.0037)	0.84
8	17.32 (11.86)	0.168 (0.354)	−0.0025 (0.0025)	5.69
9	11.27 (13.80)	0.304 (0.545)	−0.0027 (0.0051)	0.85
10	−12.94 (21.88)	0.987 (0.624)	−0.0078 (0.0044)	3.27
11	17.77 (0.87)	0.147 (0.057)	−0.0023 (0.0007)	3.38
12	17.98 (2.35)	0.192 (0.101)	−0.0029 (0.0009)	8.38
13	18.44 (0.52)	0.165 (0.067)	−0.0027 (0.0007)	13.33
14	15.45 (5.01)	0.255 (0.175)	−0.0035 (0.0014)	5.33
15	18.51 (1.11)	0.168 (0.056)	−0.0026 (0.0006)	15.53
16**	–	–	–	–
17	−60.65 (60.70)	2.830 (2.010)	−0.0251 (0.0165)	3.42
18**	–	–	–	–
19	19.11 (0.89)	0.061 (0.123)	−0.0009 (0.0022)	0.07
20	15.55 (4.93)	0.233 (0.333)	−0.0024 (0.0057)	0.53
21**	–	–	–	–
22**	–	–	–	–
23	18.20 (1.92)	0.213 (0.090)	−0.0033 (0.0010)	2.20
24	12.88 (6.25)	0.264 (0.194)	−0.0029 (0.0015)	6.80

* $H_0 : \hat{\alpha} = 18.94$, $\hat{\beta} = 0.079$, $\hat{\gamma} = -0.0011$; $F_{.95}(3,72) = 2.15$.

** Subjects 5, 16, 18, 21, and 22 could not have estimated a concave payoff function.

Table A.4: Delta Method

payoff function: $\pi = \alpha + \beta e + \gamma e^2$, e: effort level

true optimal effort level: $e^* = -\beta/2\gamma$

estimated optimal effort level: $\hat{e}* = -\hat{\beta}/2\hat{\gamma}$

estimate's standard error: $\sigma_{\hat{e}*} = \left[(1/4\hat{\gamma}^2)\sigma^2_{\hat{\beta}} - (\hat{\beta}/2\hat{\gamma}^3)\sigma_{\hat{\beta}\hat{\gamma}} + (\hat{\beta}^2/2\hat{\gamma}^4)\sigma^2_{\hat{\gamma}}\right]^{1/2}$

APPENDIX B: Instructions for Experiment 1a

Introduction

This is an experiment about decision making. The instructions are simple, and if you follow them carefully and make good decisions, you could earn a considerable amount of money, which will be paid to you in cash.

Specific Instructions

As you read these instructions you will be in a room with a number of other subjects. Each subject has been randomly assigned an ID number and a computer terminal. The experiment consists of 75 decision rounds. In each decision round you will be paired with a computerized subject which has been programmed to make the same decision in every round. The computerized subject randomly matched with you will be called your pair member. Your computerized pair member will remain the same throughout the entire experiment.

Experimental Procedure

In the experiment you will perform a simple task. Attached to these instructions is a sheet called your "Decision Cost Table". This sheet shows 101 numbers from 0 to 100 in column A. These are your decision numbers. Associated with each decision number is a decision cost, which is listed in column B. Note that the higher the decision number chosen, the greater is the associated cost. Your computer screen should look as follows as you entered the lab:

PLAYER #__

ROUND	DECISION #	RANDOM #	TOTAL #	COST	EARNINGS

In each decision round the computer will ask you to choose a decision number. Your computerized pair member will also choose a decision number. Remember that it will always choose the same decision number, which will be 37 in each decision round. You of course, are free to choose any number you wish among those listed in column A of your "Decision Cost Table". Therefore, in each round of the experiment, you and your computerized pair member will each select a decision number separately (and you know that it will always choose 37). Using the number keys, you will enter your selected number and then hit the Return (Enter) key. To verify your selection, the computer will then ask you the following question:

Is ___ your decision number ? [Y/N]

If the number shown is the one you desire, hit the Y key. If not, hit the N key and the computer will ask you to select a number again. After you have selected and verified your number, this number will be recorded on the screen in column 2, and its associated cost will be recorded in column 5. After you have selected your decision number, the computer will ask you to generate a random number. You do this by hitting the space bar (the long key at the bottom of the keyboard). Hitting the bar causes the computer to select one of the 81 numbers that fall between −40 and +40 (including 0). Each of these 81 numbers has an equally likely chance of being chosen when you hit the space bar. Hence, the probability that the computer selects, say, +40 is the same as the probability that it selects −40, 0, −12 or +27. Another random number (again between −40 and +40) will be automatically generated for your computerized pair member as well. The processes that generate your random number and the random number assigned to your computerized pair member are independent – i.e., you should not expect any relationship between the two random numbers generated to exist. After you hit the space bar, the computer will record your random number on the screen in column 3.

Calculation of Payoffs
Although the experiment will have 75 rounds in total, your actual payment will just be determined by your earnings in the last round of the experiment (round 75). In each of the first 74 rounds, however, you will have the chance of observing what your payoff would have been as a consequence of your choice, the choice of your computerized pair member (37), and the two random numbers generated. These hypothetical payments will not be paid, however. Only your period 75 payment will count. Your payment (either real or hypothethical) in each decision round will be computed as follows. After you select a decision number and generate a random number, the computer will add these two numbers and record the sum on the screen in column 4. We will call the number in column 4 your "Total Number". The computer will do the same computation for your computerized pair member as well. The computer will then compare your Total Number to that of your computerized pair member. If your Total Number is greater than your computerized pair member's Total Number, then you will receive the high fixed payment of 29 Fr., in a fictitious currency called Francs. If not, then you will receive the low fixed payment 17.2 Fr. Whether you receive the fixed payment 29 Fr. or the fixed payment 17.2 Fr. only depends on whether your Total Number is greater than your computerized pair member's Total Number. It does not depend on how much bigger it is. The Francs will be converted into dollars at the conversion rate to be stated below. The computer will record (on the screen in column 6) which fixed payment you receive. If you receive the high fixed payment (29 Fr.), then "M" will appear in column 6. If you receive the low fixed payment (17.2 Fr.), "m" will appear. After indicating which fixed payment you receive, the computer will subtract your associated decision cost (column 5) from this fixed payment. This difference represents your (actual or hypothetical) earnings for the round. The amount of your earnings will be recorded on the screen in column 6, right next to the letter ("M" or "m") showing your fixed payment.

Continuing Rounds
After round 1 is over, you will perform the same procedure for round 2, and so on for 75 rounds. In each round you will choose a decision number and generate a random number by pressing the space bar. Your Total Number will be compared to the Total Number of your computerized pair member, and the computer will calculate your earnings for the round. Your final earnings will depend only on your decisions in round 75. When that round is completed, the computer will ask you to press any key on its keyboard. After you do this, the computer will convert your Francs earnings in round 75 in Dollars at the rate of $ .75 per Franc. We will then pay you this amount.

Example of Payoff Calculations
Suppose that the following occurs during one round: pair member A_2 chooses a decision number of 60 and generates a random number of 10, while computerized pair member A_1 selects a decision number of 37 and gets a random number of 5. Pair member A_2 would then receive the high fixed payment of 29 Fr. From this fixed payment, A_2 would subtract 7.2 Fr. (the cost of decision number 60). A_2's earnings for that round would then be 21.8 Fr. (i.e., 29 Fr. – 7.2 Fr.). Note that the decision cost subtracted in column 5 is a function only of your decision number; i.e., your random number does not affect the amount subtracted. Also, note that your earnings depend on the following: the decision number you select (both because it contributes to your Total Number and because it determines the amount – i.e., your Decision Cost – to be subtracted from your fixed payment), your computerized pair member's pre–selected decision number (37), your generated random number, and your computerized pair member's generated random number.

Table 3.5: Descriptive Statistics of Subjects' Choices — Experiment 1*b*

Subject	Mean Choice	Standard Deviation	Median Choice	Interquantile Range
1	57.00	11.03	60	15
2	47.73	5.68	50	5
3	73.79	19.37	78	37
4	61.93	27.47	67	44
5	0.00	0.00	0	0
6	44.11	4.74	45	4
7	9.00	19.47	1	4
8	64.72	14.38	65	20
9	48.12	11.33	45	15
10	63.93	6.97	65	5
11	26.99	27.80	9	40
12	53.17	28.22	50	46
13	16.75	34.39	0	0
14	51.86	17.58	45	14
15	59.92	33.05	76	50.5
16	23.47	15.94	35	35
17	61.12	5.32	60	7
18	76.97	17.50	77	22
19	27.49	25.64	35	50
20	39.95	7.57	40	7
21	38.89	11.95	45	2.5
22	50.00	0.00	50	0
23	45.61	21.09	44	21.5
24	69.74	13.05	70	15.5
Average	46.34	15.81	46.33	19.17

Table 3.6: Descriptive Statistics of Subjects' Choices — Experiment 1*c*

Subject	Mean Choice	Standard Deviation	Median Choice	Interquantile Range
1	48.15	29.46	49	50
2	49.99	27.75	53	52
3	54.53	28.94	51	51
4	47.47	32.16	38	64
5	47.32	25.68	44	44
6	44.72	27.82	46	48
7	51.08	29.36	52	54
8	54.61	26.72	54	45
9	45.77	26.89	46	45
10	48.03	27.46	49	43
11	50.81	31.59	52	57
12	50.72	28.37	53	49
13	47.75	28.15	47	50
14	41.47	27.72	37	42
15	49.96	29.67	46	51
16	49.01	27.48	47	43
17	46.79	27.83	45	43
18	52.08	28.29	52	44
19	44.89	26.01	43	39
20	54.88	30.16	61	54
21	46.37	29.04	46	59
22	52.39	29.00	53	48
23	49.71	27.74	45	45
24	53.80	30.68	59	55
Average	49.26	28.50	48.67	48.96

Decision Cost Table

Column A Decision Number	**Column B** Cost of Decision (Francs)	**Column A** Decision Number	**Column B** Cost of Decision (Francs)	**Column A** Decision Number	**Column B** Cost of Decision (Francs)
0	0.000	36	2.59	72	10.37
1	0.002	37	2.74	73	10.66
2	0.008	38	2.89	74	10.95
3	0.02	39	3.04	75	11.25
4	0.03	40	3.20	76	11.55
5	0.05	41	3.36	77	11.86
6	0.07	42	3.53	78	12.17
7	0.10	43	3.70	79	12.48
8	0.13	44	3.87	80	12.80
9	0.16	45	4.05	81	13.12
10	0.20	46	4.23	82	13.45
11	0.24	47	4.42	83	13.78
12	0.29	48	4.61	84	14.11
13	0.34	49	4.80	85	14.45
14	0.39	50	5.00	86	14.79
15	0.45	51	5.20	87	15.14
16	0.51	52	5.41	88	15.49
17	0.58	53	5.62	89	15.84
18	0.65	54	5.83	90	16.20
19	0.72	55	6.05	91	16.56
20	0.80	56	6.27	92	16.93
21	0.88	57	6.50	93	17.30
22	0.97	58	6.73	94	17.67
23	1.06	59	6.96	95	18.05
24	1.15	60	7.20	96	18.43
25	1.25	61	7.44	97	18.82
26	1.35	62	7.69	98	19.21
27	1.46	63	7.94	99	19.60
28	1.57	64	8.19	100	20.00
29	1.68	65	8.45		
30	1.80	66	8.71		
31	1.92	67	8.98		
32	2.05	68	9.25		
33	2.18	69	9.52		
34	2.31	70	9.80		
35	2.45	71	10.08		