Games and Economic Behavior **28,** 25–54 (1999)
Article ID game.1998.0681, available online at http://www.idealibrary.com on IDEAL®

A Surprise-Quiz View of Learning in Economic Experiments*

Antonio Merlo

Department of Economics and Department of Politics, New York University, 269 Mercer Street, New York, New York 10003

and

Andrew Schotter[†]

Department of Economics, New York University, 269 Mercer Street, New York, New York 10003

Received February 6, 1996

We use experimental data to investigate what subjects have learned after participating in an economic experiment for many periods. At the end of our experiments subjects are given another experiment to perform which functions like a surprise quiz since they had not been informed about its existence before. The results of this surprise-quiz round, along with the time-series of their responses before the quiz, allow us to judge what they have learned. We find that depending on the manner in which subjects are paid, they attempt to learn different aspects of the experimental task and perform differently. *Journal of Economic Literature* Classification Numbers: C91, D83.

1. INTRODUCTION

Consider the following economic experiment. I subjects indexed $i = 1, 2, \dots, I$ are brought into a laboratory and asked to play a game $\Gamma(N, S_j, \Pi_j;\ j \in N)$, where N is the set of players, S_j is the strategy space

*The authors thank Jess Benhabib, Tim Cason, Yaw Nyarko, Roy Radner, Aldo Rustichini, two anonymous referees, and the participants of the Microeconomics Workshop at New York University, the Conference on Individual Rationality and Social Norms, The Economic Science Association Meetings, and the Russell Sage Foundation Conference on Behavioral Economics for their helpful comments. The research assistance of Mark Breton, Vicky Myroni, Sangeeta Pratap, Ken Rogoza, Blaine Snyder, and Tao Wang as well as the financial support of the C.V. Starr Center for Applied Economics are gratefully acknowledged. The usual disclaimer applies.
[†] E-mail: schotter@fasecon.econ.nyu.edu.

0899-8256/99 $30.00

for player j, and Π_j is player j's payoff function.[1] The game Γ is played T times in succession with payoff π_{it} for subject i in period $t = 1, 2, \dots, T$. At the end of the experiment each subject is paid $\Sigma_t \pi_{it}$ and sent home. The question asked by the experimenters is whether the subjects learn to play a particular static equilibrium of the game repeated in the experiment.

In answering this question it is common practice to watch the time series of decisions that subjects make and look for some type of convergence to the predicted equilibrium. If convergence is observed, then the theory is supported; if not, then the theory is one step closer to falsification. But is the convergence of behavior toward equilibrium evidence of learning? If so, what is it assumed these subjects have learned?

In this paper, we use experimental data to investigate what exactly it is that subjects have learned after participating in an economic experiment for 75 periods. It is our hypothesis that depending on the manner in which subjects are paid, they will attempt to learn different aspects of the experiment they are placed in and perform differently.

We call this paper a surprise-quiz view of learning, since at the end of our experiments subjects are given another experiment to perform which functions like a surprise quiz since they had not been informed about its existence before.[2] The results of this surprise-quiz round, along with the time series of their responses before the quiz, allow us to judge what they have learned. To explain, we run a simple 75-round one-person decision-making experiment using two different payoff structures. In one, which we call the Learn-While-You-Earn environment (LWYE), we pay subjects the way they are typically paid by having them play a game 75 times for small payoffs and then pay them the sum of their earnings over the 75 rounds. When this experiment is over, we then give these subjects a ‘‘surprise quiz’’ by having them play the same game one more time for ‘‘big stakes,’’ in fact, payoffs that are 75 times the one round payoffs in the game they have just played. If subjects have learned the structure of the decision problem they were just engaged in, then we would expect them to make a choice in this extra round which is close to the optimum defined by the problem (i.e., at the top of the payoff function defined by the experiment). In short, this decision should be their best guess as to what optimal behavior is here since it is a period in which a significant amount of money is on the table.

[1] If a computer plays the role of N-1 nonstrategic players, leaving only one live subject to make decisions, then the game becomes a one-person maximization problem. Even in that case we will continue to call the situation a game.

[2] One should not consider this surprise equivalent to lying to subjects since we never indicated that the experiment they were performing would be the only task we ask them to do when recruited. We left that possibility ambiguous.

In our other payoff condition, which we call the Learn-Before-You-Earn (LBYE) environment and which we use as our control environment, subjects play the same game and again play it for 75 rounds. However, here they play 74 rounds for free and only in the 75th round do they receive a payoff. The subjects are informed about this at the beginning of the experiment and the payoff they receive in the last round is of the exact size as that of the big-stakes extra round in the Learn-While-You-Earn experiment. Hence, supposedly subjects have the same incentive to choose correctly in the last round of the Learn-Before-You-Earn experiment as they did in the extra round of the Learn-While-You-Earn experiment (abstracting from any income effects we consider to be fairly irrelevant).

What we find is that the decisions made in the Learn-Before-You-Earn environment are significantly better (i.e., closer to the optimal decision defined by the problem) than are the decisions made in the extra round of the Learn-While-You-Earn experiment. More surprisingly, these differences in the last and extra rounds of these experiments cannot be attributed to differences in the data subjects accumulated throughout the experiment before the surprise quiz or last round was administered. The subjects appeared to have the same information at their disposal when they were asked to choose for big stakes. As a consequence, we hypothesize that because of the different payoff environments in which these experiments were run, the subjects chose to learn about different objects, with the subjects in the Learn-Before-You-Earn experiment concentrating on learning the optimal choice for the big-stakes 75th round, which they knew was approaching, and the subjects in the Learn-While-You-Earn experiment concentrating on learning some type of myopic stimulus-response rule not knowing that a big-stakes surprise quiz was coming up. This conjecture is supported by the results of simple non-parametric tests as well as linear regressions run to approximate the subjects' decision rules.

These results have potential consequences for experimental methodology since they indicate that the manner in which subjects are paid may have a direct impact on what they choose to learn about. More critically, they indicate that the theory of learning in markets may need reevaluation since markets place subjects in exactly those payoff environments (LWYE) which are shown here to be the least friendly to learning optimal decisions. Finally, we hope to present surprise quizzes as a potentially useful technique that can be used in other experiments where the task facing agents is to learn the static equilibrium or optimum of the decision situation they are engaged in.

This does not mean, however, that current experimental practices are wrong or have been used incorrectly in the past. For example, take the case of the double-oral-auction experiments which have a long and glorious history. In these experiments convergence and learning are well known stylized facts. Since these are Learn-While-You-Earn environments, how can our criticism be relevant? The answer is that these happen to be good environments for Learn-While-You-Earn payoff structures since such environments tend to focus attention of subjects on the feedback of the experiment and, in the case of the oral auctions, such feedback provides a clear and unambiguous path to the competitive equilibrium—those buyers not making a transaction must bid higher while those sellers not making a transaction must ask lower prices. Hence, Learn-While-You-Earn payoff structures, since they foster learning about feedback or stimulus-response rules, are expected to exhibit convergence in those experiments where the feedback subjects get leads them efficiently to the equilibrium (see Milgrom and Roberts, 1991, for a discussion of how adaptive impulse-response rules can lead to convergence to Nash equilibria). In environments where the feedback is noisy and harder to interpret, as in the experiments performed here and perhaps many others, such a payoff structure would not be optimal.

In this paper we proceed as follows. In Section 2 we explain the experiments performed and the experimental design. In Section 3 we present our results, while in Section 4 we offer a possible explanation of the results. Finally, in Section 5 we offer some conclusions.

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 *et al.* (1987) and Schotter and Weigelt (1992).[3] 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

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

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 of winning the big prize but also imply a higher decision cost.[4] 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. The objective function in this problem is the conditional expected payoff function obtained from our two-person tournament game after restricting the choice of one of the players to be equal to 37.[5]

We consider this experiment to be a good one for our purposes for at least two reasons. First, although it presents subjects with a complete information maximization problem for which the optimal action could be calculated *a priori*, such a problem is 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 problem it involves, this experiment is simple to describe to subjects and to understand. This feature is appealing since it should reduce the noise in the data.

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 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"

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

[5] The analytical derivation of this quadratic function, which is specified in Merlo and Schotter (1995), involves straightforward calculations using the formulas reported in Schotter and Weigelt (1992).

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.

While the LBYE environment is used here strictly as an artificial control environment, it could be given a real-world interpretation. For example, some trading houses give new traders "paper money" training periods in which they make on-paper trades to see how well they do before they start trading using the firm's money. One of the points of this exercise is to explicitly allow traders to make mistakes and learn from them while they do not involve actual losses.[6] Also, consider an infinitely lived firm which interacts repeatedly in a market and has an extremely low (possibly zero) discount rate. In such an environment, the firm 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 two different experiments that were conducted at the Experimental Lab of the C.V. Starr Center for Applied Economics at New York University, using 47 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 these experiments, subjects played the tournament game described above 75 times consecutively against a computer whose strategy was known

[6] We thank one the referees for bringing this example to our attention. The same referee also suggested that many extracurricular activities in high school and college can be interpreted as LBYE experiments. For instance, in student governments, students make very minor decisions with surface features of real political activity such as voting, lobbying, etc.

TABLE I
Experimental Design

Experiment	Opponent	Opponent's strategy	Payoff environment	Number of subjects
1	computer	fixed and known to be 37	LBYE	23
2	computer	fixed and known to be 37	LWYE	24

to be that of choosing 37 in each period. These experiments then presented our subjects with a one-person decision problem under stochastic uncertainty.[7] The experiments were performed under two payoff regimes. In Experiment 1 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 2 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 2 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 2 performed the experiment twice: once for 75 rounds with small payoffs in each round, and once for one extra round (the surprise-quiz 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 LBYE Experiment 1.

Our experimental design is summarized in Table I.

[7] For results of the two-person version of these experiments, see Merlo and Schotter (1994).

3. 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 1 (the LBYE experiment) and the extra round of Experiment 2 (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 and whether they knew there would be a test round. What we find is very different behavior in the last (extra) round of these two experiments despite the fact that subjects appear to have accumulated similar data at the end of the experiments and hence to have the same information at their disposal to guide their final decision.

Tables II and III present the last-round and extra-round choices of subjects in Experiments 1 and 2, respectively, while Fig. 1 presents histograms of the absolute deviations of these choices from the optimal choice of 37. 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 the subjects in the surprise-quiz round of the LWYE experiment. More precisely, while the mean (median) last period choice for subjects in LBYE Experiment 1 was 42.61 (40.00)—5.61 (3) units away from the optimal choice of 37—the mean (median) surprise-quiz round choice for subjects in LWYE Experiment 2 was 51.33 (50.00)—14.33 (13) units away from the optimal choice of 37.[8] In addition, while 11 subjects in the LWYE experiment made surprise-quiz choices which were 30 units or more away from 37, in the LBYE experiment only two such choices were made. A Kolmogorov–Smirnov test of equality of the distributions of last- and extra-round choices rejects the null hypothesis at conventional significance levels (P value 0.03).

In terms of a money metric, the choices of our subjects in the LBYE environment of Experiment 1 led to a mean (median) expected payoff for the experiment of $14.68 ($15.21) which was $1.11 ($1.03) higher than the mean (median) expected payoff of subjects in the extra round of LWYE Experiment 2 of $13.57 ($14.18).[9] In addition, while 13 subjects in LWYE

[8]A median test rejects the null hypothesis that the median of the LWYE group was equal to 37 at the 5% level of significance (P value 0.00), while such a hypothesis cannot be rejected for the LBYE group at the 5% level (P value 0.11).

[9]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.

TABLE II
Last Period Choices—LBYE 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
Median	40.00	15.21	7.00	0.06

[a] An * indicates that the payoff loss is smaller than $0.10.

Experiment 2 had payoff losses (measured by the difference in their expected payoffs when choosing the optimal choice of 37 and their actual last-round choice) greater than $1.00, only 2 subjects in LBYE Experiment 1 had such big losses. Conversely, while 12 subjects in LBYE Experiment 1 had losses of less than $0.10, only 2 subjects had such small losses in the LWYE environment of Experiment 2. The histograms of these losses are presented in Fig. 2.

Finally, looking at the median absolute difference between the last-round choices and the optimal choice of 37 for both groups reveals that while this difference is 28 in the LWYE experiment, it is only 7 in the LBYE experiment. An Epps–Singleton test using all data finds a significant difference in the distributions of absolute deviations from 37 for the two experiments at the 5% level of significance (P value 0.02).

TABLE III
Extra Period Choices—LWYE Experiment 2[a]

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
Median	50.00	14.18	28.00	1.09

[a] An * indicates that the payoff loss is smaller than $0.10.

We use medians instead of means for these comparisons, to help demonstrate that our results are not artifacts of the five subjects in the LWYE experiment whose extra-round choice was either 0 (subjects 5, 13, and 16) or 100 (subjects 3 and 24). In the LWYE experiment, 10 of 24 extra-round choices were dominated choices, i.e., 65 or more—these choices are dominated by 0−and three additional choices were 0. Hence, 13 of the 24 choices were either dominated or 0. On the other hand, in the LBYE experiment only 2 extra-round choices were greater than 65 and none were less than 15. We take this as evidence that subjects' learning capabilities in the LWYE experiment were severely inhibited by their payoff environment since it led to such a high frequency of irrational choices.

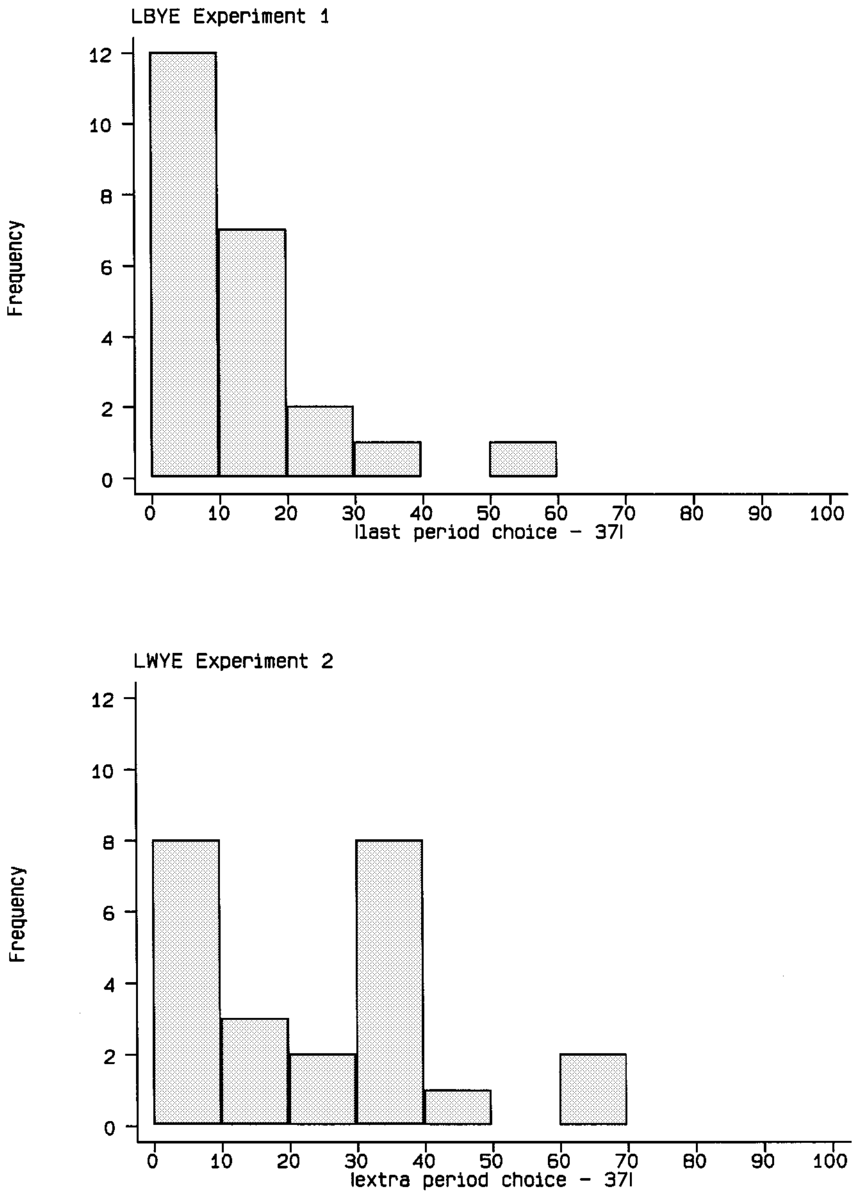

FIG. 1. Histograms of absolute deviations from 37.

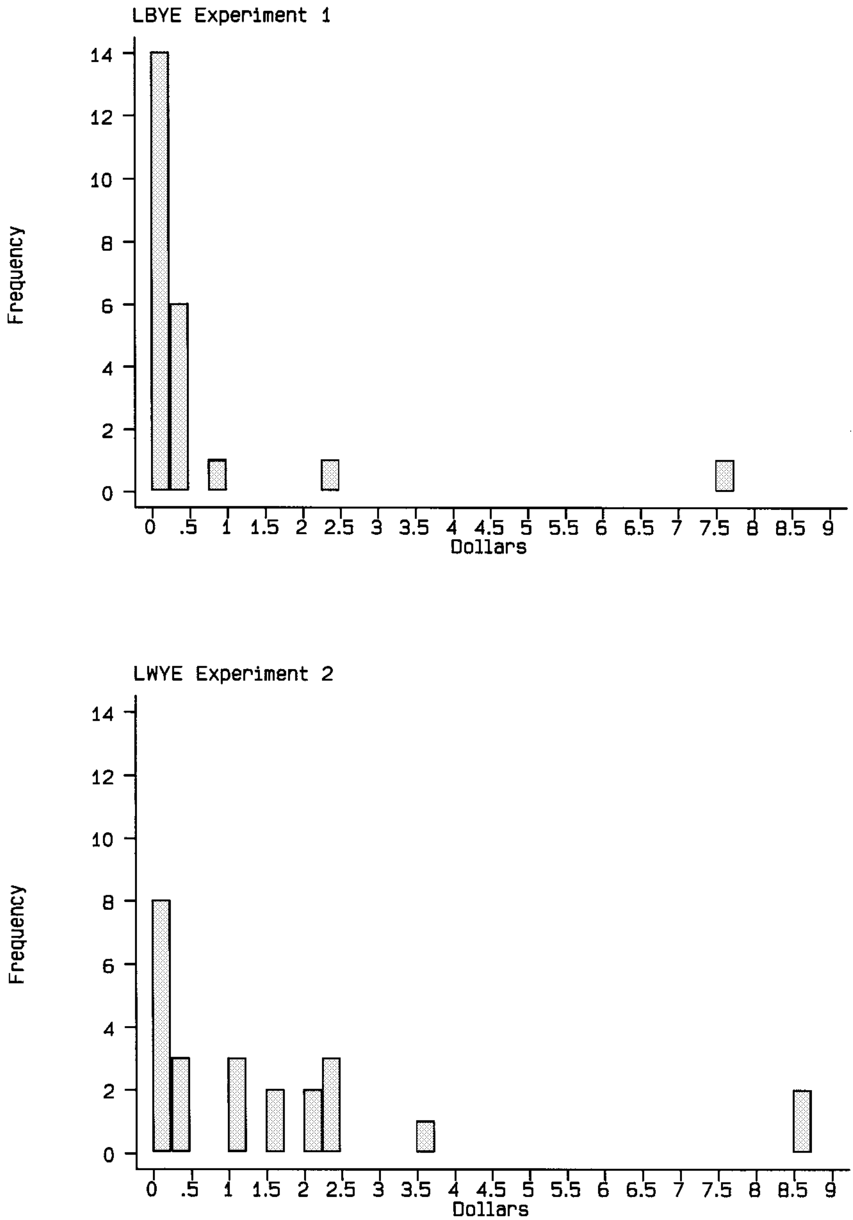

FIG. 2. Histograms of payoff losses.

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 their choices in the last or extra round they do so using very different information sets. Our data reject this interpretation. Our major source of refutation is presented in Tables IV and V which contain descriptive statistics of the distributions of the decision numbers chosen by our subjects in Experiments 1 and 2 over the 74 or 75 rounds of their history prior to making their last choice for big stakes. This was the data at their disposal (i.e., the stock of knowledge they accumulated throughout the experiment) when they made their extra-round or last-round choices. Kolmogorov–Smirnov tests of equality of the distributions (across subjects) of such descriptive statistics in the two experiments find no statistically significant differences at the 5% level between

TABLE IV
Descriptive Statistics of Subjects' Choices—LBYE Experiment 1

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 V
Descriptive Statistics of Subjects' Choices—LWYE Experiment 2

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

the distributions of either the mean, median, standard deviation, or interquantile range of these samples in LBYE Experiment 1 and LWYE Experiment 2 (P values 0.10, 0.06, 0.36, and 0.20, respectively).[10]

These statistics are static snapshots of subject search patterns. However, it might be that the dynamic search sample paths from which these are derived could be quite different with the LWYE subjects narrowing their search over time while the LBYE subjects continue to search broadly or vice versa.

[10] Note that the hypothesis of equality of the distributions of either the mean or the median of the subjects' choices in the two experiments is rejected at the 10% level. The important point, however, is that the hypothesis of equality cannot be rejected at any conventional significance level for the distributions of either measure of dispersion of subjects' choices in the two experiments. This implies that the conjecture that subjects in a LWYE environment sample more conservatively than their counterparts in a LBYE environment for the duration of the experiment is not supported by the data.

To explore this point further we calculated the standard deviation of subject choices separately for each individual over the first, second, and third 25 rounds of their experience. What we find is that while the mean standard deviation over these three 25-round aggregates were 21.06, 15.96, and 13.62 for the LBYE subjects and 17.99, 12.46, and 11.27 for the LWYE subjects, respectively, group differences are insignificant using Epps−Singleton tests run pair-wise on the distributions of standard deviations in each of the three 25-round data sets (P value 0.10, 0.30, and 0.69, respectively, for the first, second, and third 25-round segments of the experiment). Hence, inter-temporal search patterns do not appear to differ across these two groups with no group searching more narrowly than the other.

Still, there is a phenomenon in the LWYE experiment of some subjects opting not to explore their payoff domain either at all or insignificantly. More precisely, in the LWYE experiment seven subjects (subjects 2, 5, 6, 10, 17, 20, and 22) had a mean standard deviation in their choices over the entire 75 rounds of 10 or less.[11] No such subjects exist in the LBYE experiment. Our point, however, concerns comparisons between the LWYE and LBYE subjects who have explored the domain of their payoff function equivalently. We claim that it is precisely for these subjects that learning differentials will be observed. If we eliminate those seven subjects in the LWYE sample who explored their payoff function insufficiently, the mean extra-round choice in the LWYE sample would increase to 53.35 (from 51.33), their average payoff would drop to 12.80 (from 13.57), and their mean absolute difference from 37 would increase to 26.11 (from 24.33). While their average payoff loss would drop from 1.70 to 1.66, this drop is behaviorally insignificant.[12] All of these facts strengthen our point.

Finally, some might suspect that our results are the consequence of the fact that since the LWYE subjects are paid each period while the LBYE subjects are not, they are subject to a wealth effect so that when they arrive at their extra-round choice they have acquired a different attitude toward risk than their LBYE counterparts. This conjecture implies that there should be a correlation between the choice of a subject in the extra round of the LWYE experiment and the amount of money accumulated by subjects in their first 75 rounds. It also implies that, since the extra-round choices of the LWYE group tended to be greater than the last-round choices of the LBYE group, the correlation should be positive (remember

[11] Five of these subjects searched so narrowly that they could not have estimated a quadratic payoff function using the observations they gathered over their 75-round experience (see Table VII).

[12] An Epps−Singleton test rejects the null hypothesis that the distributions of these two modified samples of last- and extra-round choices are the same at the 5% level (P value 0.01).

no money was paid to subjects in the LBYE experiment).[13] Such a positive correlation implies increasing absolute risk aversion on the part of these subjects. In the LBYE experiment no such correlation should exist since, as we just stated, no money was accumulated.

To test this hypothesis we calculated a correlation coefficient for the LWYE and LBYE groups correlating the amount of money they had earned (LWYE) or could have earned (LBYE) prior to making their last- or extra-round choice and the choice they actually made. We find that for the LWYE and LBYE groups, respectively, the correlation coefficients are -0.549 and -0.510. Note that counter to our expectations, the LWYE correlation coefficient is negative and not positive as we predicted. More importantly, however, the two coefficients are extremely close in magnitude, which again is counter to the wealth-effect hypothesis which would predict a difference. Hence, while we expect no wealth effect at all in the LBYE experiment, we actually see a correlation of equal strength to that of the LWYE experiment, indicating that there is no differential impact of wealth on subjects across treatments.

The point to be learned from these correlations is that while two factors might influence the choice of subjects in their last or extra rounds—their wealth and their experience before the last round—only one, their experience, seems to be important.

To summarize, from the data presented in Tables II–V and Figs. 1 and 2 it would appear that subjects in the LBYE environment of Experiment 1 were significantly better at the decision task presented to them than were their counterparts in the LWYE environment of Experiment 2 as measured both by the distance of their last- (extra-) period choices from the optimal choice and by their payoffs. Furthermore, these differences cannot be explained by fundamental differences in the data generated by subjects prior to their last- (extra-) period choices in 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.

4. A POSSIBLE EXPLANATION OF THE RESULTS

Our surprise-quiz and big-stakes rounds test one thing and one thing only, which is how well the subjects in our experiments learn to locate the maximum of the payoff function they faced. The fact that subjects in our LBYE environment did better on this task may simply indicate that that payoff environment, with its big payoff round looming on the horizon,

[13] Clearly, if subjects had utility functions which exhibited constant absolute risk aversion, wealth effects would be unimportant.

focuses attention on this maximum as the correct object of learning. In the LWYE environment, however, where small payoffs come every period, subjects might think it more natural to attempt to learn an impulse-response or adaptive rule and not concentrate on the static payoff function and its maximum (remember, these subjects had no idea that a big payoff experiment was awaiting them in the extra round). In short, what we are saying is that different payoff environments serve as different framing devices which lead subjects to attempt to learn about different aspects of the problem they face. While the LBYE environment focused attention on the static payoff function and its maximum, the LWYE environment focused attention on a more myopic-reactive behavior to the data generated period by period.

To substantiate this conjecture we perform two calculations. To explain our first calculation, note that one way of testing the hypothesis that subjects in the LBYE environment focus their attention on the static optimization problem they face while those in the LWYE environment do not, is to ask whether the subjects in LBYE Experiment 1 use the data they have generated to estimate the quadratic payoff function they face in the experiment and then make a last-round choice which is close to the estimated maximum while subjects in LWYE Experiment 2 do not. Hence, we should observe choices in the last (big-stakes) round of the LBYE experiment closer to the estimated payoff maximizing choices than are the surprise-quiz round choices in the LWYE experiment.

Second, we use the data generated by Experiments 1 and 2 to estimate (reduced form) linear approximations of the decision rules used by subjects in these two experiments. In particular, we regress the decision number chosen by a subject in each round (dec_t^i, $t = 3, 4, \dots, 74$ or 75) on his or her lagged decision numbers ($\mathrm{dec}_{t-\tau}^i$, $\tau = 1, 2, 3$), lagged payoffs ($\mathrm{pay}_{t-\tau}^i$), lagged dummy variables ($\mathrm{win}_{t-\tau}^i$) denoting a win (1) or a loss (0) in the tournament, and a local estimate of the sign of the gradient of the payoff function (slope_t^i).[14] In order to identify behavioral regularities in the subjects' population, we run one such regression for each of the two pooled samples obtained by combining all individual histories in each of the two experiments. To control for individual heterogeneity, we also include among the covariates player-specific dummy variables (α^i).

If subjects in the LWYE experiment were merely responding to the data they were generating in a myopic manner and attempting to learn an appropriate impulse-response rule, then we would expect that they would not look more than one period back in formulating their next period choice and would not have a significant slope coefficient since they were

[14] $\mathrm{slope}_t^i \equiv \mathrm{sign}([\mathrm{pay}_{t-1}^i - \mathrm{pay}_{t-2}^i]/[\mathrm{dec}_{t-1}^i - \mathrm{dec}_{t-2}^i])$.

TABLE VI
Estimated Optimal Choices—LBYE Experiment 1[a]

Subject	Estimated optimal choice	Estimate's standard error	Last period choice	|(4) − (2)|	Estimated payoff loss
1*	43	16.9	40	3	0.03
2*	34	5.9	38	4	0.02
3	52	4.5	58	6	0.36
4	33	5.0	50	17	0.49
5	27	21.3	39	12	0.28
6	34	5.3	96	62	7.95
7	37	4.9	30	7	0.20
8	24	8.1	33	9	0.20
9*	19	11.5	23	4	0.03
10	34	5.4	50	16	0.42
11	24	10.7	42	18	0.64
12	28	8.2	15	13	0.39
13	18	25.1	69	51	4.08
14*	38	4.1	36	2	0.01
15*	35	7.5	32	3	0.02
16	30	20.0	22	8	0.12
17*	39	5.8	42	3	0.03
18	50	4.7	37	13	0.64
19*	45	9.5	50	5	0.05
20	38	7.4	50	12	0.32
21*	34	9.4	38	4	0.03
22*	41	4.2	40	1	0.00
23	42	10.0	50	8	0.16
Average	34.7	9.4	42.6	12.2	0.72
Median	34	7.5	40	8	0.20

[a] An * indicates that the estimated payoff loss is smaller than $0.10.

not attempting to locate the peak of their payoff function. Subjects in the LBYE experiments should behave differently in that their slope coefficient should be significant and they might be more likely to look further back at the data in formulating their next move.

Looking at our first calculation, we present Tables VI and VII which report the maxima of the payoff functions that could have been estimated by the subjects in the LBYE (Table VI) and LWYE (Table VII) experiments using the data they generated before the last or extra round, together with the estimates' standard errors, the subjects' final choices, their absolute differences from the estimated payoff maximizing choices, and the subjects' payoff losses computed using their estimated payoff

TABLE VII
Estimated Optimal Choices—LWYE Experiment 2[a,b]

Subject	Estimated optimal choice	Estimate's standard error	Extra period choice	|(4) − (2)|	Estimated payoff loss
1	46	0.6	65	19	1.12
2	38	20.4	45	7	0.11
3	43	8.6	100	57	8.31
4	32	7.8	77	45	3.96
6*	43	2.1	45	2	0.10
7	31	3.8	41	10	0.39
8	34	36.9	68	34	2.22
9	57	1.6	70	13	0.38
10*	63	8.3	65	2	0.02
11	32	5.1	45	13	0.29
12*	34	7.5	35	1	0.01
13	30	5.1	0	30	1.89
14	37	10.7	44	7	0.15
15	32	4.1	68	36	2.48
17	56	3.4	70	14	3.49
19	32	15.7	50	18	0.17
20*	48	50.7	50	2	0.00
23*	33	5.4	30	3	0.01
24	46	10.7	100	54	6.46
Average	40.4	10.9	56.2	19.3	1.66
Median	37	7.5	50	13	0.38

[a] Subjects 5, 16, 18, 21 and 22 could not have estimated a concave function.
[b] An * indicates that the estimated payoff loss is smaller than $0.10.

functions.[15] As we can see, the last-round choices of subjects in the LBYE experiment were closer to their estimated payoff maximizing choices than were the choices of subjects in the LWYE experiment. For example, while the mean (median) absolute deviation of last-round choices from the estimated peak of the payoff function was 12.2 (8) in the LBYE experiment, it was 19.3 (13) in the LWYE experiment. An Epps-Singleton test finds statistically significant differences in the distributions of absolute deviations from the estimated peak of the payoff function for the two experiments both using all data (P value 0.05) and also after controlling for

[15] The estimates of the coefficients of the payoff function for each subject can be found in Merlo and Schotter (1995). Note that five subjects in LWYE Experiment 2 were eliminated from the subject pool because they could not have estimated a concave payoff function.

outliers—i.e., after excluding subjects 3, 13, and 24 in the LWYE experiment who chose 0 or 100 in their extra rounds (P value 0.05). Also, while 14 out of 19 deviations in the LBYE experiment were less than 10 units, in the LWYE experiment only 7 out of 19 deviations were this small. A Kolmogorov–Smirnov test of equality of the distribution of last-round or extra-round choices and the distribution of estimated optimal choices cannot reject the null hypothesis at conventional significance levels for the LBYE experiment (P value 0.16), while it clearly rejects such a hypothesis for the LWYE experiment (P value 0.02).

In terms of aggregate statistics, both in terms of the mean and the median, the estimated payoff maxima for the LWYE group was remarkably close to the true optimum of 37. In fact, the median optimum was 37. Still on an individual-by-individual basis, an F test rejects the hypothesis that the coefficients of the estimated payoff function were equal to the true (theoretical) payoff function for 13 of the 19 subjects in the LWYE experiment who were capable of estimating a concave payoff function, while this was true for only 9 of 24 subjects in the LBYE experiment.[16] This seems to indicate that despite the aggregate means and medians, subjects in the LBYE experiment did a better job of learning the true parameters of the payoff function they faced.

In terms of our money metric, in Fig. 3 we display histograms of the subjects' payoff losses based on their estimated payoff functions.[17] Note that as it was true in Fig. 2, these losses are on average greater in the LWYE experiment than in the LBYE experiment. For example, 8 subjects had estimated losses greater than $1.0 in the LWYE experiment while only 2 had such large losses in the LBYE experiment. Conversely, only 8/19 (42.1%) of the subjects had losses less than $0.25 in the LWYE experiment while 13/23 (56.5%) of the subjects in the LBYE experiment had such small losses. In short, subjects in the LBYE experiment acted as if they were making choices in the last round which were close to the peak of their estimated payoff function far more consistently than did subjects in the LWYE experiment.

The results of our second calculation are summarized by the following regression results. (To economize on space we present our results without reporting the estimated individual intercept terms.)

[16] See Merlo and Schotter (1995).

[17] Such losses are measured by the differences in subjects' expected payoffs when choosing their estimated optimal choices and their actual last-round or extra-round choices.

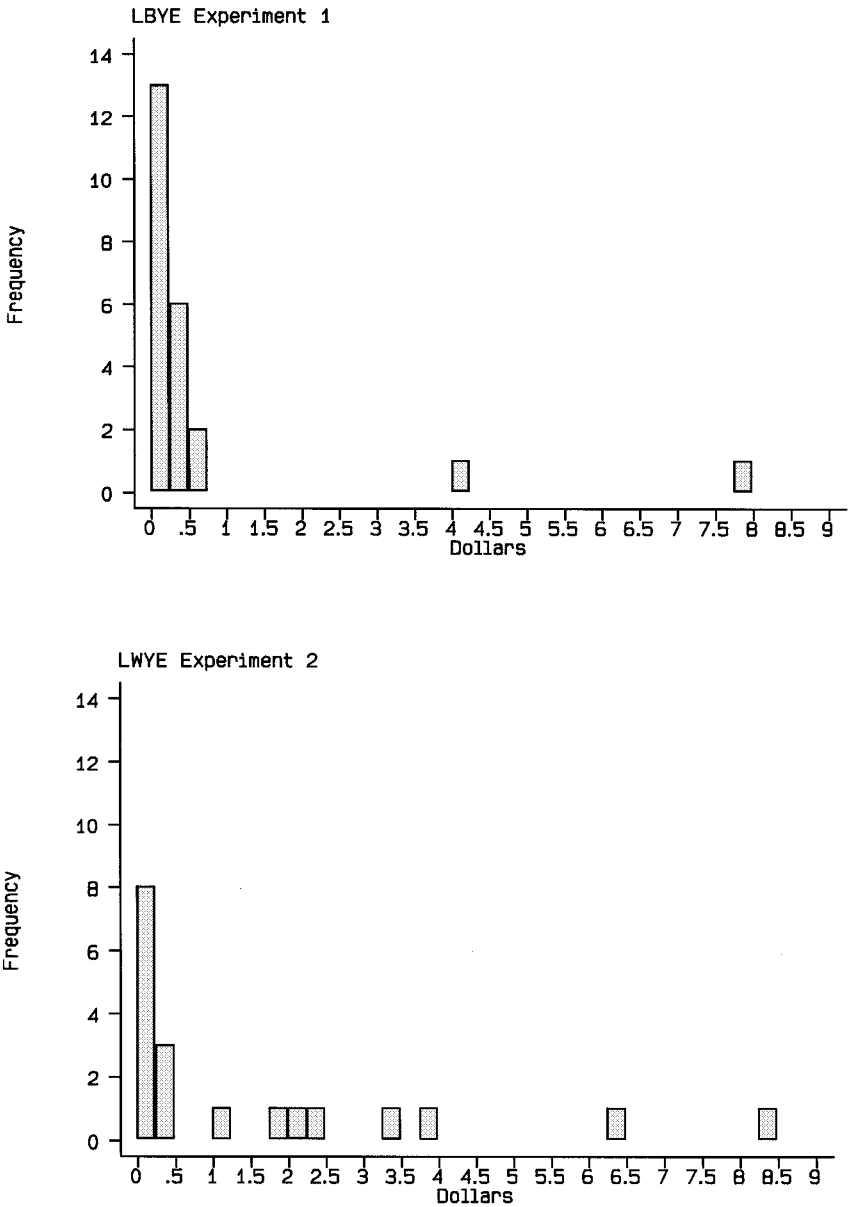

FIG. 3. Histograms of estimated payoff losses.

LBYE Experiment 1: (Number of Observations = 1633)

$$\text{(1)}\quad \text{dec}_t^i = \alpha^i + \underset{(0.128)}{\underline{\mathbf{0.765}}} \cdot \text{dec}_{t-1}^i + \underset{(0.127)}{\underline{\mathbf{0.257}}} \cdot \text{dec}_{t-2}^i - \underset{(0.099)}{0.031} \cdot \text{dec}_{t-3}^i$$

$$+ \underset{(0.631)}{\underline{\mathbf{1.801}}} \cdot \text{pay}_{t-1}^i + \underset{(0.604)}{0.531} \cdot \text{pay}_{t-2}^i - \underset{(0.470)}{0.326} \cdot \text{pay}_{t-3}^i$$

$$- \underset{(7.506)}{\underline{\mathbf{22.412}}} \cdot \text{win}_{t-1}^i - \underset{(7.317)}{7.763} \cdot \text{win}_{t-2}^i + \underset{(5.680)}{2.973} \cdot \text{win}_{t-3}^i$$

$$+ \underset{(0.612)}{\underline{\mathbf{1.683}}} \cdot \text{slope}_t^i, \ R^2 = 0.399.$$

LWYE Experiment 2: (Number of Observations = 1728)

$$\text{(2)}\quad \text{dec}_t^i = \alpha^i + \underset{(0.111)}{\underline{\mathbf{0.471}}} \cdot \text{dec}_{t-1}^i + \underset{(0.113)}{0.102} \cdot \text{dec}_{t-2}^i + \underset{(0.112)}{0.105} \cdot \text{dec}_{t-3}^i$$

$$+ \underset{(0.532)}{\mathbf{1.004}} \cdot \text{pay}_{t-1}^i - \underset{(0.538)}{0.073} \cdot \text{pay}_{t-2}^i - \underset{(0.524)}{0.121} \cdot \text{pay}_{t-3}^i$$

$$- \underset{(6.370)}{\underline{\mathbf{13.461}}} \cdot \text{win}_{t-1}^i - \underset{(6.468)}{0.194} \cdot \text{win}_{t-2}^i + \underset{(6.350)}{0.525} \cdot \text{win}_{t-3}^i$$

$$+ \underset{(0.663)}{1.021} \cdot \text{slope}_t^i, \ R^2 = 0.610.$$

The numbers in parentheses are heteroskedasticity consistent standard errors (White, 1980). Estimates in bold typeface indicate that the coefficients are statistically significant at the 10% level, while estimates that are also underlined indicate that the coefficients are statistically significant at the 5% level.

These results clearly show that the decision rules used by subjects in the two experiments appear to be quite different.[18] In particular, note that in regression (2) (which refers to LWYE Experiment 2) the only coefficients that are significant at either the 5% or the 10% confidence level are the ones associated with one-period lagged variables. In contrast, in regression (1) (which refers to LBYE Experiment 1) the coefficients associated with the slope variable and the two-period lagged decision number are also significantly different from zero at the 5% level. Such findings are consistent with our conjecture that subjects in the LWYE environment only look one period back to determine their decision in any given period and react

[18]An F test of equality of the coefficients in the two regressions (excluding the intercept terms) rejects the null hypothesis at conventional significance levels (P value 0.00).

in a purely adaptive way to its outcome by increasing their choice if either they lose the tournament or if they win with a relatively low decision number—i.e., they receive a ‘‘relatively high’’ payoff. Subjects in the LBYE control environment instead appear to use gradient information about the payoff function and a two-period adjustment rule to guide their search for the peak of the payoff function.[19]

Note that the coefficient associated with win_{t-1} in both regressions is negative and significant. One might think that the sign should be reversed, since reinforcement learning would suggest that when a choice has been rewarded by a win, the subject should increase the use of the choice in the future leading to a positive coefficient. The reason why a negative coefficient makes sense here is because of the fact that higher choices have higher costs associated with them and the decision cost function is convex. More precisely, say that I choose a high number and win. Since costs are convex, reducing my choice a little can be relatively lucrative especially if it does not lead to a big reduction in my probability of winning. Hence, subjects who choose high numbers and win have an incentive to lower their choice in the next period in an effort to economize on decision costs.

An attempt to run these regressions on a more disaggregated level was not successful in demonstrating any drastic differences between these two groups as the aggregate regressions did. This is a common occurrence in experimental data since variation on the individual level is sufficiently noisy as to swamp many of the regularities uncovered when the data is aggregated.

Another important observation is that the differences in the decision rules employed by subjects in the two experiments led to substantial differences in their earnings (with the subjects in LBYE Experiment 1 doing better than their counterparts in LWYE Experiment 2) not only in the last (extra) round of the experiment but also during the experiment. To illustrate this point, in Table VIII we report the average earnings of subjects in the first 75 rounds of LWYE Experiment 2 and the average earnings that subjects in LBYE Experiment 1 would have made in the first 74 rounds if they were paid. As we can see from this table, average ‘‘short-term’’ earnings of subjects in the LWYE environment ($13.87) are smaller than average (hypothetical) ‘‘short-term’’ earnings of subjects in the LBYE environment ($14.50). This finding indicates that the myopic learning of subjects in the LWYE environment is clearly sub-optimal.

[19] While it is logically possible for subjects using a myopic stimulus-response rule to arrive at the peak of the experiment’s payoff function without utilizing information about its slope, such convergence is unlikely.

TABLE VIII
Average Round-by-Round Payoffs

LBYE Experiment 1		LWYE Experiment 2	
Subject	Average payoff (rounds 1–74)	Subject	Average payoff (rounds 1–75)
1	15.09	1	13.63
2	14.30	2	14.98
3	14.03	3	11.14
4	13.31	4	12.64
5	15.77	5	15.50
6	12.69	6	15.49
7	15.22	7	14.34
8	15.31	8	13.04
9	14.48	9	14.55
10	12.63	10	13.42
11	15.46	11	13.73
12	13.61	12	13.38
13	12.80	13	12.94
14	13.62	14	13.70
15	15.53	15	12.27
16	15.22	16	14.87
17	15.27	17	13.53
18	14.42	18	10.96
19	15.10	19	14.59
20	14.45	20	15.63
21	14.73	21	15.86
22	15.18	22	15.17
23	15.30	23	14.73
		24	12.65
Average	14.50	**Average**	13.87

5. 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 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 can be problematic. While such environments can foster convergence when the feedback subjects receive unambiguously pushes them in the direction of the equilibrium (as in the double-oral-auction mechanism), when the feedback inherent in the experiment is noisier, experimentalists might think of using a LBYE payoff structure. Of course, when using a LBYE payoff structure in game-theoretic situations, caution must be used since a ''free-learning'' period might induce subjects to engage in a sort of reputation building that they might not find economically attractive if payoffs accrue in every round.[20]

Furthermore, our experiments imply a simple way for experimentalists to proceed. If the experiment performed requires subjects to learn the properties of a static equilibrium, then the extent to which this learning has occurred can be tested using a surprise quiz as performed here. In short, if you want to know what people have learned, we suggest simply asking them using a surprise quiz. The only other alternative would be to try to infer their learning (or the behavioral rule of thumb they were using in the experiment) using some type of maximum-likelihood technique as developed by El-Gamal and Grether (1994) (see also Cox *et al.*, 1995, for an application of this technique to a problem in learning). While such a technique is quite elegant and useful, we are suggesting that there might be instances where no inference needs to be made since we can elicit the actual rules people are using through an appropriately defined surprise quiz. If such a quiz cannot be formulated, then clearly inference is the only path left.

Our results also 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, 1994). This view has been recently summarized by Milgrom and Roberts (1991, p. 84) 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.''

[20] Depending on the nature of the experiment, this concern may be lessened by adopting an experimental design where the identity of one's opponent changes every period.

In our view, it is not only that people learn how to learn, but they also learn what to learn about and what they learn about is institution specific. If learning is situation or institution dependent, however, it 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.

APPENDIX: INSTRUCTIONS FOR EXPERIMENT 1

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’’ (Table IX). 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 decision 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 into Dollars at the rate of $0.75 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.

TABLE IX
Decision Cost

A Decision number	B Cost of Decision (Francs)	A Decision number	B Cost of Decision (Francs)	A Decision number	B Cost of Decision (Francs)	A Decision number	B Cost of Decision (Francs)
0	0.00	26	1.35	52	5.41	78	12.17
1	0.00	27	1.46	53	5.62	79	12.48
2	0.01	28	1.57	54	5.83	80	12.80
3	0.02	29	1.68	55	6.05	81	13.12
4	0.03	30	1.80	56	6.27	82	13.45
5	0.05	31	1.92	57	6.50	83	13.78
6	0.07	32	2.05	58	6.73	84	14.11
7	0.10	33	2.18	59	6.96	85	14.45
8	0.13	34	2.31	60	7.20	86	14.79
9	0.16	35	2.45	61	7.44	87	15.14
10	0.20	36	2.50	62	7.69	88	15.49
11	0.24	37	2.74	63	7.94	89	15.84
12	0.29	38	2.89	64	8.19	90	16.20
13	0.34	39	3.04	65	8.45	91	16.56
14	0.39	40	3.20	66	8.71	92	16.93
15	0.45	41	3.36	67	8.98	93	17.30
16	0.51	42	3.53	68	9.25	94	17.67
17	0.58	43	3.70	69	9.52	95	18.05
18	0.65	44	3.87	70	9.80	96	18.43
19	0.72	45	4.06	71	10.08	97	18.82
20	0.80	46	4.23	72	10.37	98	19.21
21	0.88	47	4.42	73	10.66	99	19.60
22	0.97	48	4.61	74	10.95	100	20.00
23	1.06	49	4.80	75	11.25		
24	1.15	50	5.00	76	11.55		
25	1.25	51	5.20	77	11.86		

(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.

REFERENCES

Bull, C., Schotter, A., and Weigelt, K. (1987). "Tournaments and Piece Rates: An Experimental Study," *J. Polit. Econ.* **95**, 1–33.

Cox, J., Shachat, J., and Walker, M. (1995). "An Experimental Test of Bayesian vs. Adaptive Learning in Normal Form Games," University of Arizona, unpublished manuscript.

El-Gamal, M. A., and Grether, D. M. (1994). "Uncovering Behavioral Strategies: Likelihood-Based Experimental Data-Mining," California Institute of Technology, unpublished manuscript.

Lazear, E. P., and Rosen, S. (1981). "Rank Order Tournaments as Optimal Labor Contracts," *J. Polit. Econ.* **89**, 841–864.

Merlo, A., and Schotter, A. (1994). "An Experimental Study of Learning in One and Two-Person Games," New York University, C. V. Starr Center for Applied Economics, Research Report No. 94-17.

Merlo, A., and Schotter, A. (1995). "A Surprize-Quiz View of Learning in Economic Experiments," New York University, C. V. Starr Center for Applied Economics, Research Report No. 95-32.

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

Mookherjee, D., and Sopher, B. (1994). "Learning Behavior in an Experimental Matching Pennies Game," *Games Econ. Behav.* **7**, 62–91.

Schotter, A., and Weigelt, K. (1992). "Asymmetric Tournaments, Equal Opportunity Laws and Affirmative Action: Some Experimental Results," *Q. J. Econ.* **106**, 513–539.

White, H. (1980). "A Heteroskedasticity-Consistent Covariance Matrix Estimator and A Direct Test for Heteroskedasticity," *Econometrica* **48**, 817–838.