CHAPTER 1

Introduction and overview

ALVIN E. ROTH

Each of the chapters in this volume concerns some aspect of economists' use of controlled experiments. Since the mid-1970s this kind of work has been transformed from a seldom encountered curiosity to a small but well-established and growing part of the economic literature. This transformation has been rapid. For example, when I began my own experimental work about a dozen years ago, it was most convenient to publish the results in journals of psychology and business. Today it is no longer unusual for controlled experiments to be reported in any of the major American economics journals. Experimental work has become well enough represented in the literature so that, in 1985, the *Journal of Economic Literature* established a separate bibliographic category, "Experimental Economic Methods."

However, as might be expected of any newly developing field of scientific endeavor, there are at least as many points of view about the role of experiments in economics as there are economists who conduct them. One of the reasons for this is that "economics" encompasses quite a diversity of activities and methodologies, and controlled experimentation appears to have the potential to play at least a supporting role, and in some cases a far larger part, in many of these.

At the time that I was organizing the conference that led to the publication of this volume, I was preparing a paper on experimental economics to present at one of the symposia of the Fifth World Congress of the Econometric Society, which was held in the summer of 1985. (World congresses are held every five years and are organized around a group of symposia on significant recent advances. A sign of the times, and of the distance experimental economics has come in so short a time, is the fact that this was the first such symposium to be devoted to experimentation.) In that paper (Roth, 1985) I tried to organize some of the experiments that have been conducted by classifying them according to the kinds of dialogues they were a part of, that is, according to how they were motivated and who they were

trying to persuade. Although it became increasingly evident that no classification scheme would be adequate to the task of demonstrating the relationships as well as the differences between different bodies of work, I organized the paper around three principal kinds of dialogues, which I referred to as "Speaking to Theorists," "Searching for Facts," and "Whispering in the Ears of Princes."

The category "Speaking to Theorists" was meant to include experiments motivated by well-articulated formal theories. These experiments are designed to test the predictions of those theories, as well as to reveal unpredicted regularities, in a controlled environment where observations can be unambiguously interpreted in relation to theory and incorporated if necessary into the construction of new theory. The requirement that the observations be interpretable in relation to a particular theory or group of theories imposes constraints on the experimental designs that are perhaps the chief characteristic of these experiments. Such experiments are typically motivated by the theoretical literature and are intended to feed back into that literature; that is, they are part of a dialogue between experimenters and theorists.

The category "Searching for Facts" includes experiments on the effects of variables about which existing theory may have little to say. These experiments are consequently often designed without reference to a specific body of theory, but might be motivated instead by some interesting unexplained phenomenon. They tend to be part of the dialogue that experimenters carry on with one another, and indeed many experiments of this sort are designed to help us understand earlier experimental observations.

The category "Whispering in the Ears of Princes" deals with the dialogue between experimenters and policymakers. These experiments might be motivated, for example, by questions raised by government regulatory agencies about the effect of changes in the organization of some market. Their characteristic feature is that the experimental environment is designed to resemble closely the natural environment that is the focus of the policy question at hand. In my symposium paper, I wrote, "These investigations offer the possibility of bringing scientific methods to bear on one of the traditional nonscientific vocations of economists, which is whispering in the ears of princes who require advice about pressing practical questions whose answers lie beyond the reliable scientific knowledge of the profession."

Although these categories seem useful for distinguishing different kinds of work, most extended experimental studies cannot be confined to a single category. Experiments that test formal theories may identify unanticipated phenomena for which existing theory offers no explana-

tion and suggest other experiments designed to reveal more about those phenomena. Experiments that are not motivated primarily by theoretical considerations may eventually lead to the construction of new theories and to further experiments that test them. And experiments that are motivated by questions of policy may nevertheless have some (perhaps informal) theoretical motivation and may uncover empirical regularities that require further experimentation.

In general, no simple classification scheme can do full justice to the variety of uses to which experimentation is presently being put in economics. This is at least in part because economics does not have a long-established tradition of experimental work, and experimenters are forced to develop their methodology and philosophy at the same time they investigate particular phenomena. So all these matters are in flux: Not only may different experimenters have different points of view, but any given experimental economist may approach experimentation differently over time or in different situations.

My purpose in organizing this volume and the conference that preceded it was to provide an opportunity for a number of "veteran" experimental economists to discuss and compare their views. In my initial letter to one of the participants, I wrote, "I'm hoping to get papers that exemplify the different points of view from which the six of us have conducted experiments." Although a number of these investigators have engaged in a wide variety of experimental work, I solicited papers from each one on an aspect of his work that would otherwise not have been adequately represented. Regarding this book, I wrote, "My hope is that the volume will be useful to two quite different audiences. For economists already involved in experimentation, I hope that this project will make it easier to discern the essential differences and similarities among the several approaches to experimentation now becoming well represented in the literature. And for economists and others who are not involved in this kind of investigation, such a volume should provide an otherwise unavailable kind of window on this literature."

Thus the hope for this book is that the whole will be more than the sum of the parts. Though each speaks about a particular group of experiments, together these contributions convey a preliminary picture – a snapshot in time – of the shape of experimental economics. That is not to say that the work here covers all, or even a large part, of the experimental work to date. On the contrary, no broad surveys have been included; each chapter is concerned with a specific study or series of related studies. But the *kinds* of work discussed go a long way toward exhibiting the different types of experimental work that econ-

omists are doing. Also, the contributors were encouraged to express their opinions – something that is sometimes difficult to do in journals – and they have availed themselves of this opportunity. So a reader of this book should come away with a good idea of the variety of reasons that lead economists to the laboratory.

With this in mind, readers might find the following questions useful in their reading of each of the chapters. What was it about the phenomena being studied that called for an experimental approach? What were the chief problems of design and implementation? What was the character of the data, and what difficulties were encountered in analyzing and interpreting them? How do the results change and inform our understanding of the problem that motivated the experiment and give rise to new problems?

Chapter 2 describes a series of experiments that my colleagues and I have conducted, experiments directly motivated by a body of formal theory about bargaining. These are by no means the first bargaining experiments to have been published in the economics literature, nor the first experiments concerned at least in part with the same body of theory (see, e.g., Stone, 1958; Siegel and Fouraker, 1960; Rapoport and Perner, 1974). But they do appear to be the first experiments designed so that the results can be interpreted unambiguously in terms of the parameters in which the theory is expressed, namely the preferences and risk postures of the bargainers. The fact that these data are virtually impossible to observe in natural bargaining situations makes experimental methods particularly appropriate for testing such theories, and obtaining the necessary control even in a laboratory environment presented the chief problem in experimental design. Unpredicted effects, due to variables that the theory predicts should have had no influence, were observed early in the experiments and further explored in later experiments. However, experiments designed to test predictions about those variables that the theory predicts *are* important yielded positive conclusions. So the experimental results lend support to some of the predictions of existing theory and tend to disconfirm others, while yielding a body of empirical evidence about systematic bargaining phenomena that seems likely to lead to new theory.

A point about how experimental evidence accumulates struck me as I was preparing this chapter. One of the clearest bargaining phenomena to emerge from these experiments (the "deadline effect"), which can be observed in the data of virtually all these experiments, escaped serious examination until a great deal of evidence had accumulated. Although my colleagues and I had informally observed this phenome-

non, we had not formally analyzed it and had consequently failed to realize how pronounced it is in the data. For example, I did not refer to this aspect of the data at all in my discussion (Roth, 1983) of the initial experiments in this series. Yet it now seems possible that this will prove to be a very significant bargaining phenomenon, for two reasons. The first is simply the fact that it appears so plainly and robustly in the data. The second reason (which is both more subtle and more speculative) is that this phenomenon, although unpredicted by existing theory, may lend itself to modeling with familiar theoretical tools. Thus it may help to provide a bridge between existing theoretical models and the body of empirical regularities (both predicted and unpredicted) that are emerging from this work.

Chapter 3, by Reinhard Selten, starts from a different point in the dialogue between theorists and experimenters. Selten's concern is with coalition formation in three-person games of the sort that game theorists call "characteristic function games." These are games in which, loosely speaking, the actions available to each group of players are independent of what other players may do and in which the available gains can be "monetized" and freely divided among the members of the coalition, so that the options available to a coalition can be summarized by a single number equal to the maximum amount of money available to it. These games have attracted considerable theoretical interest, because they are in some sense the simplest environments in which coalition formation can be studied. By the same token, natural economic environments can, at best, be modeled only very approximately as characteristic function games, so that artificial laboratory environments provide an otherwise unavailable opportunity to test theories of coalition formation in their simplest form. Selten considers a body of data from a number of experiments involving such games conducted under various experimental conditions by different experimenters. He identifies some empirical regularities in this diverse set of experimental data and proposes a formal theory to describe and explain them. He also proposes novel statistical tests with which to compare the descriptive accuracy of the new theory with that of existing theories applicable to the same data. (The problem of developing appropriate statistical tests for comparing the descriptive accuracy of alternative hypotheses arises in many economic experiments.) These comparisons seem promising for the new theory, which therefore seems likely to suggest further experiments. (For some related work, see Rapoport et al., 1979, and the papers referenced there.)

Two remarks about the background of this chapter seem in order. First, experimental economics in Europe seems to have developed for

a time separately from similar work in the United States. Some of Selten's early work on this topic was published in a continuing series of volumes, *Contributions to Experimental Economics*, the first of which, edited by Heinz Sauermann, appeared in 1967. (Selten earlier collaborated with Sauermann on an experimental study of oligopolistic market behavior; see Sauermann and Selten, 1959.) Second, Selten is perhaps best known among economic theorists for his theoretical work on perfect equilibria, which forms the basis for much of the current theoretical work on rational and "hyperrational" behavior. It is therefore noteworthy that the theory of coalition formation that he proposes here is a theory of *limited* rationality. His chapter emphasizes his conclusion that limitations of human rationality must be taken seriously in descriptive theories of behavior. In my 1985 symposium paper I wrote in this regard that "it is the mark of a committed scientist to be able to adjust his theoretical ideas in the face of compelling evidence, and I think that a characteristic of experimental evidence is that it will often have the power to compel such adjustments in economic theories."

The next chapter, by Richard Thaler, deals with the assumptions about individual choice behavior that implicitly or explicitly underlie most of contemporary economic theory. Specifically, Thaler is concerned with systematic deviations of individual choice behavior from the predictions of subjective expected utility theory and from predictions derived from utility theory in conjunction with various auxiliary assumptions that are often used in applying it to economic models. The study of systematic departures of observed individual choice behavior from predicted behavior began not long after the introduction of formal theories of individual behavior – see, for example, the famous papers of Maurice Allais (1953) and Daniel Ellsberg (1961). Some attempts to observe and explain certain of these departures were made in a systematic way, and theoretical explanations were offered for at least some of the observed phenomena (see, e.g., the work of Kenneth May, 1954, and Amos Tversky, 1969, on intransitive preferences). Tversky and Daniel Kahneman, among others, subsequently turned to the investigation of a wider range of individual choice phenomena that are anomalous from the point of view of existing theory, some of which they attempted to codify into a descriptive theory (Kahneman and Tversky, 1979). (An alternative approach being pursued by other investigators attempts to deal with observed empirical anomalies by extending and "reconstructing" utility theory so as to preserve many of the appealing properties of the original versions of the theory; see, e.g., Machina, 1982.)

Thaler is interested in exploring some of the implications for economic theory that would follow if we adopted the position that, because of the weight of the empirical evidence, utility theory will have to be abandoned as a useful model of individual choice behavior. He begins with a description of what he believes are the most telling phenomena that cannot be accommodated by traditional economic theory and goes on to discuss the experimental methodology that has given rise to these conclusions.

The discussion of methodology is particularly interesting, because many of the data come from subjects' responses to hypothetical questions. This contrasts with the methodology adopted by most experimental economists, who typically take pains to control for subjects' economic incentives when making choices in the laboratory, in order to guard against the possibility that subjects' verbal responses to hypothetical questions about their choice behavior might differ systematically from the choices they would make if actually faced with the indicated situation. (In this regard it seems to me that the kinds of hypothetical questions put to experimental subjects in these studies place different demands on their imaginations. Some of the questions are relatively straightforward – for example, "Which of the following options do you prefer?" Some require a little more imagination – "Imagine that you are about to purchase a jacket . . ." and some call for considerable insight – "Assume yourself richer by $500 than you are today.") However, a number of the phenomena initially identified via hypothetical questions have also been observed in experiments in which incentives were controlled. Thaler's remarks on this subject are thought provoking (as are his reflections on the trade-off between gathering data from experienced vs. inexperienced subjects). It seems likely that more experimental work will be needed to clarify the potential strengths and weaknesses of these methods for different categories of choice phenomena. In the meantime, the survey methods Thaler discusses seem to offer the possibility of obtaining at least certain kinds of data inexpensively and quickly.

The chapter by Marc Knez and Vernon Smith arose out of the discussion that followed Thaler's presentation at the conference. Smith, who had earlier spoken about another experimental study, remarked that he had some data that might be related to the evaluation of the significance, for general economic theory, of the kind of individual choice phenomena Thaler had discussed. The data Smith was referring to had to do with the way subjects responded when they were asked how much they would be willing to pay to acquire a certain object that they did not have or how much they would be willing to

accept in exchange for a certain object that they did have. One of the observed phenomena was that the price individuals said they would be willing to accept (WTA) for some object was often much higher than the price they said they would be willing to pay (WTP) for the same object. Another observed phenomenon, called preference reversal, is that subjects sometimes report a higher WTP (or WTA) for one object than for another but nevertheless say that they prefer the second object to the first.

The question that motivated the Knez and Smith chapter was, How would these phenomena be reflected in transactions occurring in a market environment? To examine this question, they allowed buyers and sellers to trade with one another in a double-auction market, after soliciting WTP information from the buyers and WTA information from the sellers. This process was carried out several times, allowing the same buyers and sellers to interact with one another repeatedly in the same market environment. Knez and Smith observed that, although anomalous WTPs and WTAs were often reported and although the associated preference reversal phenomenon persisted, these reported prices did not seem to be a reliable indicator of subjects' market trading behavior, in that traders were often observed to sell below their reported WTA and buy above their reported WTP. Although numerous WTPs violated the predictions of expected utility theory, all but a few market transactions were at prices consistent with the theory. Knez and Smith conclude that these results "call into question the interpretation, reliability, and robustness of preference reversal phenomena in the joint context of repetitive responses and market trading."

There is now a considerable body of experimental evidence about various aspects of market behavior. (An early market experiment was reported in Chamberlin, 1948.) Smith has been a prolific experimenter, and this study is representative of much of his work, with its emphasis on the function of markets as economic institutions and its close attention to the "texture" of the complex observations that result from a repeated market interaction by a group of subjects. (An account of much of Smith's work can be found in Smith, 1982, and a discussion of some of the dialogue among experimenters of which it is a part can be found in Roth, in press.)

Chapter 6, by John Kagel, is different from the others in this volume in that the experimental subjects are rats and pigeons rather than people. Kagel and his colleagues have conducted many experiments on individual choice behavior of laboratory animals in which animals make choices (pigeons by pecking keys, rats by pressing levers) that influence how much food and water is delivered to them. The "price"

of goods is controlled by varying the amount of effort needed (e.g., the number of times a lever must be pressed) to obtain a given amount of a given commodity. (Thus Kagel observes that "rats clearly prefer root beer to water at equal effort price.") Experiments of this kind have long been conducted by psychologists, and Kagel and his colleagues have found that the predictive power of economic theories of consumer demand compare favorably with those of theories found in the behavioral psychology literature. Also, economic theories have been a fruitful source of conjectures about animal behavior, which then motivate novel kinds of animal experiments. So one role of animal experiments motivated by economic theory that should be relatively uncontroversial (at least among economists) is that of establishing the extent to which the predictions of economic theory apply to, and further the study of, animal behavior.

It is another use, however, that Kagel primarily addresses in his chapter, in which he discusses the implications of animal experiments for economic theory as it applies to people. The underlying view is that many aspects of human choice behavior have a biological component that might be shared with other species of animals. Kagel reports that animals can be experimentally observed to exhibit transitive and risk-averse preferences and to share with people the kind of anomalous choice behavior observed in the Allais paradox.

He further notes that animal experiments can be designed to avoid one of the criticisms frequently leveled at economic experiments conducted with humans – that the incentives offered to subjects may not be sufficient to command their utmost attention. (Kagel cites a number of studies suggesting that experiments based on hypothetical questions may be unreliable.) However, in the experiments he discusses, the animals are maintained at approximately 80% of their natural body weight to ensure that the edible commodities involved in the experiment will be highly desirable. Thus phenomena in which strong incentives might play a critical role may sometimes be easier to study in animals than in humans. Kagel reports the results of two experiments motivated by hypotheses about the effects of poverty. The "welfare trap hypothesis" is that agents who have unearned income will get "hooked on leisure" and subsequently reduce their labor supply. Kagel reports observing a small effect in this direction among pigeons who had earlier been given "unearned income" in the form of free access to food and water. The "cycle of poverty hypothesis" is that low-income agents tend to discount the future more heavily than high-income agents, and Kagel reports the opposite effect among liquid-deprived rats.

Kagel notes that many economists question whether direct implications about human behavior can be drawn from observations of animal behavior. (Cognitive psychologists may also find this controversial for some kinds of behavior.) But he argues that this question is not logically different from the question of whether an experiment with people will generalize from one subject population (e.g., college students) to another (e.g., business executives) and that both questions lend themselves to further empirical investigation.

The last chapter, by Charles Plott, deals with experiments designed to help make and evaluate policy decisions, mainly involving economic activity subject to government regulation. Some of these laboratory experiments are thus related in spirit to the kinds of field experiments that have occasionally been conducted to evaluate policy questions. (Field experiments, which are often viewed as demonstration projects, have concerned peak-load pricing of electricity and the operation of various forms of public assistance. See Ferber and Hirsch, 1982, for a review of such experiments, the most famous of which is the New Jersey Income-Maintenance Experiment.) Plott discusses a number of experiments that he and his colleagues have conducted. His discussion is organized around the different experimental "strategies" that were adopted in response to the different ways these experiments were intended to influence the process of making and evaluating policy.

For example, Plott discusses experiments that were designed to demonstrate to policymakers, who would normally find the academic literature inaccessible or unpersuasive, some point that economists already regarded as relatively well established. Other experiments were meant to "shift the burden of proof" in an adversarial debate (e.g., court testimony), and still others were meant to guide the design of policies for new situations (such as those arising after deregulation of the airlines).

These experiments differ from the others discussed in this volume in that each was interpreted with respect to some "target" market. Consequently, Plott considers very carefully the kinds of conclusions that can be drawn about the target market from the laboratory market, which is necessarily different in many respects and inevitably much simpler than the natural target market. Such questions of interpretation are of more than academic concern, since the experiments that were meant to be introduced into adversarial debates were conducted with the anticipation that their interpretation might be actively disputed. These concerns are reflected in the design of the experiments. Since it might not be clear which features of the economic environment would be germane to the question at hand, the experimental markets were often constructed to resemble scale models of the target market. In this