ECONOMIC RESEARCH REPORTS

EXPERIMENTATION AND LEARNING IN LABORATORY EXPERIMENTS:

Harrison's Criticism Revisited

by

Antonio Merlo

and

Andrew Schotter

R.R.# 90-23

June, 1990

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

Experimentation and Learning in Laboratory Experiments: Harrison's Criticism Revisited *

by

Antonio Merlo and Andrew Schotter

Department of Economics
New York University

* The authors would like to thank the participants of the NYU Informal Theory Workshop for their comments. In addition, Roy Radner, Yaw Nyarko and Franco Peracchi also offered helpful advice. The financial support of the C.V. Starr Center for Applied Economics at NYU as well as National Science Foundation Grant SES-871277 is gratefully aknowledged.

Abstract

In a recent paper Glenn Harrison (1989) has raised an interesting and significant point concerning the methodology of experimental economics. Specifically, Harrison has warned economists to take care in designing their experiments to give subjects sufficient incentives to overcome the significant calculation and decision costs that exist in experiments in order not to loose control over their actions. One way that Harrison claims that control can be lost is if the payoff function faced by any subject is flat, conditional on the others taking their equilibrium actions. Harrison's ideas, while undoubtedly correct, have relatively more of less force depending on the model one uses to describe how experimental subjects go about playing experimental games with monetary payoffs. In this paper we define two types of experimental subjects -- "experimenters" who learn about the game being played and its payoff function by choosing actions that are very different from each other during the experiment and observing their payoff at each of these actions or choices, and "theorists" who experiment very little by choosing actions far away from their theoretical best choice. Our claim is that the Harrison criticism is not applicable when experimental subjects behave as theorists, since their lack of experimentation will not permit them to experience the shape of the payoff function. Furthermore, even if experimental subjects act as experimenters, it is still possible that the Harrison criticism may not have strength since it may be observationally impossible for them to realize that the payoff function they face is flat. We explore these ideas using data generated by the experiments performed by Bull Schotter and Weigelt (1987), one claim from which was recently criticized by Drago and Heywood (1989) using the Harrison criticism.

Section 1: Introduction

In a recent paper Glenn Harrison (1989) has raised an interesting and significant point concerning the methodology of experimental economics. Specifically, Harrison has warned economists to take care in designing their experiments to give subjects sufficient incentives to overcome the significant calculation and decision costs that exist in experiments. If these costs are not overcome the fear is that the investigator may lose control over the actions of his subjects. One way that Harrison claims that control can be lost is if the payoff function faced by any subject, conditional on the others taking their equilibrium actions, is flat. Such a "flatness" would give any player little incentive to find his best response since in a money metric non-optimal decisions would pay very little less than the truly optimal decision. Hence good experimental design would require that payoff functions be "steep" around the equilibrium or optimal action. A recent paper by Kagel and Roth (1990) supports this view[1].

Harrison's ideas, while undoubtedly correct, have relatively more of less force depending on the model one uses to describe how experimental subjects, and by this we mean college undergraduates, go about playing experimental games with monetary payoffs. For example, Harrison's criticism implies that subjects are "experimenters" in their approach to their laboratory decision tasks. His model of experimental subject behavior must be one in which subjects read the instructions, are not capable of understanding them totally, and hence proceed in a boundedly rational manner to choose

---

[1] The problem of offering experimental subjects sufficient incentives is already considered in Vernon Smith's (1982) paper; Smith, however, does not directly address the question of the slope of the payoff function.

different actions until he or she discovers the shape of the payoff function they face. At this point this experimentation stops and an action is chosen. When the payoff function is flat, people will ultimately learn that their choice has little influence on their payoff and hence we can expect a wide variety of behaviors all of which yield relatively the same payoff. When the payoff function is steep, we can expect that subjects will learn this through their experimentation and hence will converge to the Nash equilibrium decision underlying the experiment, since steep payoff functions imply severe penalties for non-optimal choices and the rewards offered are enough to overcome the mental costs of calculation.

We say this must be the model Harrison has in his mind because if all agents were fully rational and fully capable of making the calculations called for in the experiment, then, except for conjectural or coordination problems existing in experiments with multiple equilibria, they would all choose the optimal decision. Hence either they are experimenters as we have described them or they are "theorists" but not very good ones. By this we mean the following. Say that experimental subjects come into a lab, read the instructions, think they have figured out the "right" thing to do in the experiment, and do it for the entire length of the laboratory session. In short, say they experiment very little by choosing actions far away from their theoretical best choice. In this case, they would not experience or discover the relative flatness or steepness of the payoff function at all and hence it would be very unlikely that such considerations could influence their choice. Hence, if subjects are theorists and not experimenters during a laboratory session, the shape of the payoff function will not be important since they will never experience it. In such cases, we would expect that Harrison's criticism would hold less force.

When subjects act like theorists, the complexity of the calculation problem posed by the experiment would determine how dispersed the actions of the subjects are. For example, if an experiment posed a trivial problem for subjects to solve (i.e. an obvious dominant strategy for all subjects), we would expect that all subjects would solve the problem in the first round of the experiment. No variance in behavior would be observed. When the problem posed is complex and subjects are theorists, we would expect greater variance since most subjects will not be able to solve the problem correctly and hence they will reach different conclusions as to their optimal action.

Even if the subjects acted as experimenters, however, it is still possible that the Harrison criticism may not have strength since it may be observationally impossible for the subjects to realize that the payoff function they face is flat. More precisely Harrison's criticism is based on the ex ante expected payoff function that a subject would face if all other subjects chose their Nash equilibrium action in the game defined by the experiment. If subjects are not sufficiently sophisticated theorists, however, they could not calculate this ex ante payoff function. Rather they will only <u>experience</u> the payoff function by choosing actions in the experiment and observing what happens. Hence depending upon the way the subject experimented (and by the realization of the random variables if nature is a player in the experiment), the data generated will differ greatly. Empirically it may be impossible, given the meager data generated by the experiment for subjects to conclude that in fact the payoff function they face is flat. We will show the difficulties involved here later. Hence even if subjects did function like experimenters, they may not be able to estimate the curvature of the ex post payoff function they face in a way that will allow them to reject the hypothesis that their payoff function is steep.

In this paper we explore these ideas further using data generated by the experiments performed by Bull Schotter and Weigelt (1987), one claim from which was recently criticized by Drago and Heywood (1989) using the Harrison criticism. It will be our claim here that this criticism is not applicable to that paper and perhaps to others because, as we will show, subjects often behave as if they were theorists and not as the experimenters required by Harrison (1989) in his criticism[2]. It is important to point out, however, that we are using the Bull et al. (1987) paper merely as an example of a larger point about Harrison's criticism and not as a rebuttal to Drago and Heywood (1989).

In this paper we will proceed as follows: In Section 2 we will summarize the claim of Bull et al. (1987) which was criticized by Drago and Heywood (1989). In Section 3 we present the data from Bull et al. (1987) which illustrate our main point that since subjects in that experiment behaved as theorists, the experiment is not subject to the Harrison criticism. In Section 4 we investigate the difficulties involved in empirically assessing the curvature of the payoff function faced by experimental subjects. In Section 5 we offer some conclusions and comments for future work.

Section 2: Bull, Schotter and Weigelt (1987) and the Harrison Criticism

Bull, Schotter, and Weigelt (1987) present an experimental analysis of two-person rank order tournaments. In these experimental tournaments each subject chose a number between 0 and 100 representing their effort level in

---

[2] It would be interesting to apply the same analysis we use here to study the data of the Cox, Smith and Walker (1983; 1985; 1988) and Cox Roberson and Smith (1982) papers to which Harrison's criticism was originally aimed.

the tournament. A monotonically increasing cost is associated with each effort level. After these efforts are chosen each is augmented by an additive shock or realization of a random variable drawn independently for each subject from a uniform distribution. The rules of the tournament dictate that the subject with the highest total of effort plus random realization will be awarded a "big" monetary prize while the one with the lowest will be awarded a "small" monetary prize. The final payoff for subjects will be the prize they win minus their cost of effort. Each such tournament is run twelve times for each pair of subjects who play against each other repeatedly. Subjects are then paid the sum of their payoffs in each of the twelve tournaments. Bull et al. (1987) chose parameters for which the unique Nash equilibrium of their tournaments were 37.

In their analysis of the results, Bull et al. (1987) observe that while the mean effort chosen by the subjects converges to the theoretical Nash equilibrium of the game defined by these tournaments, there is a wide variance across subjects in the final round of the experiments and no tendency shown for this variance to decrease as the experiment progresses. Such a large variance was not observed when subjects performed in a piece-rate experiment which had an identical optimal effort choice as the tournament but presented the subjects with a one-person maximization problem rather than a two-person non-cooperative game, even though the game they played was the simplest game analogue to the piece-rate optimization problem. The mystery created by these experiments is why does the mean effort level observed here reliably move towards their theoretical levels while the variance remains large (the theory would predict a zero variance), and why are these results different for games as opposed to one-person optimization problems?

One possible explanation of this variance is that subjects in games have trouble inferring either what their partners have done in the past or are about to do in the future (i.e. they may have either inferential problems or conjectural problems), and since these problems are not present in one-person optimization problems, that could explain the qualitative difference in the results. To investigate this hypothesis Bull et al. (1987) had some subjects play against a computerized partner called an automaton who was programmed to play the Nash equilibrium effort level in each period and this fact was told to the subjects. Hence while they theoretically were playing a game, subjects were in fact performing a one-person optimization problem since the inferential and conjectural problems mentioned above were absent in this experiment. Bull et al. (1987) claim that the fact that this automaton experiment still exhibited significantly greater variances than the associated piece-rate experiment demonstrates that after the conjectural and inferential components of the observed variance in behavior is eliminated, the game-like payoff structure of the automaton experiment presented subjects with a computationally more complex problem then the piece-rate and therefore the remaining observed variance was a function of the increased complexity of games over piece-rates.

This hypothesis was criticized by Drago and Heywood (1989) who claim, using the Harrison criticism, that the real reason for the increased variance of behavior in the automaton experiment as opposed to the piece-rate experiment is the flatness of the ex ante payoff function around the equilibrium point. The ex ante expected payoff functions for the piece-rate and the automaton experiment are depicted in Appendix A. Hence, they claim that if the automaton payoff function is flat there is little incentive for subjects to zero in on the optimal choice since choices far from the optimum in the "effort metric" are not far from the optimum in the "money metric". Therefore, subjects chose all over the place. To demonstrate their point Drago and Heywood (1989) re-run the Bull et al. (1987) experiments with parameters which made the steepness of the ex ante payoff function equivalent across their piece-rate and automaton experiments. This treatment significantly decreased the variance observed by their subjects as the following table indicates:

TABLE 1

EXPERIMENTAL RESULTS: MEANS AND VARIANCES

Experiment	Mean Decision Number		Mean Variance in Decision Numbers		Mean Decision Number (Round 12)	Variance in Decisions Numbers (Round 12)
	Rounds 1 - 6	Rounds 7 - 12	Rounds 1 - 6	Rounds 7 - 12		
	1.a) Results from Bull, Schotter and Weigelt (1987)					
Piece-rate	40.44	38.91	103.61	87.38	37.38	33.66
Automaton tournament	44.25	42.07	325.96	265.49	40.41	275.30
	1.b) Results from Drago and Heywood (1989)					
Piece-rate	39.32	37.21	206.78	98.61	37.19	51.31
Automaton tournament	38.49	36.48	163.70	61.56	37.46	50.10

Based on this table, Drago and Heywood (1989) conclude that the variance in the Bull et al. (1987) experiments was an artifact of the flatness of the payoff function there.

Section 3: Experimenters vs Theorists in Laboratory Experiments

In order to operationalize our definition of an experimenter and theorist we will first define the notions of an experiment[3] and experimenter. Anyone not defined as an experimenter will be classified as a theorist. Intuitively, an experimenter will be someone who learns about the game being played and its payoff function by choosing actions that are very different from each other during the experiment and observing his payoff at each of these actions or choices. Experimenters are empiricists who learn from trial and error. To formalize this consider a set of laboratory experiments run with J independent sets of subjects. In each experiment the experimental stage games are repeated T times and the I_j players in each experiment j $(j=1,\dots,J)$ have simply to choose a number in each period t $(t=1,\dots,T)$. Let x_{it} denote the choice of player i at time t (there are JxT such observations); let $\bar{x}_i$ denote the average choice of a player i over T rounds of the experiment he partakes in and let σ_i denote the standard deviation of player i's choices over these same T rounds. The mean standard deviation of all of our subjects is

$$\bar{\sigma} = (1/I_1+\dots+I_J) * \Sigma_i\sigma_i,$$

where this mean is taken by averaging the standard deviations of all players in our J experiments. We will say that player i is making an experiment in period t if he chooses an action sufficiently far away from his mean action over the course of the experiment. We operationalize this by saying that agent i in period t is making an experiment if,

---

[3] Note that the definition of an experiment we use here is not the same as that used by Fudenberg and Kreps (1988), who define an experiment as an out-of-equilibrium action in the play of a repeated game.

$|x_{it} - \bar{x}| > 1.5 * \bar{\sigma}.$

In words, an experiment for player i at time t is a choice x_{it} whose absolute deviation from his mean choice over the T periods is greater than one and a half times the average standard deviation across all of the players involved in the laboratory experiments we are considering. We compare each choice to the mean standard deviation of all subjects to standardize the deviation. If we did not do this we might run the risk of labelling a subject with a very small σ_i an experimenter since more than 1/4 of his choices may be more than $1.5\sigma_i$ units away from their mean despite the fact that absolutely the choices are all almost identical. Likewise, our normalization allows us to avoid the problem of not labelling a subject with a large σ_i an experimenter because although he made widely diverse choices, his σ_i was so large that not more than 1/4 of the choices fell 1.5 standard deviations away from their mean (the standard deviation of a subject may be determined by 1 or 2 huge deviations). We will call an experimenter a player who makes int(T/4) or more experiments in a T-period repeated laboratory experiment - where int(y) is the largest integer $n \leq y$. A theorist is a player who makes less than int(T/4) experiments in a T-period repeated laboratory experiment.

3.1: Are Subjects experimenters or Theorist?

If the Harrison's criticism is to be in force, it must be that subjects have experimented enough with the payoff function to discover its shape. To demonstrate that this is not the case let us first present histograms indicating the number of instances where any of our subjects in either the piece-rate or the automaton experiment chose an action which was a given distance from their mean action.

In figures 1a and 1b we see the deviations from the mean actions of subjects when all observations are pooled over subjects in each experiment. As

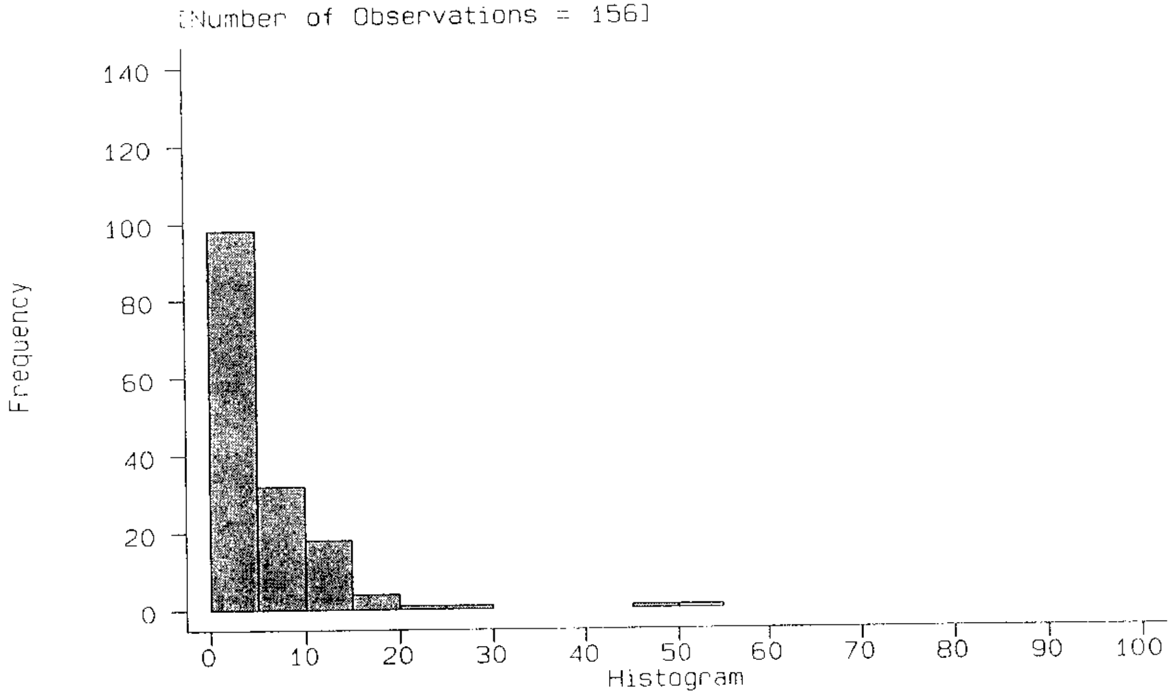

FIGURE 1.a - Piece-Rate Experiment

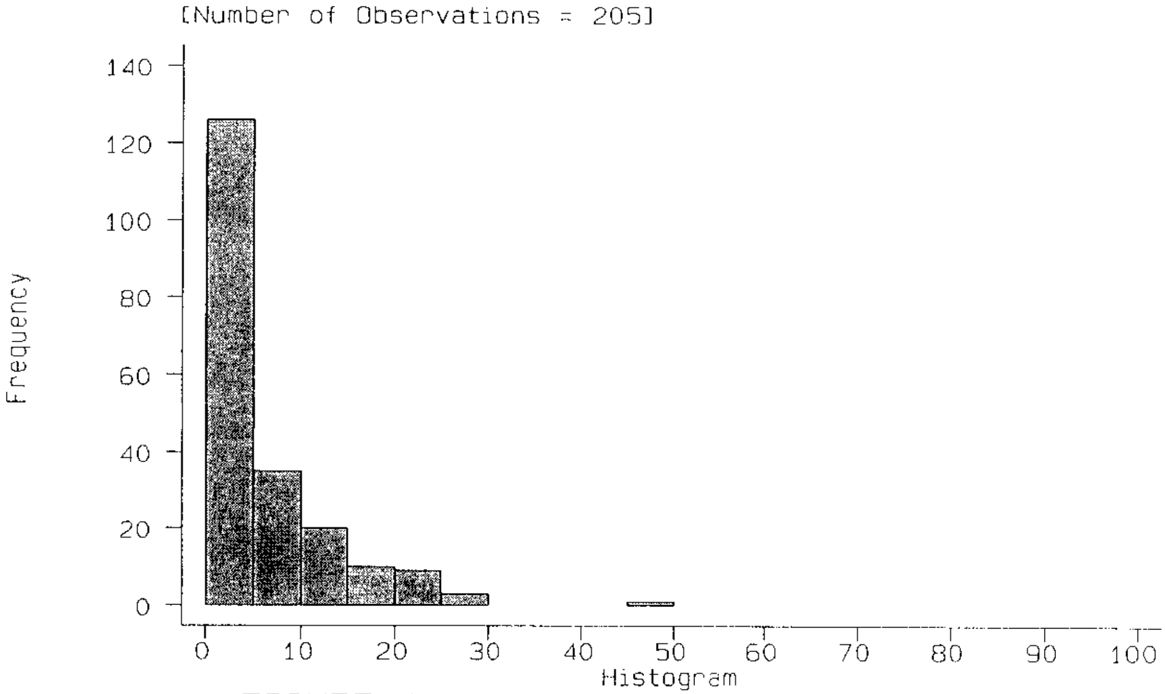

FIGURE 1.b - Automaton Experiment

FIGURE 1

Absolute Differences of Individual Choices
from Their 12 Period Mean

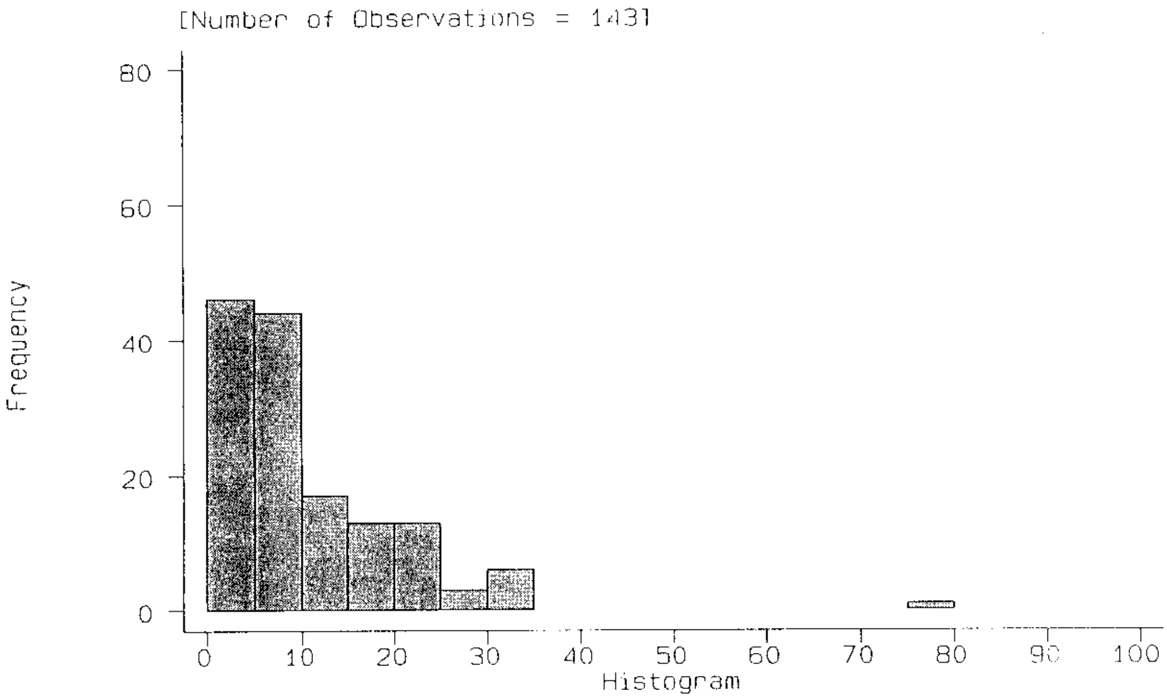

FIGURE 2.a - Piece-Rate Experiment

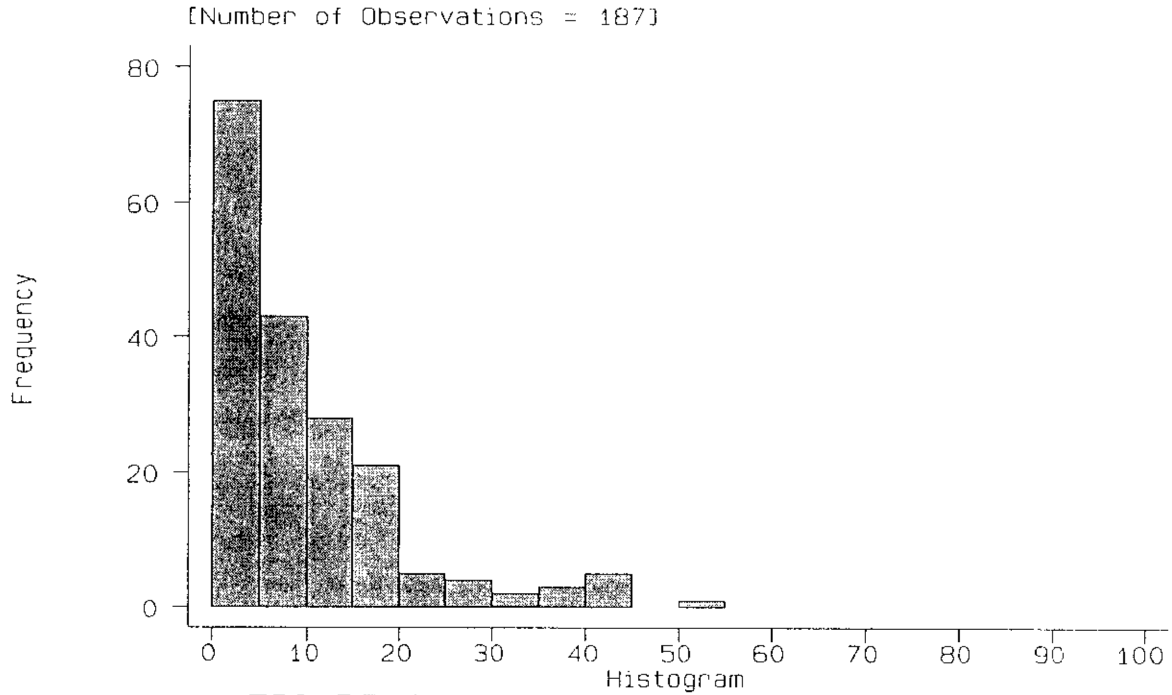

FIGURE 2.b - Automaton Experiment

FIGURE 2

Absolute Differences of Individual Choices
from Their First Period Choice

we see in these figures, in both the piece-rate and the automaton experiment, subjects tended to choose an action and persist in that action throughout the entire experiment, never deviating greatly. Note that of the 156 (204) choices made in the piece-rate (automaton) experiment, 115 (160) were not more than 10 units away from the subject's mean choice over the 12 rounds of the experiments (85 (125) of them were not more than 5 away from their mean). As it turns out 10 happens to be approximately 1.5 mean standard deviations so as we will see, these histograms reinforce the idea that subjects in these experiments were not experimenters.

If subjects were truly theorists, their first period action in the experiment would indicate their theoretical best guess as to what the optimal thing to do in the experiment was since it is untainted by the experience they gain during the laboratory session. If we look at the histograms generated by taking the deviations from a player's first-period choice (figures 2a and 2b), we see a similar pattern as that depicted in figures 1a and 1b.

As figures 2a and 2b indicate, of the 143 (187) choices made by all subjects from round 2 to round 12 of the piece-rate (automaton) experiment 84 (110) were not more than 10 units away from their first period choices (46 (75) were less than 5 units away). This demonstrates that subjects did not really experience the shape of the payoff function since they were making choices in a fairly circumscribed neighborhood of their first period choice. One explanation of this is that subjects are theorists who construct a model of the experiment in their mind before they choose a first period action. This model predicts an optimal choice and an expected payoff and, except for some minor adjustment, radical movements from this theoretically optimal choice are taken only when the payoff received from the experiment deviates substantially from the expected one. When a problem is easy to solve, as is the piece-rate problem, most subjects make choices close to the optimum and hence the observed last period variance is small. As the problem increases in complexity, theorists begin to differ about what the optimal choice is and hence the variance of the last round choices increases. In this sense we feel that our piece-rate and automaton experiments were nested in their level of complexity and hence serve well in testing this hypothesis.

In table 2 we present this argument in a slightly different manner using the notion of an experiment and experimenter.

TABLE 2

EXPERIMENTATION IN THE BULL ET AL. (1987) EXPERIMENTS

2.a) Piece-rate experiment

players	1	2	3	4	5	6	7	8	9	10	11	12	13	mean
# of experiments	0	6	1	0	1	2	0	1	0	4	0	4	1	1.54

2.b) Automaton experiment

players	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	mean
# of experiments	1	2	1	0	2	0	0	1	1	0	5	6	5	7	2	0	1	2.00

In table 2 we see the number of rounds for each subject in which he could be labeled an experimenter. For example, in the piece-rate experiment we see that out of the 13 subjects who participated in the experiment only subjects 2, 10 and 12 (who made 6, 4 and 4 "experiments" respectively) can be considered "experimenters". All of the other subjects are labeled theorists using our criteria. In the automaton experiment there are 4 out of 17 subjects labeled experimenters (subjects 11, 12, 13,and 14). Note also that the mean number of experiments in the piece-rate experiment was 1.54 while it was 2.0 in the automaton experiment.

Again, the interpretation here is clear. Subjects do not experiment enough and over a sufficient range of the payoff function to discover the shape of the payoff function they face. Hence it can not be the shape of the payoff function which determined their actions.

Section 4: Estimating the Shape of the Payoff Function

It was our point above that the overwhelming majority of the subjects in the Bull et al. (1987) experiments were what we called theorists and not experimenters. Hence since theorists sample the shape of the payoff function only locally, they never become aware of its global shape. But what about the few subjects whom we labeled as experimenters? Were they able to successfully discover the shape of the payoff function they faced? We think not and to demonstrate our doubts we present figures 3, 4 and 5. In figure 3.a we present the observed data points for each of our 3 experimenters in the piece-rate experiment where the data points indicate the choice made and the resulting payoff. Figure 3.b does the same for the 4 subjects we designated as experimenters in the automaton experiment. In each we also draw the best quadratic approximation of the data points for a subject. In figure 4 we draw for each of our experimenters in the piece-rate and in the automaton experiments the ex ante expected payoff function for a subject at the theoretical equilibrium and the observed ex post payoff function based on the best quadratic approximation of the data points. In figure 5 then, we

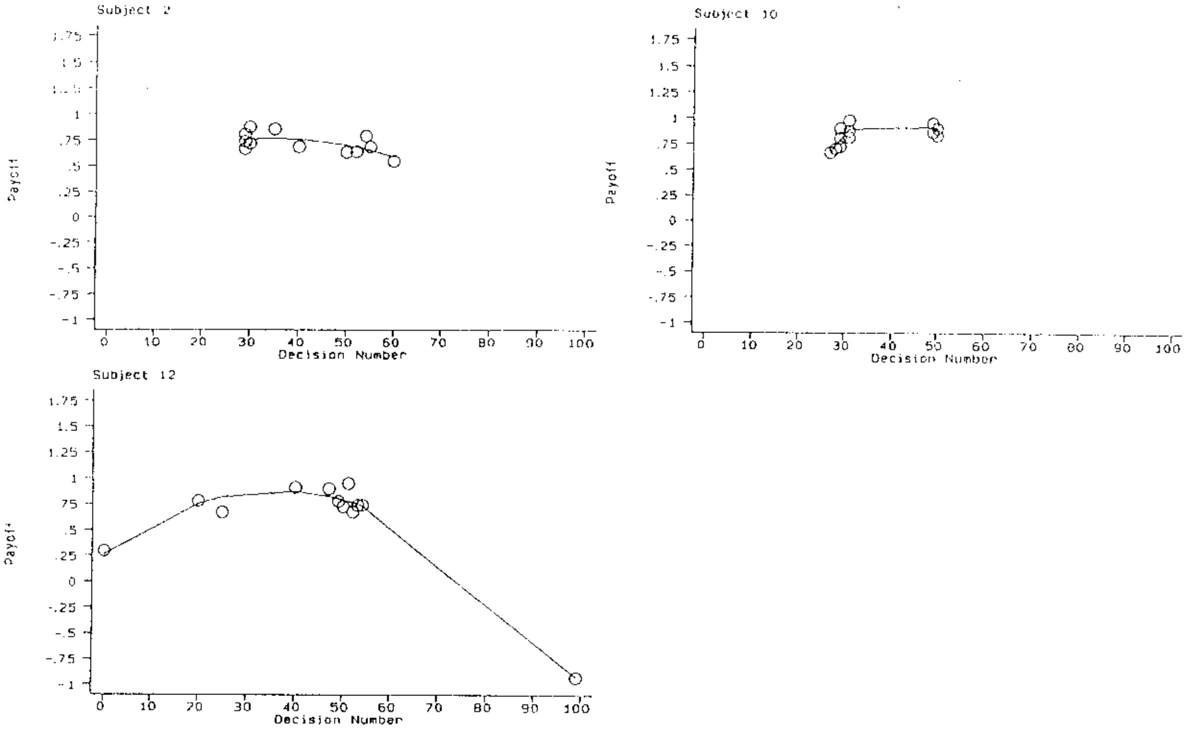

FIGURE 3.a — Piece-Rate Experiment

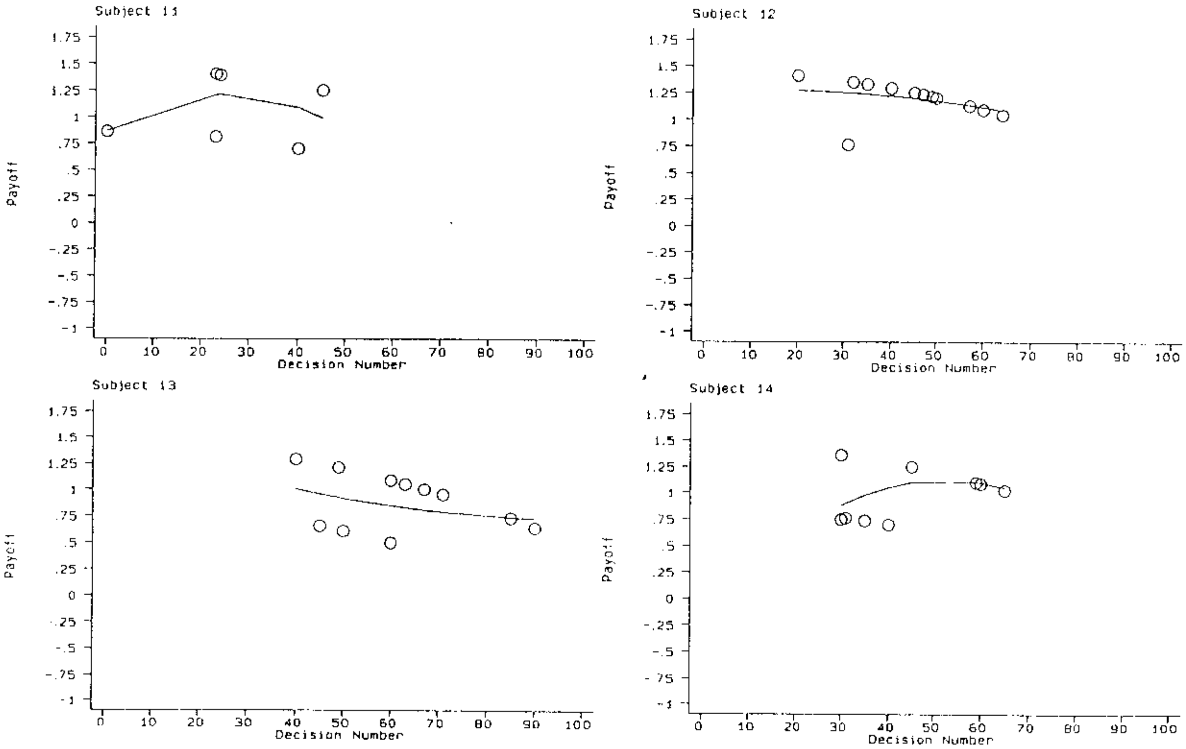

FIGURE 3.b — Automaton Experiment

FIGURE 3

Realized Payoffs

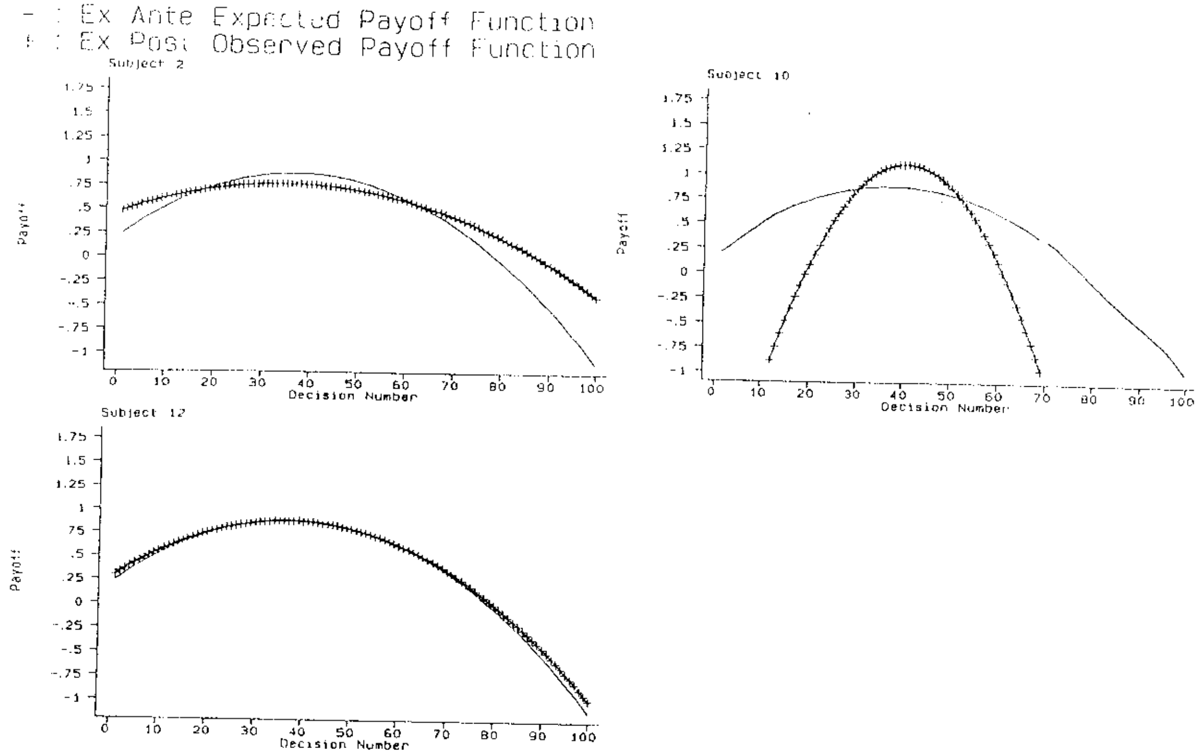

FIGURE 4.a – Piece-Rate Experiment

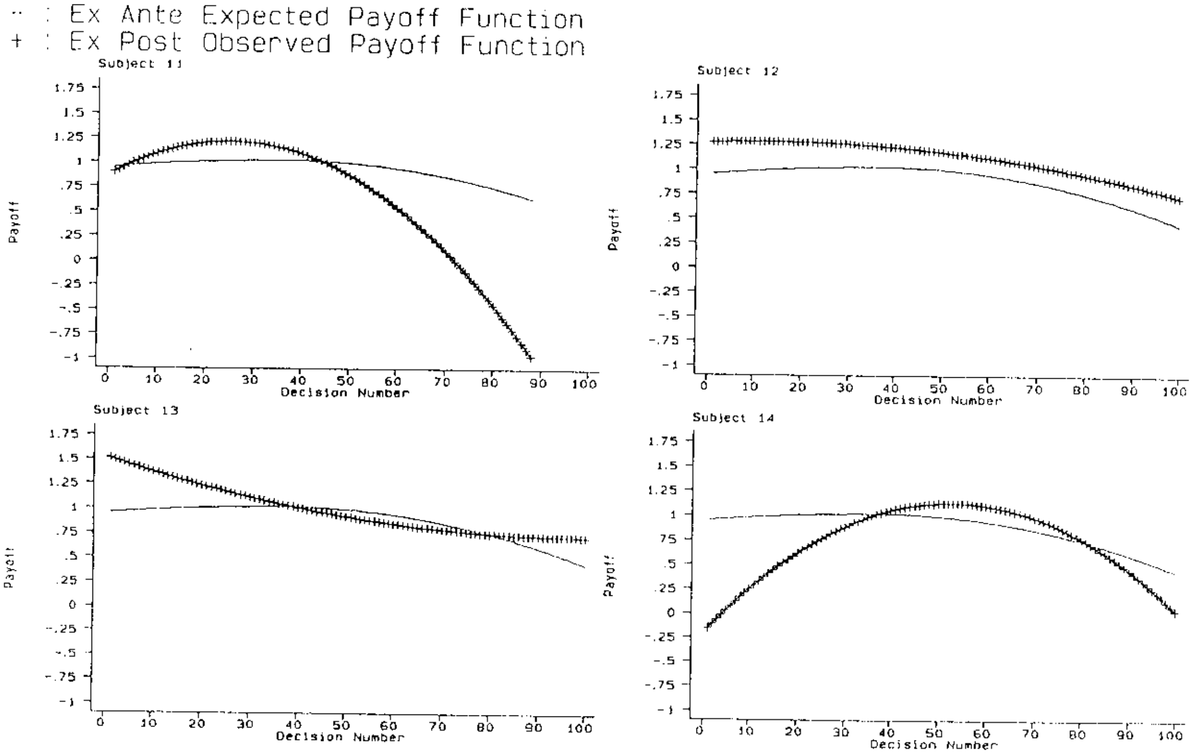

FIGURE 4.b – Automaton Experiment

FIGURE 4

Ex Ante & Ex Post Payoff Functions

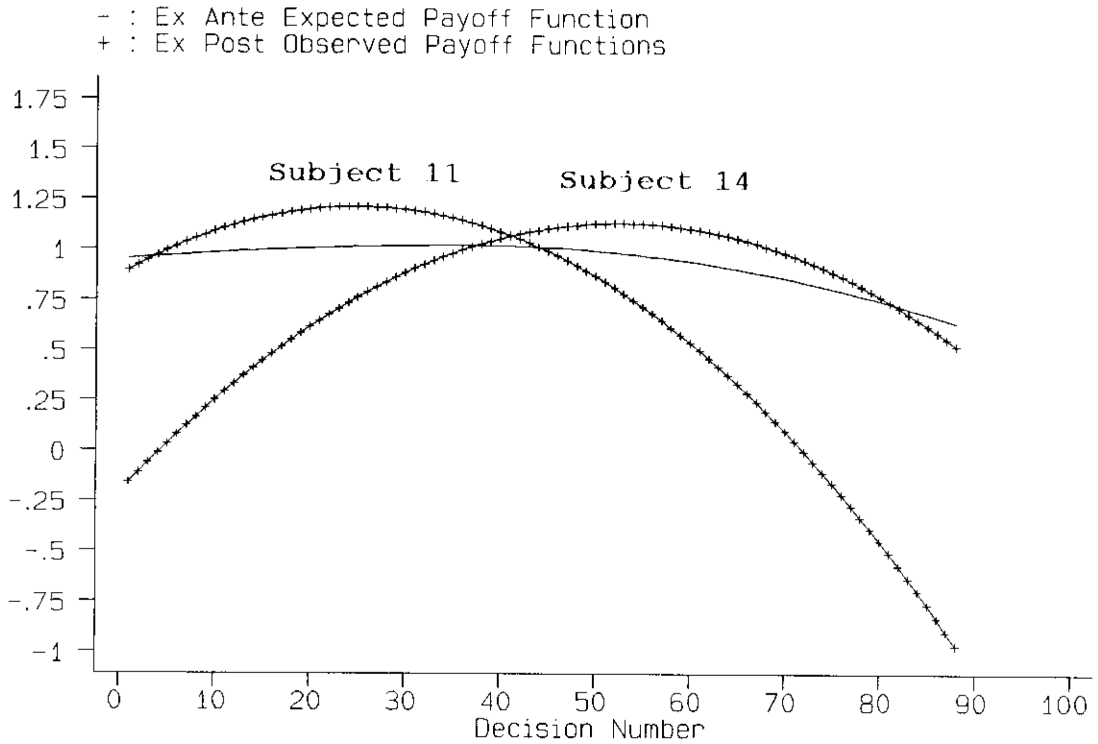

FIGURE 5 - Automaton Experiment

Ex Ante & Ex Post Payoff Functions

superimpose the ex post observed payoff function for two experimenters on the ex ante expected payoff function in the automaton experiment.

As we can see from figure 3, there is little in these diagrams to indicate to an experimenter-subject that the payoff function was flatter in the piece-rate than in the automaton experiment. In fact to the eye one might even claim that the piece-rate payoff function was flatter. To make these descriptive ideas more concrete, we estimated the best quadratic approximation through the scatter of points for each of these seven subjects in an effort to see if the "best" inference they could make about the shape of the payoff function - using the information on their realized payoffs - could lead them to the "right" picture. The estimated coefficients of the regressions used to draw the ex post observed payoff functions in figure 4 are presented in Appendix B together with the coefficients of the ex ante expected payoff functions. By looking at figure 4.a we can observe that just one subject in the piece-rate experiment would in principle be able to discover the shape of the payoff function, but only through a very costly experimentation strategy (i.e this subject lost .93$ when he chose an effort level of 100). As we can see from figure 4.b however, it seems unlikely that any of the four experimenters in the automaton experiment could successfully identify the shape of the payoff function of the experiment. Furthermore, by looking at figure 5 we can see that there is the possibility that experimenters get "trapped" into a wrong region. For example after experimentation it appears to one experimenter (subject 11) that the payoff function he faces is steep around 23, and to another experimenter (subject 14) that this function attains its maximum at 55[4]. This result would then lead us to think that even with

---

[4] Whereas the ex ante expected payoff function is peaked at 37.

experimentation subjects may fail to perceive the true shape of the experimental payoff function. In fact they may convince themselves that its shape is steep and peaked at a multitude of values, leading to a wide variance of choices in the last period. Note that the steepness of the ex post observed payoff functions in figure 5 is due to the fact that the sampling realizations of the randomness in the experiments may influence the shape of these functions. For example, the inference that player 11 can make about the shape of the payoff function he faces in the automaton experiment, is very likely to be affected by the fact that when he chose exactly 23 he always received the big monetary prize, whereas choices in the neighborhood of 23 were not as succesfull (i.e. he sometimes won the big prize and sometimes got only the small prize). Here then in the neighborhood around 23 the payoff function appears to be steep.

Section 5: Conclusion

Before one can apply the Harrison's criticism one must actually look at the data of the experiment on an individual-by-individual basis to discover whether in fact subjects acted like experimenters or theorists in their experiment. Whatever one concludes about this first step, one must then look at the ex post observed payoff function and see if it is sufficiently flat to warrant the criticism leveled by Harrison (1989). If subjects are not good theorists, then they will only infer the shape of the payoff function by trial and error experimentation during the experiment. If this experimentation does not yield a set of data points which indicate a flat payoff function, then despite the shape of the theoretical ex ante function, it can not be claimed that the subjects actually experienced such flatness. If subjects are good theorists and the problem is not too hard, then no matter what the shape of the payoff function they face, they should choose an optimal action. This is not to mean that Harrison's criticism is ill guided. Quite the contrary, we think it is a fundamental point that needs to be made. In addition, a good experimental design must worry about just this point and other things being equal, one should always strive to make the payoff function subjects face at equilibrium sufficiently steep to satisfy it. However, in designing an experiment we are faced with innumerable tradeoffs placed on us by the model we are working with, the statistical tests needed to test it, and financial constraints. This may constrain the design and produce payoff function which are not as steep as one would like. It is our point here that such flatness in the payoff function is not per se evidence that the experiment is invalid. In a world of sophisticated theorists, flatness does not matter at all. In a world of naive experimenters, it matters only if it is observationally evident to the subjects and this may not be possible in an experiment run even for as many times as 25, which is typically greater than the number of trials experimentalists use.

Appendix A

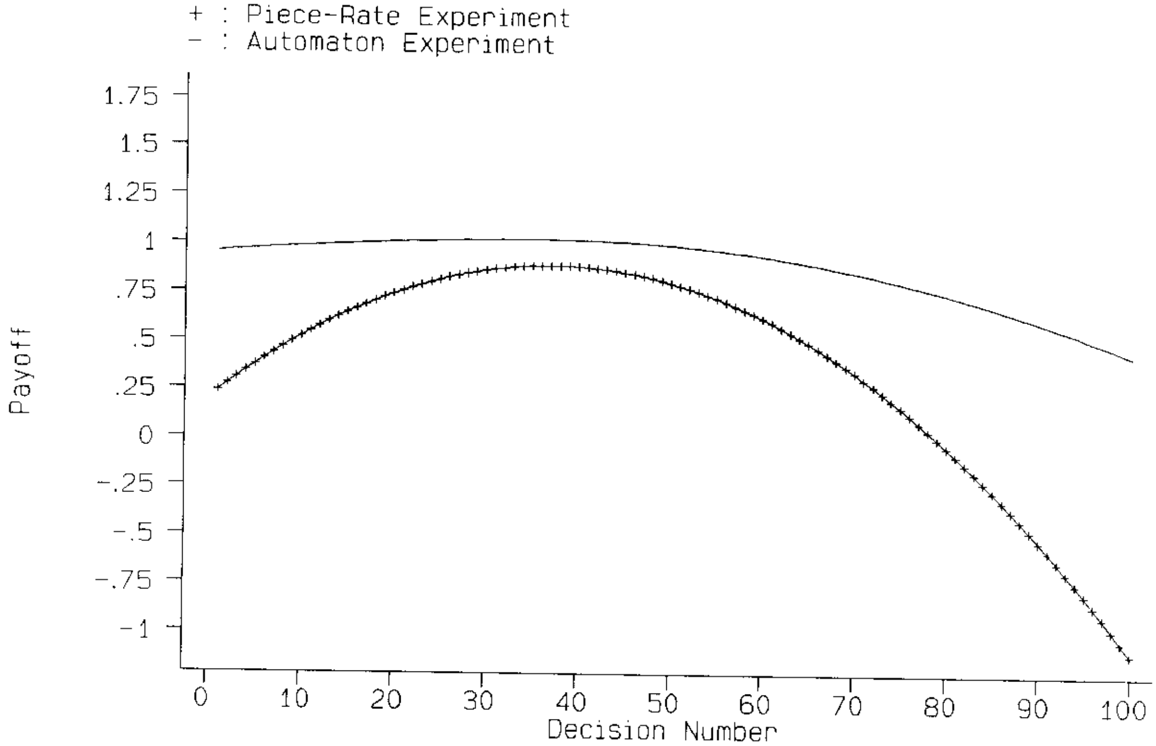

FIGURE 6

Ex Ante Expected Payoff Functions

Appendix B

TABLE 3

COEFFICIENTS OF THE EX ANTE EXPECTED PAYOFF FUNCTIONS

payoff function: $\pi = a + b_1 e + b_2 e^2$, e: effort level

		a	b_1	b_2
Piece-Rate Experiment		.2	.037	-.0005
Automaton Experiment*	(if $e \leq 37$)	.95	.004	-.00005
	(if $e \geq 37$)	.85	.01	-.00014

* The probability of winning the big monetary prize conditional on e is in fact different in the two regions

TABLE 4

COEFFICIENT ESTIMATES OF THE EX POST OBSERVED PAYOFF FUNCTIONS

payoff function: $\pi = a + b_1 e + b_2 e^2$, e: effort level

	a	b_1	b_2
Piece-rate experiment			
subject 2	.455129	.018557	-.00027
subject 10	-2.934331	.201309	-.00250
subject 12	.265038	.033577	-.00046
Automaton experiment			
subject 11	.869337	.027318	-.00055
subject 12	1.268799	.001564	-.00007
subject 13	1.525669	-.016198	.00008
subject 14	-.209029	.050888	-.00048

References

Bull, Clive, Schotter, Andrew and Weigelt, Keith, "Tournaments and Piece Rates: An Experimental Study," Journal of Political Economy, February 1987, 95, 1-33.

Cox, James C., Roberson, Bruce and Smith, Vernon L., "Theory and Behavior of Single Object Auctions," in Vernon L. Smith, ed., Research in Experimental Economics, Greenwich: JAI Press, 1982.

_______, Smith, Vernon L. and Walker, James M., "Test of a Heterogeneous Bidder's Theory of First Price Auctions," Economic Letters, 1983, 12, 207-12.

_______, _______, and _______, "Experimental Development of Sealed-Bid Auction Theory: Calibrating Controls for Risk Aversion," American Economic Review, May 1985, 75, 160-65.

_______, _______, and _______, "Theory and Individual Behavior of First-Price Auctions," Journal of Risk and Uncertainty, March 1988, 1, 61-99.

Drago, Robert and Heywood, John S., "Tournaments, Piece Rates, and the Shape of the Payoff Function," Journal of Political Economy, August 1989, 97, 992-8.

Fudenberg, Drew and Kreps, David M., "A Theory of Learning, Experimentation, and Equilibrium in Games," Mimeographed, MIT and Stanford Un., July 1988.

Harrison, Glenn W., "Theory and Misbehavior of First-Price Auctions," American Economic Review, September 1989, 79, 749-62.

Kagel, John H. and Roth, Alvin E., "Comment on Harrison versus Cox, Smith and Walker: 'Theory and Misbehavior in First-Price Auctions'," Mimeographed, Un. Pittsburgh, March 1990.

Smith, Vernon L., "Microeconomic Systems as an Experimental Science," American Economic Review, December 1982, 72, 923-55.