ELSEVIER
European Economic Review 45 (2001) 615–628
EUROPEAN ECONOMIC REVIEW
www.elsevier.com/locate/econbase

Joseph Schupeter Lecture

A theorist's view of experiments☆

Ariel Rubinstein[a,b,*]

[a]Department of Economics, Tel Aviv University, 69978 Tel Aviv, Israel
[b]Department of Economics, Princeton University, Princeton, NJ, USA

Abstract

The paper springs from a position that economic theory is an abstract investigation of the concepts and considerations involved in real life economic decision making rather than a tool for predicting or describing real behavior. It is argued that when experimental economics is motivated by theory, it should not look to verify the predictions of theory but instead should focus on verifying that the considerations contained in the economic model are sound and in common use. It is argued that when theory is motivated by experiments, the theorist should not be hasty in adopting new functional forms but should try to identify the basic psychological themes which are revealed exposed by the experiment. Finally, some critical comments on the methodology of experimental economics are presented.

JEL classification: B0; C9

Keywords: Experimental economics; Hyperbolic discounting; Economic theory; Economic modeling

1. A view of economic theory

I am a pure theorist whom became interested in experimental economics only in the past few years. In this lecture I want to share with you some of my

☆This paper was delivered as the Schumpeter lecture at the European Economic Association meeting, Bolzano, September 2000 and at the European Economic Association meeting, New York, June 2000. My thanks to Jim Mirrlees and Andy Schotter for their invitations and encouragement.
* Corresponding author. Department of Economics, Tel Aviv University, 69978 Tel Aviv, Israel. Tel.: + 972-3-640-9601.
E-mail address: rariel@post.tau.ac.il (A. Rubinstein).

PII: S0014-2921(01)00104-0

thoughts about experimental economics and economic theory following my short detour into experimental economics. Some of the comments are over controversial, but is that not the role of such lectures?

If you had asked me 10 years ago about experimental economics I probably would have said that I do not see the point of conducting experiments in economics. This is not because I thought that economics is invariant to facts in the real world or because I found experimental economic to be boring. In fact, I found many of the experiments in decision theory and economics to be fascinating. However, I believed then that experiments do not need to actually be carried out. I believed that our intuition provides the test. If a phenomenon is robust, we intuitively recognize it as such. It strikes a chord upon us. If we are honest with ourselves we can feel that it is true. Look at philosophy. Almost the entire philosophical literature (so I am told) is based on introspection; very rarely have philosophers used experiments to verify their intuition. If this is sufficient for discussions of the great problems of our existence such as "what is good and what is bad" or "what is just and what is wrong", should it not be sufficient for the discussion of materialistic matters such as the choice between two simple lotteries with monetary rewards?

Ten years have passed. During that period, I began doing experiments myself. I rediscovered the obvious: To criticize something, you need to know it intimately; the best way to know something intimately is to do it yourself. Once you have done that, you do not want to criticize it anymore …

Actually, I was influenced by the numerous conversations I had with the late Amos Tversky, who was an exceptional teacher and a good friend. I enjoyed teasing Amos that I do not see why so much NSF money should be wasted on experiments since the results of his experiments were obvious and did not require confirmation. However, I came to understand (once we had worked together on several game theoretic experiments) that even his intuition (not to mention my own) was often wrong. I learned to appreciate how difficult it is to produce an experiment that cleanly demonstrates the point one wants to make.

Thus, during the last few years I began more and more to like "experimental economics". This lecture is in a sense my concluding comments on my brief detour into the world of experiments. This detour will lead me back, at the end of the lecture, to the original question: Can we make do with introspection?

My view on experimental economics is derived from my view of economic theory. I should mention, that when I speak of economic theory, I am referring primarily to decision theory and game theory. This is not because they are the most important fields in economic theory but simply because I feel more familiar with them. It probably seems strange that we need to talk about something so fundamental as the meaning of economic theory. But let us face it, there is quite a bit of confusion about the goals of economic theory even within the profession. Economic theory lacks an agreed-upon objective and interpretation. Again and again we find economic theorists asking themselves the uneasy question: "What is it that we are doing?"

It is generally agreed that economic theory is a part of economics and that economics is a study of the real world. It is not a branch of abstract mathematics even though economic theory uses mathematical tools. Some people look to game theory and decision theory to provide them with a "guide for behavior". They look at theory as the basis for "essential training in making choices and weighing possibilities … not only in business but also in daily life". Some game theorists think that theory can "help you improve your skills at discovering and using effective strategies". Two very distinguished economic theorists whom I highly regard have even declared that: "Our aim is to improve your strategy IQ".

I doubt that economic theory delivers the goods. We do not make predictions anything like those made in the natural sciences. The link between economic theory and practical advice is tenuous, if it exists at all. Academic economists like to emphasize the usefulness and applicability of what they are doing. This might be a result of guilt, or because we honestly want to save the world, or that we have a vested interest in this position or perhaps because it is indeed useful. However, let me say that after so many years in the profession I have yet to see a case in which a game theory or choice theory model (including of course my own…) contributed an insight that clearly should have influenced the real world. I cannot think of a case that I can use to convince the skeptics. Even if they do exist I doubt their benefits "justify" the investment societies make in our profession.

By the way, though I am rather doubtful about the applicability of economic theory, it is not a cause for my concern. In other words, I am not sure that applicability is desirable. If microeconomics is useful, the first to benefit will be the MBA students who are among the last people in the world I feel obliged to assist. If game theory were indeed useful it could be used for military purposes. Game theorists were employed to assist in thinking of military strategies long before telecommunication companies began hiring them. Some of the military advice given by game theorists does not make me particularly proud.

So allow me briefly present my view of economic theory (the rest of the section is based on chapter 5 in Rubinstein, 2000a). First, note a critical point about economic models (see Backhouse, 1998 for a related point): An economic model differs from a purely mathematical model in that it is a *combination* of mathematical structures and interpretation. The names of the mathematical objects are an *integral* part of an economic model. When mathematicians use everyday concepts such as "group" or "ring", it is only for the sake of convenience. When they define a collection of sets as a "filter", they do so in an associative manner. In principle, they could call such a set an "ice cream cone". When they use the term "good ordering", they are not assigning any ethical value with the word "good". In economic theory, the interpretation of a model is an essential ingredient of that model. A "game" varies according to whether the players are human beings, different "selves" of the same person or bees. A strategic game changes entirely when payoffs are switched from utility numbers representing von Neumann and Morgenstern preferences to sums of money or to measures of evolutionary fitness. The formal model is identical with any of these interpretations but the models per se are not. Assumptions we find plausible with a certain interpretation may be absurd with another.

The concepts which we attach to the mathematical symbols are the targets of economic theory. These are not the formal models that rather the concepts which appear in the interpretation. Economic theory is about the real world in the sense that these concepts are taken from our deliberations on the world. A good model is realistic in the sense that it orders our perception of real life social phenomena. It is realistic if it describes a situation as it is perceived by decision makers rather than as a presentation of the physical world. According to this approach models are not meant to be isomorphic with respect to reality but rather to the way in which the world is perceived by its inhabitants. And as economic theorists our goal is to clarify the connections between different types of concepts and arguments and patterns of reasoning. We attempt to "draw links" (a phrase often used by Aumann) and "understand", rather than "predict".

To draw an analogy, I do not believe that the study of formal logic can help people become "more logical", and I am not aware of any evidence showing that the study of probability theory significantly improves people's ability to think in probabilistic terms. Game theory, decision theory and much of economic theory in general are similar to logic and probability theory and I doubt that they could prove useful to economic agents in concrete situations. In fact, I suspect that they could even be misleading since people often ignore the subtleties of theoretical arguments and treat equilibrium notions (especially game theoretic solution concepts) as instructions. Nevertheless, I feel obliged to mention that we (me in collaboration with several graduate students in Tel Aviv) failed to demonstrate this in an experiment carried out on a group of Tel Aviv University students. We presented students before and after a course in game theory with several scenarios. I conjectured that the course in game theory actually had a negative influence on students' thinking: for example, they might use mixed strategies even in situations where it is clearly suboptimal, or, they might feel that they have less of an obligation to truthfully express their view in a ballot and that they have a license to be selfish. I am not convinced that my conjecture was wrong, however, the results of my experiments were far from conclusive.

Having briefly presented this perhaps controversial view of economic theory, I now return to experimental economics. My view about the meaning of economic theory leads me to find it hopeless and, more importantly, pointless to test the predictions of models in economic theory. However, when an economic model is based on intuitions about how people reason, experimental economics can verify that these intuitions are not extrinsic. Experiments serve as a test of the plausibility of assumptions and not conclusions. When experimental economics feeds economic models it can suggest new ideas about human reasoning in economic situations. In any case, experimental economics should relate to the plausibility of assumptions we make on human reasoning rather than trying to accurately predict human behavior.

2. From theory to experiments

I wish to discuss the following two "issues": Nash bargaining theory and hyperbolic discounting. Regarding the former, theory has led the experimental work. Regarding the latter, experiments have motivated new theory. I will use these models as the basis for making some critical comments on experimental economics and economic theory.

Consider Nash bargaining theory. It has often been said that Nash's theory is designed to provide a prediction of the bargaining outcome based on two elements: (i) the bargainers' preferences which are defined on the set of possible agreements (including the event of disagreement) and (ii) the bargainers' attitudes towards risk. Nash based the construction of his model on nothing but his own intuition that these are the important factors in the determination of real life bargaining outcomes.

The primitives of Nash's (two-person) *bargaining problem* are the "*feasible set*" and a "*disagreement point*". In Nash's formalization of the bargaining problem each element of the feasible set corresponds to the pair of numbers interpreted as the utility levels obtained by the two bargainers in at least one possible agreement. The utilities are von Neumann–Morgenstern in that they represent the bargainers' preference relations over lotteries that satisfy the expected utility assumptions. The disagreement point is modeled as one of possible agreements. A *Bargaining solution*, is defined as a function which assigns a unique pair of utility levels to each bargaining problem. Thus, a bargaining solution is meant to provide a unique "prediction" of the bargaining outcome (in utility terms) for each of the problems in it's domain.

Nash (1950) showed that the unique bargaining solution satisfying four well known axioms (Invariance to Positive Affine Transformations, Symmetry, Pareto Optimality and Independence of Irrelevant Alternatives) is (after normalizing the disagreement point to be "zero") the argmax of the product of utilities over all agreements which are better for the two bargainers than the disagreement point.

If the task of bargaining theory is to provide "clear-cut" numerical predictions for a wide range of bargaining problems, then Nash certainly achieved this goal. The fact that it is defined by a simple formula is a significant advantage of the theory, especially when we intend to embed it in larger economic models that contain a bargaining component. But can this "prediction" be tested like a prediction in the sciences? Does the formula approximately predict how people will share a pie?

I think that many experimentalists naively interpreted Nash's theory to be predictive. Researchers tried to design experimental environments in which the main factors influencing the outcome were the attitudes of the players towards risk. They tried to test whether the outcome of the Nash bargaining solution is indeed obtained. In my opinion, it would be a miracle if the Nash formula could provide a prediction of the complex activity we call bargaining. I do not find a sign of a miracle here.

The fact that John Nash formulated elegant axioms which analytically deduce his solution does not mean that human beings behave according to them. A claim that elegant axioms are true is a matter for theology. Nash created a very convenient analytical tool but try telling someone outside the economic profession that "economic theory" predicts that the outcome of bargaining will be the agreement which maximizes the product of utilities. I am afraid that they would not think very highly of economic theory once they heard that this is a jewel in the crown of economic theory… Nash's formula lacks any natural meaning. What is the interpretation of the product of two von Neumann–Morgenstern utility numbers? What is the meaning of the maximization of that product? Can we consider the maximization of the product of utilities to be a "useful" principle for resolving conflicts or predicting the way that conflicts will be resolved?

By the way, I think some people (including myself) mistakenly interpret the execution of the so-called Nash Program as a "justification" for the Nash bargaining solution. The fact that in some well-specified way, the unique subgame perfect equilibrium of the alternating offers model with risk of breakdown converges to the Nash bargaining solution where the probability of the breakdown is small … is indeed a nice result. It provides a link between two different ways to reason about bargaining. However, it in no way persuades me that the Nash bargaining solution is a good predictor of real life bargaining. It does not provide any justification for advising someone to follow the Nash bargaining solution as part of a strategy in real life negotiations, even if I knew what it meant in that situation. It does form links between abstract ways to think about bargaining, but no more than that!

Let us return to experiments. My view is that in order to approach experimentally Nash bargaining solution we need to have a better interpretation. It is the interpretation, rather than the predictions of the solution which should be the subject of the experimentation. In Rubinstein et al. (1992) we rewrote Nash theory to provide an alternative definition of the Nash bargaining solution: A solution is a "convention" which attaches a unique agreement to any bargaining problem. It embodies the assumption that players are aware that when they raise an objection to an alternative, they risk that the negotiations will end in disagreement. The Nash solution agreement is the one which satisfies the following property:

If it is worthwhile for one of the players to demand an improvement in the convention, even at the risk of a breakdown in negotiations, *then* it is optimal for the other player to insist on following the *convention* even at the risk of a breakdown in negotiations.

Or alternatively, the Nash bargaining solution is an agreement y^* satisfying the condition that any argument of one party of the type "I request x without delay; I so much want x that I am ready to take the risk that there will be a breakdown in the negotiation with probability $1 - p$", will be rebuffed by the counterargument: "Well, if we agree on x, then I would be able to make the same kind of argument against x and demanding y^*".

The Gulf War provides a concrete example of this definition: The bargainers were Iraq and the US. The set of agreements contained the various possible partitions of the land in that region. The disagreement event was a war. When Saddam Hussein moved his troops he deliberately took the chance that the situation would deteriorate into war before the US capitulated. Apparently, he preferred the risk of war with a certain probability while hoping that the US would agree to his request to annex Kuwait. In contrast to his expectation, the US preferred to take the risk of war and basically demanded a return to the status quo rather than give in to Iraq's demands. If the US had yielded to Iraq, it would have meant that the pre-invasion borders were not part of a Nash bargaining outcome.

I cannot see how this type of argument could serve as the basis for a universal bargaining theory. However, it is possible that this kind of argument makes sense to people. It is probably one type of argument, from among many, which plays a role in negotiations. This is where experimental economics should be called into action. I would like to see experimental economics test whether the type of consideration buried in the above interpretation rings true. It would be interesting to test whether the logic of the argument and counterargument described in the above definition is indeed acceptable to a wide range of people. I am not aware that such experiments have been done. This is the kind of testing we need in order to challenge theory.

3. From experiments to theory

A recent spate of papers have replaced the standard "constant discount utility function" with the following particular form of the hyperbolic discount utility function: $v(x_0) + \beta\sum_{t=1,2,\ldots} \delta^t v(x_t)$. In this functional form the rate of substitution between today and tomorrow ($\beta\delta$) is smaller than that between any other pair of successive periods (δ). In the wake of Phelps and Pollak (1968) and Laibson (1997) this functional form has been applied to a wide range of issues.

Phenomena which cannot be explained by standard discounting utility functions appear as equilibrium outcomes once the decision maker is assumed to use hyperbolic discounting. Policy questions have been freely discussed in these models even though welfare assessment is particularly tricky in the presence of time inconsistency.

The justification for the abandonment of constant discounting utility functions and the adoption of hyperbolic discount functions can be found in statements like the following one from Laibson (1996): "Research on animal and human behavior has led psychologists to conclude (see …) that discount functions are generalized hyperbolas … ". A few years later, Brocas and Carrilo (1999) wrote: "There is well documented literature both in psychology and more recently in economics showing that individuals' discount rates are best approximated by hyperbolic rather than the traditional exponential functions. We refer the reader to … for empirical support of this theory both in animals and humans … ".

In a recent paper (Rubinstein, 2000a, b) I briefly discuss the so-called "empirical support" from animal studies. In any case, the main justification for the adoption of the hyperbolic discounting utility function is empirical evidence in the cognitive psychology literature. Typical observations were first discussed by Thaler (1981): some people prefer "one apple today" to "two apples tomorrow" but at the same time they prefer "two apples in one year plus one day" to "one apple in one year". Or, Ainsley and Haslam (1992) reports: "A majority of subjects say they would prefer to have a prize of a $100 certified check available immediately over a $200 certified check that could not be cashed before 2 years; the same people do not prefer a $100 certified check that could be cashed in 6 years to a $200 certified check that could be cashed in 8 years".

Experiments of this type have been replicated and more importantly they are confirmed by our intuition. Experimental results have exposed a phenomenon which theorists have modeled and applied in a wide range of economic issues. The economic paradigm of optimizing a simple functional form remains "untouched" as the hyperbolic discounting functional form is only marginally different from the standard utility function and seems to provide an "explanation" of real life evidence.

The "problem" is that the experimental findings do not justify the selection of one of the infinite number of functional forms which are consistent with them. Hyperbolic discounting, even if consistent with these experiments, does not capture the psychology behind the experimental results. The result, as we will see, is that the basic procedural element which undermines the constant discounting model also undermines hyperbolic discounting.

My own interpretation of the experimental results relies on ideas presented in Rubinstein (1988) (within the context of decision making under uncertainty). Leland (1993) was the first to apply those ideas to "decision making with time". This approach holds that the decision maker uses a procedure which attempts to simplify the choice by applying similarity relations to cancel out "similar components". The important role of similarity in decision making was emphasized by Amos Tversky (see for example Tversky, 1977). The first formalization of a notion of similarity is due to Luce (1956).

In the present context, the objects of choice are of the form (x, t) where x is a prize received with a delay of t units of time. I think that when comparing between two pairs (x, t) and (y, s), many decision makers go through a three-stage procedure using two similarity relations, one relates to the money dimension and one to the time dimension:

Stage I: The decision maker looks for dominance: If x is better than y and $t < s$ then the pair (x, t) is preferred over (y, s).

Stage II: The decision maker looks for similarities between x and y and between t and s. If he finds similarity in one dimension only, he determines his preference between the two pairs using the dimension in which there is no similarity. For example, if t is similar to s but x is not similar to y and is preferred over y, then (x, t) is preferred over (y, s).

Stage III: If the first two stages were not decisive the choice is made using a different criterion.

The experimental findings described previously are compatible with the application of this procedure. Consider, for example, a decision maker who is applying the above procedure and determines that "today" and "a year from now" are not similar while "10 years" and "11 years" are. Then, if the decision maker is indifferent between \$$x$ today and \$$y$ in a year from now, it must be that $x < y$ and that x and y are not similar (if they were similar then the subject would prefer \$$x$ today over \$$y$ in a year from now). On the other hand, if a subject is indifferent between \$$x$ in 10 years and \$$z$ in 11 years, then it must be that $x < z$ and that x and z are similar (if x and z were not similar, then since 10 and 11 years are, the above procedure would find \$$z$ in 11 years to be preferred over \$$x$ in 10 years). If y is similar to x and z is not and both are greater than x, then one would expect z to be smaller than y.

It seems that both the hyperbolic discounting approach and the "similarity-based" approach are consistent with the evidence. However, decision problems can be designed and tested to demonstrate that while the behavior of a significant number of subjects is incompatible with the hyperbolic discounting hypothesis it is consistent with a plausible application of the above procedural approach. Here I will describe briefly one such experiment taken from Rubinstein (2000a, b).

My experimental method was cheap and simple, involving no laboratories and (almost) no monetary rewards. Students were twice approached by e-mail with an interval of 14 days in between. In each round, the students were asked to enter a web site designed for the experiment and to respond on-line to several questions. A prize of \$100 was randomly awarded to one of the participants in each round and was intended only to slightly encourage the students to spend a few minutes on the experiments. One hundred and sixty-five students responded to the first message and 145 to the second; 45% of the students in the second round did not participate in the first, making it possible to say that the fact that students participated or not in the first round made no difference. Each experiment consisted of two questions, one of which has presented in each round:

Q1 You can receive the amounts of money indicated according to one of the two following schedules:

A	April 1	July 1	Oct 1	Dec 1
	$1000	$1000	$1000	$1000
B	March 1	June 1	Sept 1	Nov 1
	$997	$997	$997	$997

Which do you prefer?

Q2 You have to choose between:

A Receiving $1000 on Dec 1st.
B Receiving $997 on Nov 1st.

Your choice is:

The hyperbolic discounting approach predicts that every subject choosing B in Q2 will choose B in Q1. If a subject chooses B in Q2, then he is ready to sacrifice $3 in order to advance the payment due in December by one month. The hyperbolic discounting theory predicts that he would find the three dollar sacrifice worthwhile in order to advance any one of the other three scheduled payments by one month.

The results contradict this prediction: 54% of the subjects chose B in Q2, while only 34%, chose schedule B in Q1.

My explanation of the results in terms of the above procedural approach is as follows: In Q1 many subjects viewed the alternative as a paired sequence of dates and $ amounts. The manipulation of subjects' behavior in this experiment was accomplished by triggering the similarity relation with regard to the sequence of dates (in another experiment, the problem triggered the similarity relation with regard to the prize dimension). Many subjects prefered A over B and D over C since they viewed the sequence of dates (April 1, July 1, October 1, December 1) to be similar to the sequence (March 1, June 1, September 1, November 1) while they found the sequence of payments ($1000, $1000, $1000, $1000) less similar to ($997, $997, $997, $997) than $1000 was to $997.

The case of hyperbolic discounting is presented here only as an example of how experimental results have influenced economic analysis. The same sort of evidence which was brought against constant discounting can just as easily reject hyperbolic discounting as well. Furthermore, the procedure based on similarity better explains the observations and is more intuitive.

The hyperbolic discounting approach misses the core of the psychological process and thus constitutes only a minor modification of the standard discounting approach while suffering from some normative disadvantages. The application of similarity-based procedures (see Rubinstein, 1988) may conflict with transitivity since the transitive closure of the partial relation as determined in the first two stages of the procedure is not likely to be consistent with the third stage which employs other principles. Thus, adopting the similarity-based procedural approach requires revolutionary changes in our theories and the development of new and original modeling devices.

4. Final words

I would like to make some comments on the methodology of experimental economics. The success of experimental economics is not disputable. Experimental economics has entered the mainstream of economics. However, there are some aspects of the field, which are deserving of criticism.

The small print of any experiment is important. Minor differences in the wording of an experiment may be crucial. The method of selecting the data which is reported in a paper may affect the conclusion. Given the relatively small samples we use, even minor mistakes made by your research assistant (or … yourself) may have a critical effect on the conclusion. I suspect that the uncertainty surrounding such mistakes is of higher magnitude than that which is put into the routinely calculated "significance measures" and render many of the "significance" calculations meaningless. Yet it is difficult for an outsider to obtain the detailed information required to assess an experiment and the data which served as the basis for its conclusions (in spite of advanced information technologies). There is almost no effective regulation of published results. We rely on our own integrity to such extent that economics (so I am told) is one of the only major academic professions that lacks a code of ethics.

My impression is based on numerous experimental papers I have read, conversations with experimentalists and a few experiments which I conducted myself. It seems to me that selecting the research questions ex post, ignoring conflicting data, eliminating problematic responses from the sample and so forth are techniques that almost no one can resist using at one time or another, especially if he feels strongly about the results he wishes to reach, results.

One would think that replications would prevent erroneous conclusions. However, our profession does not reward replication of experiments. Professional rewards are given to original, new experiments. In fact, the current incentive system deters the refutation of experimental results. Let us say you are a researcher who is interested in a paper by Prof. X who claims to have found something quite interesting. Let us say that you find the results plausible but you are not sure that the experiment was done properly and that indeed conclusion is valid. Do you have any incentive to repeat the experiment? No, because no one would publish it. Yet, you are interested in the subject matter and you probably think that Prof. X's finding is sensitive to a certain key detail of the experiment. Now you are quite eager to demonstrate your point and to publish a paper. In order to do that you have to first confirm Prof. X's basic claim. If you fail to repeat Prof. X's result, your point is lost. Thus, you approach the experiment with a desire to confirm the published result. To summarize, replications are too often conducted when the experimenter has a new point to prove and needs to confirm the original experiment. Thus, the value of these replications is less than it appears.

Another problematic practice I would like to mention in passing is the sifting of results ex post, namely after the results are gathered. Obviously, if some pattern of behavior, from among an endless number of possibilities, is discovered in the data ex post, the results are much less informative. In the absence of rules of maintaining a research protocol one cannot check whether the results were conjectured before or after the results were obtained.

As to experimental techniques, I am not happy with some of the established professional standards. I have heard researchers in the field, most of whom are young, complain bitterly that their papers were rejected (or that they believe that their papers would be rejected) because they did not fulfill the requirement to conduct "laboratory" experiments and to pay their subjects monetary rewards. Systematic comparisons of experimental results (see Camerer (1999) and the papers quoted there) leave me unpersuaded as to the need to pay subjects. (As you can see I am struggling to avoid the use of popular terms such as "it was shown by", "as was demonstrated", and "it was found", which I find to be overused in the experimental literature.) Good experiments which demonstrate commonly used pattern of reasoning "work" even when we use cheap procedures that do not require paying any monetary rewards such as "class experiments" or "posting the experiments on the Web". In some cases paying subjects will change the distribution of responses, however, assuming that we only want to confirm the existence of a plausible pattern of reasoning it seems unlikely that whether or not we pay the subjects will effect results more than an infinite number of other factors (such as, gender, age, profession, time of day, mood, etc.).

While teaching game theory to undergraduates, I conducted more than 40 experiments by posting the questions on the Web (see Rubinstein, 1999) and especially the electronic second edition of the paper posted in http://www.princeton.edu/ ~ ariel/99/gt100.html). My students were required to go onto the web site and submit weekly responses to questions of the type: "What would you do if you were playing the following game?" In almost all experiments which I replicated, there were no qualitative differences between the results reported in the literature and those from the students in my course. It is true that I elected to repeat only those experiments which seemed very robust. But, would I have had any reason to believe those experimental results which were not robust?

If one believes that people maximize expected utility and one wishes to find the "real" coefficient of risk aversion in order, for example to devise optimal insurance policies, then the real distribution of this value in the population must be found. In order to do this, subjects must be presented with concrete situations in which they receive monetary rewards. The fact that most experiments are conducted on the "representative" population consisting of mainly students in psychology and economics hints at some of the difficulties in achieving this objective. However, if the purpose of experiments is simply to confirm one's intuition, we then only need to make sure that the subjects are as attentive in the experimental environment as they are in life. For this we do not need money and certainly not laboratories.

It is hard to avoid the thought that the lack of clear standards in experimental economics and the burdens placed on experimental research procedures actually serve as barriers to entry into the field. Of course, no one purposely set up these barriers, but the fact is that they did evolve.

And by the way, I am always amazed at how difficult it is for us to apply criticism to ourselves. I am sure this is true for me: I would probably sharply criticize this lecture if I were in the audience. Economists are very sensitive to the issue of incentives. On the one hand, we send the world a loud and clear message that the world must be designed to provide people with the right incentives in order for them to behave in the manner we wish them to. On the other hand, we are quite blind to the incentives we ourselves establish in economic research. Should not we consider establishing better professional standards in order to ensure that we are serving the goals of academic research and teaching, whatever they might be?

As I stated in my opening remarks. I still find experimental economics fascinating. Often it is even fun. But do we really need experiments? If we give up on the idea that we can measure the "social world" and if in any case we heavily rely on our intuition, can we not simply make do with our impressions about how human beings reason? If the significance of experimental work relies so heavily on our honesty can we not rely on our gut feelings? Can we not pursue economics as philosophers pursue philosophy and as economic theorists used to justify their models?

I am not sure. Experimental economics provides us with a safeguard which protect us from mistaken intuitions. An economist's intuition is often distorted by his own model's assumptions. We are often blinded by our wish to obtain the goals established for us by the profession's incentives. I will refrain from drawing any definite conclusions. As usual, nothing is that clear-cut in economic methodology.

References

Ainsley, G., Haslam, N., 1992. Hyperbolic discounting. In: Loewenstein, G., Elster, J. (Eds.), Choice over Time. Russell Sage Foundation.

Backhouse, R., 1998. If mathematics is informal, then perhaps we should accept that economics must be informal too?. The Economic Journal 108, 1848–1858.

Brocas, I., Carrilo, J., 1999. Entry mistakes, entrepreneurial boldness and optimism. Mimeo.

Laibson, D., 1996. Hyperbolic discount functions, undersaving, and savings plans. Working paper No. 5635, NBER, Cambridge, MA.

Laibson, D., 1997. Golden eggs and hyperbolic discounting. Quarterly Journal of Economics 112, 443–477.

Leland, J., 1993. Similarity judgments, violations of stationarity, and reflection effects in intertemporal choice. Mimeo.

Luce, R.D., 1956. Semiorders and a theory of utility discrimination. Econometrica 24, 178–191.

Nash, J., 1950. The bargaining problem. Econometrica 18, 155–162.

Phelps, E.S., Pollak, R.A., 1968. On second-best national saving and game-equilibrium growth. Review of Economic Studies 35, 201–208.

Rubinstein, A., 1988. Similarity and decision-making under risk. Journal of Economic Theory 46, 145–153.

Rubinstein, A., 1999. Experience from a course in game theory: Pre- and post-class problem sets as a didactic device. Games and Economic Behavior 28, 155–170. (Second edition is posted in http://*www.princeton.edu/* ~ ariel/99/gt100.html).

Rubinstein, A., 2000a. Economics and Language. Cambridge University Press, Cambridge.

Rubinstein, A., 2000b. Is it "Economics and Psychology"? The case of hyperbolic discounting. Mimeo.

Rubinstein, A., Safra, Z., Thomson, W., 1992. On the interpretation of the Nash bargaining solution. Econometrica 60, 1171–1186.

Thaler, R., 1981. Some empirical evidence on dynamic inconsistency. Economics Letters 8, 201–207.

Tversky, A., 1977. Features of similarity. Psychological Review 84, 327–352.