



Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
<http://www.econ.upenn.edu/pier>

PIER Working Paper 01-040

“Learning By Not Doing: An Experimental Investigation of Observational Learning”

by

Antonio Merlo and Andrew Schotter

http://papers.ssrn.com/sol3/papers.cfm?abstract_id=290084

Learning By Not Doing: An Experimental Investigation of Observational Learning*

ANTONIO MERLO
(University of Pennsylvania and CEPR)

ANDREW SCHOTTER
(New York University)

Revised, October 2001

ABSTRACT. In this paper we present experimental evidence suggesting that observational learning (i.e. learning-by-not-doing but by observing) may outperform learning-by-doing.

JEL Classification: C91, D83.

Keywords: Learning.

*We thank two anonymous referees and an Associate Editor for their useful suggestions. Seminar and conference participants at several institutions provided helpful comments. This paper was completed while Schotter was visiting the Russell Sage Foundation. Their generous hospitality and the support of the C.V. Starr Center are gratefully acknowledged. Meta Brown and Shachar Kariv provided excellent research assistance.

1. Introduction

Common sense might suggest that if two identical rational economic agents who are attempting to maximize the same objective function have the same data or information set at their disposal, they will make the same choice. Perhaps for this reason it is hardly ever asked whether the process through which this data was generated matters or is it sufficient to merely know the stock of information the agents have in their possession at the time the choice is made. This question is clearly important for many economic situations. For example, think of an environment where a task is repeatedly performed and the expertise of the decision maker increases as a function of the cumulative stock of output produced as in a learning-by-doing environment. In such an environment, does the stock of observations generated along the way summarize all the information needed to make efficient economic decisions or is the actual experience that led to these observations important? In this paper we present experimental evidence that addresses this question.

One can envision learning in distinctly different ways. Observational learning, would have a decision maker (the observer) simply looking over the shoulder of another decision maker (the doer) as he performs his learning-by-doing task (see Jovanovic and Nyarko (1995) for an application of an apprentice-craftsman model similar, in spirit, to the one we explore here). Here the observer simply keeps his eyes open but his hands in his pocket. Still after enough repetition both the observer and the doer would have the same observations and if they are rational and out to maximize the same objective function, we might actually expect the doer to perform no worse than the observer since he or she has had additional hands-on experience.

A large number of animal studies support the hypothesis that observational learning can be efficient. For example, Terkel (1996) shows that young rats seem to learn how to skin pine cones quite easily by observing their mothers when they are young. John, et al, (1969) demonstrates that cats can be trained to perform tasks by being placed in a cage divided down the middle with a glass wall through which they can observe an animal already trained in a particular task perform. After a certain amount of observation, the cat seems perfectly able to perform the task despite the

fact that it has never done it.

In this paper we investigate the hypothesis that the process through which information is gathered can affect the choices of experimental subjects drawn at random from identical populations. What we find is that subjects can sometimes learn better by not doing but by watching. A possible explanation is that when subjects perform an experiment or participate in a market repeatedly and earn a small payoff each period, as is true for most markets, the feedback of the market tends to focus attention of the decision makers on the myopic stimulus-response aspect of the problem and seems to prevent them from learning the trade-offs existing in the experiment (see Merlo and Schotter (1999), henceforth M-S, for a fuller exposition of this point). However, when such behavior is merely observed, attention is again re-focused and performance is enhanced. To some extent, this explains the value of coaches, theatrical directors, Sunday morning quarterbacks, and kibitzers in general.

Another finding which we consider of interest is the fact that not only is observation potentially performance improving, but how much better an observer does depends crucially on who they observe. Put differently, those observers watching doers who did relatively well made far better decisions than those who watched relatively poor doers. Good role models are crucial to good performance.

In the remainder of this paper we proceed as follows: In Section 2 we discuss the decision task that our subjects performed and our baseline experimental design. In Sections 3 and 4 we present the results of our baseline experiment. In Section 5 we report the results of two additional experiments run to answer potential criticisms of our baseline experiment. Finally, in Section 6 we offer some conclusions.

2. The Decision Task and the Experimental Design

As was true in M-S, all of the experiments performed to investigate our hypotheses were of the tournament variety and similar to those of Bull Schotter and Weigelt (1987) and Schotter and Weigelt (1992). In those experiments, randomly paired subjects must, in each round, 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 closed interval $[-a, +a]$. These numbers, each player’s decision number and random number, are then added together and a “total number” defined for each subject. 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 m , $M > m$. The cost of the decision number chosen, given by a convex function $c(e) = e^2/k$, is then subtracted from these fixed payments to determine a subject’s final payoff. Hence, in these experiments there is a trade-off in the choice of decision numbers: higher numbers generate higher probabilities of winning the big prize but also imply higher decision costs.¹ By letting $k = 500$, $a = 40$, $M = 29$, and $m = 17.2$, the two-person tournament has a unique Nash equilibrium at 37. By replacing one player with a computerized automaton programmed to always choose 37, and informing the live subject of that fact, we transform the problem for the remaining live player into a one-person maximization problem with a quadratic payoff function. The objective function in this problem is the conditional payoff function obtained from our two-person tournament game after restricting the automaton’s choice to equal 37.

We consider the experimental task 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. Second, in spite of the complexity of the decision problem, the experiment is simple to describe and to understand. This feature is appealing since it should reduce the noise in the data.

As in the M-S experiments, subjects in our baseline experiment repeated the decision task described above 75 times and were paid a small amount after each trial. When they were finished, they were surprised to find out that they were going to do the experiment one more time for “big stakes”, i.e. amounts that were 75 times

¹In the instructions, a sample of which is contained in the Appendix, we take great care not to use value loaded terms such as “winning” or “losing”.

the period-to-period payoffs they earned in the previous 75 rounds. The choice they make in this “surprise-quiz” round should be a snapshot (sufficient statistic) of exactly what they learned during the previous 75 rounds. Presumably, the choice they make here should be their best guess as to what is the optimal choice. M-S focus their attention on comparisons of the surprise-quiz rounds in experiments with different payoff environments.

The difference between our experiments here and those of M-S is that in the experiments we report on here as each subject performed the experiment another subject watched quietly over his or her shoulder. Strict silence was rigorously enforced (although no instance of talking occurred) and subjects were forbidden to communicate in any manner either by groaning, sighing, laughing, raising an eye-brow or in any other manner. These subjects who watched without doing the experiment (the observers) were not told what they would do after the subjects they were watching (the doers) were finished, but were told that what they were observing was relevant to the experiment they would ultimately do. They left the room before the surprise-quiz was announced to the doers. When they returned, the doers were gone, having performed a surprise-quiz round, and they were told that all they needed to do in the experiment was to play the same game they just watched being played once and only once for big stakes, i.e. amounts that were 75 times larger than the individual round payoffs of the experiment they watched.

Note that in the experiments we performed, the information of each pair of subjects (a doer and an observer) was identical, since each observer was watching over the shoulder of a doer. We simply manipulated the manner in which the information was obtained with the doer actually generating it and the observer passively watching it being generated. What we compare in our analysis are the surprise-quiz choices of these two groups of subjects. As we said before, if these subjects learned differently or learned different things, this differential learning should be detectable in their surprise-quiz choices.

2.1. Procedures

All of the experiments reported here were run in the experimental economics laboratory of the C.V. Starr Center for Applied Economics at New York University. Students were recruited and brought into a room in groups ranging in size from 10 to 20. There were a total of 62 subjects who participated in the Baseline Experiment, and an additional 124 subjects who participated in two additional experiments (Experiments 2 and 3) run to answer potential criticisms of our baseline experiment. The procedures we followed in our baseline experiment are described below and the results of the experiment are discussed in Section 3. The procedures followed in Experiments 2 and 3 and the results of these experiments are discussed in Section 5.

When subjects arrived at the lab to participate in the Baseline Experiment, they were randomly divided into two types called A and B with the A types being the doers and the B types being the observers. The instructions were handed out to all subjects. Observers were told that after the subjects they were watching were done they would partake in an experiment themselves, but they were not told the nature of the experiment. All they were told was that the experiment would, in some way, be related to the one they observed. Subjects who were actually doing the experiment completed 75 rounds and before they chose in the surprise-quiz, the observers were asked to leave the room so that they could not see what decision was ultimately made in the big-stakes round. Observers were subsequently brought back in and asked to make one and only once choice for big stakes. All payoffs in the experiments were denominated in an experimental currency called francs which was converted into dollars at the rate of 0.01:1. In the surprise-quiz rounds the conversion rate was 0.75:1.

Subjects who performed the Baseline Experiment as doers earned approximately \$30 for the one and one half hours they spent in the lab, since they in essence did the experiment twice. All other subjects earned half that amount.

2.2. Research Questions

While no theory exists upon which to base predictions for these experiments, the previous experiments of M-S do help in establishing a basis upon which a series of questions can be formulated. In M-S the authors discovered that subject choices in the surprise-quiz portion of the experiment were not particularly good in the sense that their choices were not close to the optimal choice of 37. They conjectured that this was true because in the environment in which they were searching, where they received small payoffs after each round, the object of learning shifted from discovering the peak of the payoff function they faced to discovering the appropriate manner in which to respond to the period-by-period feedback they received from the experiment. Hence subjects processed the data they generated in a particular fashion and when they were surprised at the end of the experiment by a surprise quiz, they realized that they have failed to provide themselves with a good cognitive map of the payoff function they faced. If this conjecture were true, and if the presence of an observer did not influence subject behavior, then we would expect subjects functioning as doers in our experiment to perform in a similar fashion as the subjects in M-S and fail to learn the optimal decision.

The focus of this paper is to investigate the performance of the observers relative to the doers. We do this by using the data generated by our subjects in their surprise-quiz rounds and asking the following four questions.

Question 1: Are the median surprise-quiz choices of both the doer and observer subjects equal to 37?

Question 2: Subject-pair by subject-pair, how do the surprise-quiz payoffs of observers compare to those of doers?

Question 3: Is the distribution of surprise-quiz choices of observers equal to the distribution of surprise-quiz choices of doers?

Question 4: Within subject-pairs, are the surprise-quiz choices of doers and observers highly correlated?

3. Results of the Baseline Experiment

The discussion of the results of our Baseline Experiment is presented in terms of the questions we proposed to answer. We begin by presenting some simple descriptive statistics of the choices doer and observer subjects made in their surprise-quiz rounds.

Table 1 presents the surprise-quiz choices of our doers and observers as well as their expected payoffs given the choices that they made and the distribution of random shocks assumed in the experiment. Several features of the data are salient.

- 1) While the median choice of the observers was 37 (the optimal choice), the median choice of the doers was 50. (The means were 40.65 and 51.06, respectively).
- 2) The mean and median choices of doers (51 and 50, respectively) were the same as the mean and median choices of subjects in M-S. This result both replicates our previous result and indicates that the presence of an observer did not alter the behavior of doers.
- 3) While there were 12 observers whose surprise-quiz choice deviated from the optimal choice of 37 by five or less, there were only five such deviations for the doers. Moreover, while six subjects in the observer sample chose 37 exactly, only one subjects in the doer sample did.
- 4) Although the observers had a higher mean and median payoff than did the doers, because the payoff function was relatively flat around the optimum, these difference do not seem very dramatic. Still, in 18 out of the 31 pairs the observer made a choice which was payoff increasing when compared to the doer's choice, and only in 10 out of the 31 pairs the observer made a choice which was payoff decreasing (in the remaining three cases the doer and the observer made the same choice).
- 5) For the doers, in the surprise-quiz round, 9 out of the 31 subjects made dominated choices (that is, choices of 65 and over) while for the observers there were only 5 such choices.

- 6) The correlation of doer and observer choices, conditional on observing the same information, was only 0.39 indicating that despite their shared history, the observer and doer subjects processed the information they experienced quite differently.

To summarize, there was a definite difference in the surprise-quiz choices of doers and observers.

The results in Table 1 are summarized by the histograms presented in Figures 1 and 2 where we plot the choices made in the surprise-quiz rounds by doers and observers (Figure 1) and the absolute deviations from the optimal choice of 37 for each group (Figure 2).

We now turn our attention to answering the four questions posed above. For each question, we present the results (and p-values) of simple non-parametric tests of specific hypotheses that address the question. Throughout the paper, we adopt the convention that a null hypothesis can (cannot) be rejected at conventional levels of statistical significance if the p-value of its test is smaller than (greater than or equal to) 0.10.

Question 1: Are the median surprise-quiz choices of both the doer and observer subjects equal to 37?

To answer this question we perform a simple Binomial test individually on the sample of doer and observer surprise-quiz choices testing that the probability that observations are above 37 equals the probability that observations are below 37. What we find is that while we can obviously not reject the hypothesis that the median of the observer choices was equal to 37 (remember the actual median was in fact 37), we can reject the hypothesis that the median of the doer choices (which was 50) was equal to 37 using a normal approximation (p-value 0.007).

Question 2: Subject-pair by subject-pair, how do the surprise-quiz pay-offs of observers compare to those of doers?

To answer this question we use a sign test of the following type. Take any doer subject in the experiment and his or her associated observer. Then take the surprise-

quiz choice of the doer and compare it to the surprise-quiz choice of the observer. Score the choice of the observer with a plus if it leads to a payoff improvement over that of the doer. If it leads to a decrease in the subject's payoff score it as a minus and if the subjects choose the same decision number score it as a zero and drop that observation. We then test the hypothesis that the median difference between these two groups is zero, which would imply that the probability of getting a payoff increase for observers would be exactly equal to the probability of a payoff decrease. In other words, we would expect this "increase probability," π , to equal $1/2$. Under the null that $\pi = 1/2$, we calculate the probability of observing at least 18 pluses out of 28. What we find is that, using a normal approximation, we can reject the hypothesis of equality of payoffs in favor of the one-tailed alternative that the choices made by our observers were payoff increasing (p-value 0.093).

Question 3: Is the distribution of surprise-quiz choices of observers equal to the distribution of surprise-quiz choices of doers?

The non-parametric procedure we use to answer this question is based on Pearson's chi-square test. Let $h = 1, 2, \dots, 10$ index the bins of the support of the histogram of the distribution of surprise-quiz choices (as depicted in Figure 1), where $h = 1$ denotes the $[0-10)$ bin, $h = 2$ denotes the $[10-20)$ bin, and so on. Let n_h^d and n_h^o denote the empirical frequency of the observations within bin h for the doers and the observers, respectively. Then, the test-statistic

$$Q = \sum_{h=1}^{10} \frac{(n_h^o - n_h^d)^2}{n_h^d}$$

has a chi-square distribution with 9 degrees of freedom. Note that Q is a measure of the departure of the observed choices of the observers from the observed choices of the doers. What we find is that we can reject the null hypothesis that the distribution of surprise-quiz choices of the observers is equal to the distribution of surprise-quiz choices of the doers ($Q = 22$, p-value 0.009).

Question 4: Within subject-pairs, are the surprise-quiz choices of doers and observers highly correlated?

The correlation between doer and observer choices is 0.39. We test the hypothesis that this coefficient is equal to one, which would indicate perfect correlation between the choices of doers and observers. Using a normal approximation, the test-statistic has a Student's t distribution with 29 degrees of freedom. What we find is that we can reject the hypothesis that the surprise-quiz choices of doers and observers are perfectly correlated (p-value 0.002). This implies that despite the fact that doers and observers observed identical data before their surprise-quiz rounds, their behavior differed substantially.

4. Do Good Craftspeople Make Good Apprentices?

What we have demonstrated so far is that in general and on average, observers do better than doers in the decision task we have confronted them with. However, we find an equally interesting result when we disaggregate our data slightly in the following manner. Look at the doers in our experiment and define those doers who received a payoff that was better than the median during their first 75 rounds in the experiment as “better doers”. Define those who did worse than the median as “worse doers”. Now consider those observers who watched the better doers and those who watched the worse doers and look at their decisions in the surprise-quiz round. When we do this we find that those observers who watched the better doers did significantly better than those who watched the worse doers. More precisely, if one calculates the median absolute difference between the surprise-quiz choices made by those subjects who watched the better doers and the optimal decision of 37, one finds that this difference is 3 (the mean difference is 9.1). In other words, half of the 16 subjects who watched better doers made surprise-quiz choices that were less than three units away from 37. For the observers who watched worse doers, the median absolute difference is 27 (mean 27.7). This is, we believe, a dramatic difference. In addition, of the five observers who, in the surprise-quiz round, chose exactly 37, all five were in the subset who watched better doers.

Another interesting finding is that the improvement of decisions of subjects observing better doers over the people they watched was relatively better than their cohorts

observing worse doers. For example, while the better (worse) doers had a median absolute deviation from 37 of 13.5 (32) the subjects that observed them made choices that had a median deviation from 37 of only 3 (27). Hence, the observers of the better doers had a greater percentage improvement in their performance over those whom they observed (from 13.5 to 3, an 80.7% improvement in the median) than did the observers of worse doers whose median deviation from 37 only improved from 32 to 27 (a 15.67% improvement).

The point seems clear. While watching can improve performance, watching someone good improves it even more. It appears that people have a relatively easier time learning from example than from object lesson. Put differently, if one is going to watch a craftsman work and learn at his shoulder, they would be better off watching a good one and learning from his skill than watching a bad one and learning from his mistakes.

5. Additional Experiments

While the results described above may be considered striking, the design used to generate them is not free of criticism. For example, one may say that our results are forced since while we told observers that what they were watching would be payoff relevant for them in a subsequent experiment, the doers had no such warning. Hence, it might be argued, it is no surprise that observers watched with a more careful eye than did doers.

It should be pointed out that we do not feel this to be a criticism, but rather a summary of our point. It is our contention that most economic environments are not conducive to learning since they involve settings where people are paid each period and engage in their decision-making task with no directed aim of learning the workings of the market or institutions they are functioning in. So, we would not be surprised that if we told them that they should pay extra attention to what they were doing (because they would need it in a future experiment), it might increase their performance. Our point is that in naturally occurring markets no such warning is offered.

To deal with this criticism we ran another experiment, Experiment 2, using 35 pairs of subjects. This experiment was identical to our baseline experiment described above except that here both the doers and the observers were told that after the doers did their experiment both groups would separately perform another experiment which was related to the one the doers just did. They were not told that the upcoming experiment was a simple surprise quiz, however. In fact, the exact wording was as follows: “When the 75 rounds are over, another experiment will be described which both Groups A and B will perform. The experiment is closely related to the one the A group just did, so it is important that you pay attention to the first experiment. The Group A people will do this experiment first and will stay in the lab, perform the experiment, and then leave after being paid. The Group B will initially leave the lab while the Group A finishes, and then return, perform the second experiment, be paid for their performance and then leave.”

The results of Experiment 2 are presented in Table 2 and in Figures 3 and 4. Table 2 presents the surprise-quiz choices of our doers and observers as well as their expected payoffs. Figure 3 depicts the distribution of choices for doers and observers in this experiment. Figure 4 contains the histogram of the absolute deviations of these choices from 37 for each group.

Our results here are analogous to the ones reported above for our baseline experiment. For example, while the median choices of doers and observers do not differ greatly (they are equal to 44 and 42, respectively), using the same Binomial test as the one used above we can reject the hypothesis that the median of the doer choices is equal to 37 (p-value 0.091), but we cannot reject the same hypothesis for the observers (p-value 0.170). Also, the choice distributions of the two groups of subjects are quite different. Using a Pearson’s chi-square test like the one used above, we find that we can reject the null hypothesis that the distribution of surprise-quiz choices of the observers is equal to the distribution of surprise-quiz choices of the doers ($Q = 17$, p-value 0.049). More descriptively, notice that the choices of observers were more clustered around the optimum of 37 (while observers made 23 of 35 surprise-quiz choices in the interval [25,50], doers only made 14 such choices).

Looking at how the surprise-quiz payoffs of observers compare to those of doers, we see that in 18 of the 28 subject-pairs where the surprise-quiz choices of doer and observer differed, the choices of the observer were payoff increasing as compared to their doer counterpart. Using a sign test identical to the one used above we can reject the hypothesis of equality of payoffs in favor of the one-tailed alternative that the choices made by our observers were payoff increasing (p-value 0.093).

As it was true in our baseline experiment, more doers (6 subjects) made dominated choices in their surprise-quiz round (that is, choices of 65 or more) than did observers (2 subjects). Finally, the correlation between the surprise-quiz choices of doers and observers is 0.21. Using a t-test similar to the one used above, we can reject the hypothesis that the surprise-quiz choices of doers and observers are perfectly correlated (p-value 0.000). This finding indicates that observers came to different conclusions than the doers they were watching perform as to which choice was the best in the surprise-quiz round.

In summary, while some of the differences were weaker in Experiment 2 than they were in the baseline experiment, there are substantial qualitative and quantitative differences in the choices of doers and observers even when both groups are told that the experiment that the doers were performing would be payoff relevant in a future experiment.

Another criticism of our baseline experiment is that since the doers were paid for their first 75 rounds while the observers were not, the differences in the surprise-quiz round choices might simply reflect an income effect that affected the doers and not the observers. To rectify this we ran a third experiment, Experiment 3, using 27 pairs of subjects. In this experiment, we paid each observer the same amount of money earned by the doer he or she was assigned to observe in the 75-round experiment, but told only the observers that there would be a follow-up experiment. In other words, in this experiment we eliminated the income effect but preserved the informational structure of the baseline experiment.

The results of Experiment 3 are presented in Table 3 and Figures 5 and 6. Table 3 contains the surprise-quiz choices of our doers and observers as well as their expected

payoffs. Figure 5 depicts the distribution of choices for doers and observers in this experiment. Figure 6 contains the histogram of the absolute deviations of these choices from 37 for each group.

Again, the results we obtain in this experiment are very close to the ones reported above for our baseline experiment. In particular, we conclude that the surprise-quiz choices of doers and observers were different, and that, by and large, observers made choices that were payoff-increasing relative to the choices of the doers they observed perform. For example, the median choices of doers and observers are equal to 49 and 45, respectively. While we cannot reject the hypotheses that either median is equal to the optimal choice of 37 (p-values 0.182 and 0.230, respectively), we see that the choice distributions of the two groups of subjects are different. Using a Pearson's chi-square test identical to the one used above, we find that we can reject the null hypothesis that the distribution of surprise-quiz choices of the observers is equal to the distribution of surprise-quiz choices of the doers ($Q = 18$, p-value 0.035). More descriptively, notice that the choices of observers were more clustered around the optimum of 37 (while observers made 19 of 27 surprise-quiz choices in the interval [25,50], doers only made 13 such choices).

Looking at how the surprise-quiz payoffs of observers compare to those of doers, we see that in 17 of the 25 subject-pairs where the surprise-quiz choices of doer and observer differed, the choices of the observer were payoff increasing as compared to their doer counterpart. Using a sign test like the one used above we can reject the hypothesis of equality of payoffs in favor of the one-tailed alternative that the choices made by our observers were payoff increasing (p-value 0.055).

As it was true in our baseline experiment, more doers (5 subjects) made dominated choices in their surprise-quiz round (that is, choices of 65 or more) than did observers (1 subject). Finally, the correlation between the surprise-quiz choices of doers and observers is 0.56. Using a t-test similar to the one used above, we can reject the hypothesis that the surprise-quiz choices of doers and observers are perfectly correlated (p-value 0.004).

In summation, we conclude that the presence of an income effect had little bearing

on the results of our baseline experiment. In fact, significant differences between the choices of doers and observers persist even after eliminating the income effect (as we have done in Experiment 3).

6. Conclusions

In this paper we have attempted to make one simple point. The process through which information is accumulated may have a significant impact on the quality of economic decisions made by decision makers. Moreover, observational learning can outperform learning-by-doing. We believe these results are puzzling and provocative and pose some new challenges for economic theory.

For example, it is striking that 30% of the doer subjects in the Baseline Experiment made dominated choices in the surprise-quiz round, while the subjects who observed them did not fall victim to this mistake. This fact raises an interesting question about the type of learning that doers and observers might be engaged in. More precisely, because doers are paid a small amount each period for their choice, they may tend to learn more like reinforcement learners and increase the probability of making any particular choice as the cumulative payoff generated by that choice increases. However, reinforcement learning only reinforces those choices that are actually made. If an action is not chosen, it gets no reinforcement. Hence, if a doer continually chose high actions some of which were dominated, she would continue to have only those actions reinforced and never even learn they were dominated. Observers, on the other hand, may behave more like the learners described by Camerer and Ho (1999), who also reinforce choices that are not made by calculating the hypothetical payoff associated with these choices. Observers may be more likely to behave this way since they are in a more abstract situation, one which lends itself to this type of theorizing. Hence, while one model of learning may explain the surprise-quiz choice of the doer another model may explain the behavior of the observer.

Finally, we are not claiming that observational learning would outperform learning-by-doing in all learning tasks. There are many situations where physically engaging in a task may be important. For example, we would not expect someone to be able

to fly a plane after sitting in the cockpit and watching. For abstract mental tasks it is more likely that observing can be informative. However, as indicated in M-S, even for mental tasks payoff environments may matter.

Appendix

A.1. Instructions for the Baseline Experiment

This is an experiment on decision making. Funds have been provided for this experiment by a number of research institutes and foundations and if you pay attention and make good decisions, you may earn a good amount of money which we will arrange to pay you after the experiment is over.

As you walk into the room you will randomly divided into two equal sized groups, Group A and Group B. If the upper right hand corner of this instruction sheet has an A on it your will be in the A Group, if it has a B, you will be in the B Group. The experiment will then proceed as follows: Both groups will read the attached instructions which are identical for all people. After you have finished reading the instructions we will read them out loud and answer any question you might have about them. When this instructional period is over we will proceed to the computer lab where the people in the A Group will sit in front of a terminal and perform the experiment described in the instructions for 75 rounds, i.e., they will repeat the same experiment 75 time. Each person in the B Group will be matched up with one person in the an A Group and while that person performs the experiment he/she will silently watch by looking over his/her shoulder.

As the B Group people watch there must be no communication between the A Group person and his pair member. This means, no talking, shared looks, groans, giggles, or reactions of any type. Just silence.

When the 75 rounds are over, the B Group people will leave the computer lab and another experiment will be described to them which they will perform. This experiment is closely related to the one the A Group people just did, so it is important that you pay attention to what your cohorts do. The A Group people will stay in the lab where they will receive further instructions, be paid, and leave. The B Group will then return to the lab when the A Group is gone, perform their experiment, be paid and leave.

A.1.1. *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.

A.1.2. Experimental Procedure

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

PLAYER# ----

ROUND DECISION# RANDOM# TOTAL# COST EARNINGS

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

IS ___ YOUR DECISION NUMBER? [Y/N]

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

A.1.3. Calculation of Payoffs

Your payment 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 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.

A.1.4. 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. When round 75 is completed, the computer will ask you to press any key on its keyboard. After you do this, the computer will add your Francs earnings from each of the 75 rounds and convert this sum in Dollars at the rate of \$0.01 per Franc. We will then pay you this amount.

A.1.5. Example of Payoff Calculations

Suppose that the following occurs during one round: pair member A2 chooses a decision number of 60 and generates a random number of 10, while computerized pair member A, selects a decision number of 37 and gets a random number of 5. Pair member A2 would then receive the high fixed payment of 29 Fr. From this fixed payment, A2 would subtract 7.2 Fr. (the cost of decision number 60). A2'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.

Decision Cost Table

A	B	A	B	A	B	A	B
Decision	Cost of	Decision	Cost of	Decision	Cost of	Decision	Cost of
Number	(Francs)	Number	(Francs)	Number	(Francs)	Number	(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		

References

- [1] Bull, C., Schotter, A., and Weigelt, K. (1987), "Tournament and Piece Rates: An Experimental Study," *Journal of Political Economy*, vol. 95, pp. 1-33.
- [2] Camerer, C. and Ho, T. H. (1999), "Experience-Weighted Attraction Learning in Normal Form Games" *Econometrica*, vol. 67, pp. 827-874.
- [3] John, E.R., Chesler, P., Bartlett, F., and Victor, I. (1969), "Observational Learning in Cats," *Science*, vol. 166, pp. 901-903.
- [4] Jovanovic, B. and Nyarko, Y. (1995), "The Transfer of Human Capital," *Journal of Economic Dynamics and Control*, vol. 19, pp. 1033-1064.
- [5] Merlo, A. and Schotter, A. (1999), "A Surprise-Quiz View of Learning in Economic Experiments," *Games and Economic Behavior*, vol. 28, pp. 25-54.
- [6] Schotter, A. and Weigelt, K. (1992), "Asymmetric Tournaments, Equal Opportunity Laws and Affirmative Action: Some Experimental Results," *Quarterly Journal of Economics*, vol. 106, pp. 513-539.
- [7] Terkel, J. (1996), "Cultural Transmission of Feeding Behavior in Black Rats (*Rattus rattus*)," pp. 17-47 in Cecilia Heynes and Bennett Galef, Jr., eds., *Social Learning in Animals and the Roots of Culture*, Academic Press: New York.

Table 1: Surprise-Quiz Results (Baseline Experiment)

Subject-Pair	Surprise-Quiz Choices		Expected Payoffs	
	Doer	Observer	Doer	Observer
1	50	40	20.14	20.34
2	45	40	20.27	20.34
3	78	76	18.41	18.59
4	41	57	20.33	19.87
5	38	37	20.35	20.36
6	69	97	19.15	16.25
7	69	0	19.15	18.94
8	57	34	19.87	20.35
9	36	1	20.36	19.02
10	100	37	15.84	20.36
11	29	35	20.31	20.36
12	100	20	15.84	20.08
13	20	37	20.08	20.36
14	29	40	20.31	20.34
15	60	37	19.72	20.36
16	30	30	20.32	20.32
17	25	30	20.23	20.32
18	37	54	20.36	20.00
19	88	88	17.37	17.37
20	55	68	19.96	19.23
21	52	20	20.07	20.08
22	98	77	16.12	18.50
23	13	13	19.78	19.78
24	79	50	18.32	20.14
25	40	37	20.34	20.36
26	70	37	19.08	20.36
27	22	13	20.15	19.78
28	51	57	20.11	19.87
29	44	0	20.29	18.94
30	57	60	19.87	19.72
31	1	38	19.02	20.35
Mean	51.06	40.65	19.40	19.71
Median	50.00	37.00	20.07	20.32

Table 2: Surprise-Quiz Results (Experiment 2)

Subject-Pair	Surprise-Quiz Choices		Expected Payoffs	
	Doer	Observer	Doer	Observer
1	47	45	20.22	20.27
2	80	38	18.22	20.35
3	65	66	19.43	19.36
4	21	36	20.11	20.36
5	43	37	20.30	20.36
6	50	50	20.14	20.14
7	39	39	20.35	20.35
8	99	52	15.98	20.07
9	37	37	20.36	20.36
10	59	37	19.77	20.36
11	25	25	20.23	20.23
12	70	48	19.08	20.20
13	37	41	20.36	20.33
14	60	73	19.72	18.85
15	44	45	20.29	20.27
16	38	38	20.35	20.35
17	68	50	19.23	20.14
18	55	36	19.96	20.36
19	23	37	20.18	20.36
20	63	23	19.55	20.18
21	40	44	20.34	20.29
22	35	60	20.36	19.72
23	35	22	20.36	20.15
24	60	26	19.72	20.25
25	13	55	19.78	19.96
26	39	78	20.35	18.41
27	24	47	20.20	20.22
28	55	59	19.96	19.77
29	60	54	19.72	20.00
30	55	55	19.96	19.96
31	60	60	19.72	19.72
32	37	40	20.36	20.34
33	1	40	19.02	20.34
34	1	37	19.02	20.36
35	67	42	19.30	20.32
Mean	45.86	44.91	19.77	20.09
Median	44.00	42.00	19.96	20.25

Table 3: Surprise-Quiz Results (Experiment 3)

Subject-Pair	Surprise-Quiz Choices		Expected Payoffs	
	Doer	Observer	Doer	Observer
1	20	34	20.08	20.35
2	49	49	20.17	20.17
3	74	45	18.76	20.27
4	50	38	20.14	20.35
5	50	55	20.14	19.96
6	23	31	20.18	20.33
7	48	40	20.20	20.34
8	42	48	20.32	20.20
9	32	37	20.34	20.36
10	29	38	20.31	20.35
11	60	56	19.72	19.91
12	64	36	19.49	20.36
13	50	36	20.14	20.36
14	65	63	19.43	19.55
15	40	38	20.34	20.35
16	15	1	19.88	19.02
17	28	50	20.29	20.14
18	59	63	19.77	19.55
19	68	48	19.23	20.20
20	38	45	20.35	20.27
21	67	57	19.30	19.87
22	38	25	20.35	20.23
23	45	45	20.27	20.27
24	65	67	19.43	19.30
25	60	42	19.72	20.32
26	55	51	19.96	20.11
27	0	45	18.94	20.27
Mean	45.70	43.81	19.90	20.10
Median	49.00	45.00	20.14	20.27

Figure 1: Histograms of surprise-quiz choices (Baseline Experiment)





Figure 2: Histograms of absolute differences from 37 (Baseline Experiment)





Figure 3: Histograms of surprise-quiz choices (Experiment 2)





Figure 4: Histograms of absolute differences from 37 (Experiment 2)





Figure 5: Histograms of surprise-quiz choices (Experiment 3)





Figure 6: Histograms of absolute differences from 37 (Experiment 3)

