Rational Choice in Experimental Markets
Author(s): Charles R. Plott
Source: *The Journal of Business*, Vol. 59, No. 4, Part 2: The Behavioral Foundations of Economic Theory (Oct., 1986), pp. S301-S327
Published by: The University of Chicago Press
Stable URL: http://www.jstor.org/stable/2352762
Accessed: 24/03/2014 18:50

Charles R. Plott
California Institute of Technology

Rational Choice in Experimental Markets*

The theory of rational behavior has several different uses. First, it is used at the most fundamental level of experimental methodology to induce preferences used as parameters in models. Second, it appears repeatedly in experimentally successful mathematical models of complex phenomena such as speculation, bidding, and signaling. Third, it is used as a tool to generate ex post models of results that are otherwise inexplicable. Finally, it has been used as a tool successfully to design new institutions to solve specific problems. When tested directly, the theory can be rejected. It is retained because neither an alternative theory nor an alternative general principle accomplishes so much.

The theory of rational individual choice has many different uses in experimental economics. The uses must be considered in any realistic evaluation of the theory. This paper is organized around that perspective.

If the only question posed is, Rational choice, true or false? then the answer is clearly false. Many critics of economics have claimed that the discipline is built on untestable foundations. Economists are indebted to psychologists for debunking such critics and demonstrating that the theory can indeed be tested. However, the gratitude can go only so far. During the process of demonstrating testability, the psychologists disconfirmed the theory. Preference transitivity experiments (Tversky 1969) and preference reversal experiments (Grether and Plott 1979) both demonstrate that the weakest forms of the classical preference hypothesis[1] are systematically at odds with facts.

* Financial support from the National Science Foundation and from the Caltech Program of Enterprise and Public Policy is gratefully acknowledged. I wish to thank Kemal Guler for his help in processing the data used in Sec. III. I also wish to thank Barry Weingast and Harvey Reed for their collaboration on the ideas and background data processing that form that section.

1. The classical hypothesis is taken to be that attitudes of preference can be represented by total, reflexive, negatively acyclic binary relations. For generalizations and alternatives to this hypothesis, see Aizerman (1985).

(*Journal of Business*, 1986, vol. 59, no. 4, pt. 2)

0021-9398/86/5904-0013$01.50

It follows that theories of markets for which rational individual choice is a necessary component are either discomfirmed by the same evidence or cannot be applied because the preconditions for application are not present. The logic is compelling, and an awareness of its existence has colored how experimental economists pose questions, how they do experiments, and what they conclude. If one wants only to "test" a theory in the sense of rejection, then one should examine its most suspect predictions. If, as part of its formulation of market behavior, a theory predicts acyclic individual choice behavior—as is the case with almost all economics models—then one seeking a disconfirmation of the theory knows exactly where to look and how to proceed. Existing experiments on individual choice behavior provide ample machinery.

The rejection of a theory of markets on the terms described above is not an especially challenging research objective. Those who study experimental markets tend to pose the questions in different ways. Rather than inquire whether a theory is true or false, they ask if the magnitude of error in the predictions of market phenomena is acceptable; or, if no concept of degree of acceptability is readily available, the question becomes which of several competing models is the most accurate, fully realizing that the best model might still be "poor." When confronted with data that suggest the existence of erratic or irrational individual behavior, the implications are immediately evaluated in terms of the possible implications for a market level of analysis. Of course, when unusual market behavior is observed, one might then turn to models of irrational individual behavior to see if they contain the seeds of an explanation.

In brief it is almost impossible to assess the importance of any problem with rationality postulates as found in experimental market studies without assessing the performance of the market models based on such postulates. In Section I, I will discuss hypotheses about rational behavior that are built directly into the foundations of laboratory market procedures. In Section II, three examples of laboratory experiments will be discussed. The accuracy of the models and the rationality postulates that form the structure of the models will be covered. Section III will demonstrate how ideas of rationality can be used to explain otherwise very confusing market behavior. Section IV will examine unusual phenomena that models of rational behavior suggest might exist, and Section V will discuss some pending problems for concepts of rationality as they are currently used.

I. Laboratory Market Procedures and Rationality

For the most part, laboratory markets are created as a challenge to theory. One research objective is to construct simple markets that are special cases of the complicated phenomena to which the models are ordinarily applied. The relative accuracies of models are assessed. The models are changed in light of the data from the special case. It is hoped, as a result, that the revised models will be more useful when applied to the complex. While other research strategies can be identified (Plott 1982, esp. pp. 1519–23), this particular strategy is frequently used.

The above objective demands that laboratory economics procedures permit some reasonably direct correspondence between parameters of models and what is controlled in an experiment. The important variables of almost all economic models are preferences (as opposed to sources of motivation), beliefs, resources, market organization (institutions), technology, commodities, prices, allocations, and incomes. If a model is to be evaluated, all these variables need to be observed and sometimes controlled. If a variable cannot be observed directly, then it is always suspected of having gone awry when the model itself does not fit the data. Of course, in this context, the preferences and beliefs are key because (*a*) they can be used to explain almost any pattern of the other variables (Ledyard, in press) and (*b*) they cannot be observed directly.

Laboratory techniques control preferences or, in a sense, allow them to be observed indirectly. The basic insight is that preferences are parameters to economic models, but the source of preferences is not a parameter. The key idea is to use monetary incentives to induce preferences for abstract commodities that exist only for the purpose of the experiment. Consider the following axioms, which are a combination of the precepts used by Smith (1976) and the axioms used by Plott (1979). If the following axioms are accepted, then preferences can be induced and controlled for purposes of experimentation.

1. More reward medium (money) is preferred to less, other things being equal (salience and nonsatiation).

2. Individuals place no independent value on experimental outcomes other than that provided by the reward medium (neutrality).

3. Individuals optimize.

Suppose, for example, that a commodity is the set of nonnegative integers, which are called units of the commodity X. Another commodity, Y, is simply U.S. currency. An individual, i, is assigned a function, $R^i(x)$, indicating the reward (dollar amount) he will receive from the experimenter should he acquire x units of the commodity. If postulates 1–3 are satisfied, then we can take as a parameter in a model, defined over $X \cdot Y$ where the operation $\cdot$ is a Cartesian product, the binary relation P_i, defined by $(x, y)P^i(x', y') \Leftrightarrow R^i(x) + y > R^i(x') + y'$. The relation P_i is the preference relation of i. If the axioms are satisfied, then P_i is in fact the individual's preference relation in the same sense that it will reflect actual individual choices from the pairs in $X \cdot Y$. Since the experimenter controls the functional form of $R^i(\cdot)$, the preference relation of each individual can be controlled as desired.

Carrying the example further, we could view $R^i(x) - R^i(x - 1)$ as the willingness to pay for additional units of X. In some circumstances the difference would be interpreted as an (inverse) demand function. That is, supose p is a constant price that a subject must pay for units of X. An optimizing subject would want to maximize $R^i(x) - px$. The optimum occurs (ignoring the problem caused by the discrete formulation) at the point $\hat{x}$ such that $R^i(\hat{x}) - R^i(\hat{x} - 1) = p$. Solving the equation for $\hat{x}$, we obtain a function, $\hat{x} = D_i(p)$, which can be interpreted as an individual demand function for X.

Notice that, if any of the conditions, 1, 2, or 3, is not satisfied, then a key parameter is misspecified. When asked to choose over $X \cdot Y$, the subject's choices would not be those predicted by P_i. If this occurs, and if the experimenter is not aware of the problem, a model might be discarded as inaccurate when in fact the experiment was not properly controlled. The point to be emphasized is that a theory of rationality is basic to experimental procedures and to the interpretation of the results. If rationality is not reliable behaviorally, then one would expect economic models to be poor predictors of experimental market behavior because the basic parameters of the economic models would not be controllable.

The nature of the argument just outlined suggests a first line of defense that can be used by anyone whose pet theory has been abused by experimental data. Were the payoffs of a sort that assures that postulate 1 is satisfied? For the most part economists have used money in amounts that will accumulate to amounts comparable to wage rates (for equivalent time) of employed members of the subject pools. Typically, this amount is between \$8.00 and \$20 per hour. A failure to provide adequate incentives is known to affect results at a group level of performance in ways that do not disappear with large samples.[2] Results regarding the importance of incentives when studying individual choices have been mixed. For example, Grether and Plott (1979) found no incentive effects. The most recent study is by Grether (1981), who demonstrated that the instances of seemingly confused behavior go up when incentives go down.

The second postulate substantially differentiates those who study markets from those who study individuals. Psychologists frequently

2. Only two examples exist. Once problems were detected along this line, subsequent experiments used more incentives. The committee experiments studied by Fiorina and Plott (1978) used incentives as a control. Means and variances were affected substantially. Plott and Smith (1978) demonstrate that traders tend not to trade units for which positive profits will be made. Just breaking even is not enough.

use rich descriptions of situations to elicit responses. From an economist's point of view, this practice is one that is to be viewed with suspicion.[3] Data that lead to a model's rejection can always be explained away by hypotheses that take advantage of any ambiguity that might exist over what preferences "really" existed in the experiment.

The final condition, 3, depends not only on human nature but also on whether the subject understands the relation $R^i(\cdot)$. This function is seldom simply verbally communicated to the subject. If the function involves random elements, they are made operational with real random devices (the word "probability" is not used). Subjects are given experience with the properties of $R^i(\cdot)$ and tested on their understanding of it. Sometimes the instructions of a complicated market experiment involve exercises in which subjects choose over $X \cdot Y$ or its equivalent as a check on conditions 1–3. While these precautions are taken as a defense against disgruntled theorists who might dismiss results on the (self-serving) claim that the preferences were not controlled, they also comment on implicit assumptions about the nature of rationality: intelligence is important; verbal communication is suspect; analytical and cognitive abilities are not dependable over experience. So the experiment proceeds, allowing for the possibility that individuals might be satisficers in the Simon (1979) sense and fail to explore the nature of $R^i(x)$ if left to their own devices.

Acknowledged problems with the concept of rational choice have shaped experimental market procedures in still a third way that was mentioned in the opening paragraphs. Almost all economic models postulate the existence (on an "as if" basis) of a transitive preference over lotteries. Thus transitive choice over lotteries can be viewed as a prediction made by the models. We know from Tversky's (1969) work on transitivity and from preference reversal experiments (Grether and Plott 1979) that those particular predictions of the models will be disconfirmed; that is, we know that models of this type make predictions that are wrong. Logic thus compels us to realize that the "truth" of the models is not necessarily the only goal of the research effort because we already have the answer to that question. Instead the research question becomes the degree to which one model is better than another at capturing market behavior. Experiments should be designed to make comparisons among models whenever such comparisons are

3. I am aware of one documented example of a problem caused by the descriptions of the alternatives. In Cohen, Levine, and Plott (1978) subjects were involved in a voting experiment. The objects of choice (letters of the alphabet) were labeled in humorous ways. Traditional financial incentives were also operative. The group-choice model, which had worked well in other experiments, was not working well, so subjects were asked to explain the reasons for their votes. The recorded votes and the reasons given by subjects indicated that subjects neglected the financial incentives and chose in ways they imagined reasonable in light of the humorous description of the options.

possible. Which model throws light on market behavior? Which model is true is a different question.

II. Performance of Market Models

If rationality assumptions are totally unreliable, then one would expect market models based on them to be similarly unreliable. Preferences for outcomes might be induced by the procedures outlined above, but it does not follow that the market supply and demand functions can be constructed from those preferences. An uncontrolled aspect of rationality is required to go from preferences to market demand; or demands and supplies might have been controlled, but laboratory markets are complicated and involve expectations formation, strategy, and so forth. The demand and supply model itself might not work as a predictor of price; or events in the market could override the incentives used. People simply might not be able to cope or might become irritated or frustrated so easily that no market model would work. If people are erratic and/or irrational, the induced preferences will not guarantee the accuracy of economic models.

Three different examples of market experiments are now summarized. Each relies on different features of human capacities. All are "success stories" in the sense that a mathematical model based on principles of rational choice seems to capture much of what is observed. The replications of these experiments have occurred in enough similar situations that the inferences drawn from the examples probably reliably reflect the facts as opposed to outlying or fortuitous observations.

A. Middlemen

The first example comes from a paper by Plott and Uhl (1981). The concern was with middlemen. Each of a group of suppliers was given a marginal cost function by application of induced preferences theory. If price was constant and each followed the competitive optimizing response, the market supply curve would be as shown in figure 1. Similarly, final buyers each had a derived demand. Should final buyers have responded in an optimizing fashion to a fixed price, the market demand would have been as shown in the figure. Each agent was assigned a different number to use as a name during the experiment. The numbers on the market demand and supply functions refer to the agent who had the limit value at the indicated level.

Final buyers and suppliers were in different rooms and could neither trade nor communicate. A group of four middlemen (speculators) were allowed first to visit the suppliers' room, at which time a market, A, was opened. Having acquired inventories, the middlemen were then taken to the final buyers' room, where they were able to sell in market

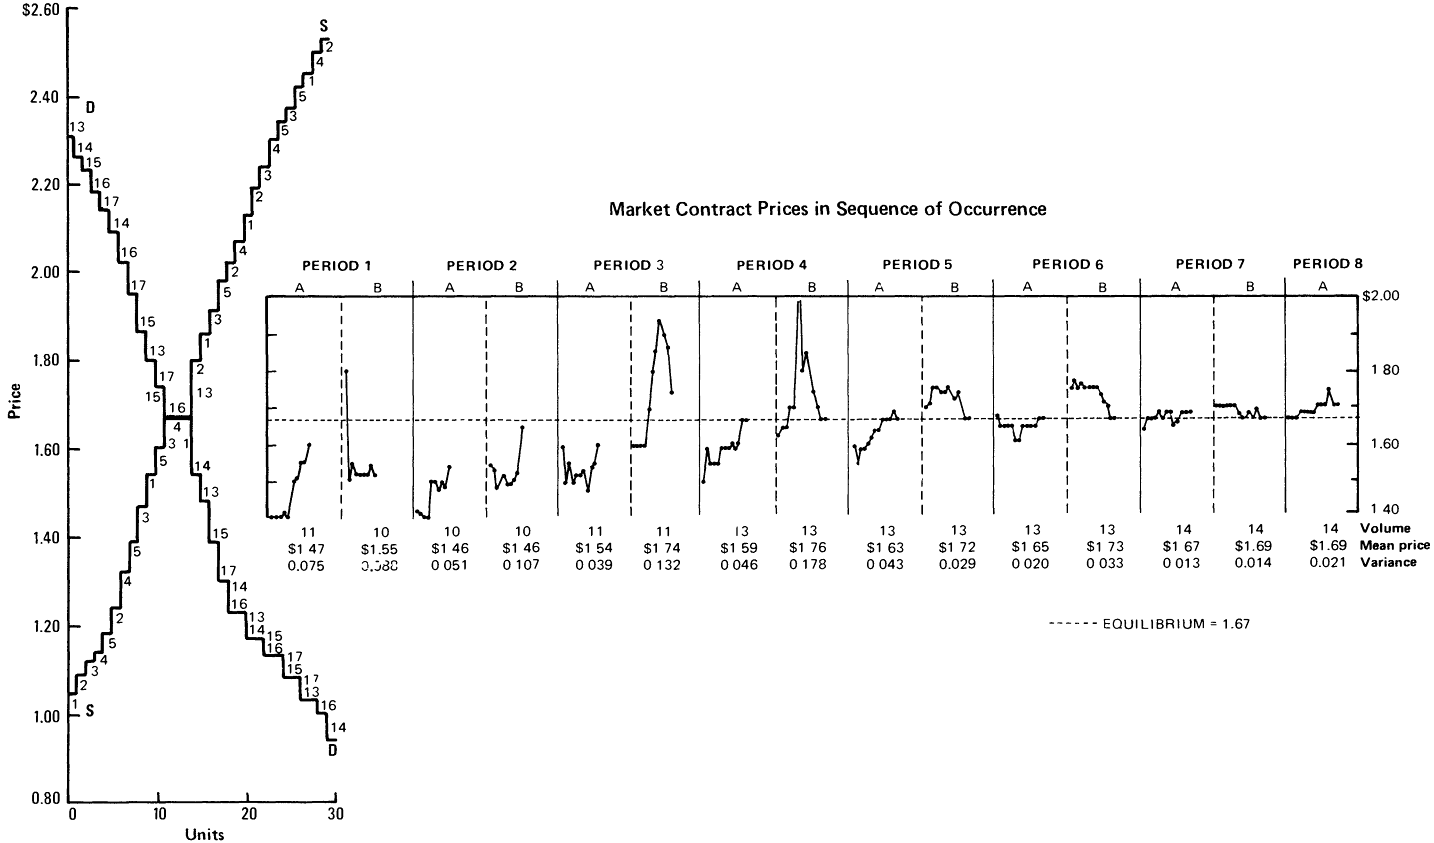

FIG. 1.—Market parameters and price time series. A = middlemen buy, and B = middlemen sell

B what they had previously purchased in market A. After market B was closed, the middlemen returned to market A to start a new period (of two markets). Inventories could not be carried forward to succeeding periods. Everything acquired in A had to be sold in B or forgotten. Both markets were organized in a manner similar to oral double auctions.

If the individuals serving the middleman function were optimizers and correctly assessed the probability of sales in market B, then the prices in markets A and B should have approached equality. The level of price should have been at the demand and supply intersection. Profits of middlemen should have approached zero. The volume in both markets should have been 14 units. It would be as if demanders and suppliers were in the same market and middlemen did not exist.

As shown in figure 1, the predictions of the model are approximately correct, and there is a time series of all contract prices in the order in which they occurred. With time and replication of periods the prices in both markets converged to the predicted equilibrium price of $1.67. The predicted carryforward of 14 units was close to the actual volume. Profits of middlemen dissipated to near zero as predicted.

From a practical perspective the competitive model works rather well when applied to the middleman markets. No model of which I am aware, based on principles other than some form of rationality, does as well as does the competitive model. Indeed, in this simple example there are many chances for things to go wrong. First, notice that the theory of derived demand is working twice removed. Derived demand theory was used to postulate the market demand of final consumers as induced by the experimenter. Derived demand for a factor of production was used by middlemen when they purchased a "resource" from the sellers in room A and transformed it into a product for sale in room B. Notice that this transformation took place under conditions of extreme uncertainty. The middlemen did not know the demand function, prices, or any other aspect of the market (or market theory). Somehow they assessed the consequences of their actions with reasonable accuracy. Once having acquired inventories, the middlemen showed no evidence of falling prey to the sunk-cost fallacy. When mistakes appeared to have been made, that is, when middlemen seem to have carried too much forward, the middlemen readily sold at a loss and recovered as much as possible. (In the first period, 1 unit was carried forward and not sold, but in subsequent periods this problem never occurred.) Notice also that we have some confirmation of the "free riding" or "prisoner's dilemma" model as applied to public goods. Middlemen had a common interest in keeping prices low in market A and high in market B. Outbidding a fellow middleman and gaining the associated personal profits is the market analogue of free riding. These participants were not characterized by such a concern for fellow middlemen that they would forgo advantages of individual gain in order that the profits of all might be higher. Not only is there support here for a "rational" perspective, but there is support for the additional proposition that these people in this setting were not naturally concerned about others; or, if they were so concerned, it was not apparent in their collective actions.

While the middleman type of market experiment suggests that elements of the rational agent model can capture much of the actual human behavior, we cannot assume that this is the end of the story. Even this simple market exhibits behavior that at best is not predicted by the model and at worst is wholly inconsistent with the model. Notice first that the model becomes accurate only after a process of convergence. The model says nothing about that. Notice that the adjustment process contains events that are hard to reconcile with rationality. In the A market, prices existed that were substantially below those observed in the previous A market, and prices tended to move up during a period. Why did sellers simply not wait and capture the higher prices? Why in period 7 did the buyer pay $2.31 when such high prices had never been necessary before? Notice that in market B of periods 1–2 an excess demand existed but that prices were below equilibrium. The model predicts equilibrium, and in periods 1–2 this did not occur.

B. Auctions

Some of the most extensive use of the precision afforded by rationality postulates is found in the auction literature. This example is of special importance because it is the only example of which I am aware that the full implications of rationality axioms have been deduced in operational terms in a form that can be examined by an experimenter. Put another way, this is the only example in all of economics where a reasonably complete theory about rational behavior in markets exists.

Compare two types of sealed-bid auctions in which a single item is to be sold. Each bidder tenders a single bid in private that is collected and examined (privately) by the market (auctioneer). The object will be awarded to the highest bidder. If the auction is a first-price auction, the winning bidder will pay the amount of his own bid. If the auction is a second-price auction, the winner will pay the amount of the second-highest bid.

The scientific challenge is to compare the bids tendered in each type of market and, more ambitiously, to predict the bids tendered. Suppose that N agents are participating and that all participants know that v_i, the value of the object to each bidder i, is drawn from a probability distribution with support on the interval $[\underline{V}, \overline{V}]$. Notice three aspects of the challenge. First, the institution can be viewed as a treatment variable, so, even if the theory fails to predict individual agent behavior, it still might add insight about market behavior. When dealing with economics, the role of the market as an aspect of inquiry should always be kept in focus. Second, it is the actions taken by agents and not their thoughts, thought processes, feelings, or attitudes that are to be studied. Finally, the concepts of value and probability that are frequently a cause for concern by critics of economics are built into the theory at the outset.

An experimental approach to the problem was first developed by Coppinger, Smith, and Titus (1980) and has since expanded dramatically. To appreciate the role of rationality in this investigation we will consider only a simple case. The values v_i are independently drawn from a constant density on [0, 1], so, by expressing bids as a fraction of the largest possible value, any interval can be considered. Each agent knows his own value before bidding but not the value of others. The above facts are public knowledge and can be controlled for experimental purposes; that is, auctions can actually be created that objectively have the requisite properties.

How might one go about developing a model of the situation? The auction theory literature suggests that the system will behave as if the following are true. (*a*) Agents choose in accord with the expected utility hypothesis. To obtain a model that can be solved we will assume each player has a utility function of wealth, $U_i(y) = y^r$, where r is distributed across the population by a publicly known probability distribution, ϕ, on [0, 1]. The constant r is a risk-aversion factor. This assumption will be treated as a maintained hypothesis for purposes of analyzing the data and testing the theory. (*b*) At the time of choice each agent, i, knows (v_i, r_i), his own value and risk parameter, but knows only the probability distribution from which those of others were drawn. (*c*) Each individual follows Bayes's law in forming expectations. (*d*) Each individual will choose a Nash equilibrium bidding function. (*e*) There are N agents.

Under all the above assumptions the symmetric Nash equilibrium bidding functions are

$$b_i = \begin{cases} v_i, & \text{for all } i \text{ if the second-price auction is used;} \\ \dfrac{(N-1)v_i}{N-1+r_i}, & \text{for all } i \text{ if the first-price auction is used.} \end{cases}$$

The comparative institutional prediction is that the expected price under the first-price auction is greater than the expected price under the second-price auction. Table 1 reproduces the results of some of Smith's experiments. The range of the support function $[0, \bar{V}]$ was varied with N to keep expected profits, as calculated by the model, the same as N increased. First, notice that the model is very accurate when applied to the second-price auction for $N > 3$. For example, if $N = 6$, the model predicts a mean price of 12.1, and the actual price averaged

TABLE 1 **Theoretical Predictions and Means and Variances Pooled over *N* Markets**

	First-Price Auction		Second-Price Auction	
Number and Statistics	Observed Price	Risk Neutral Theoretical ($r \equiv 1$)	Observed Price	Theoretical
3:*				
Mean	2.44	2.5	1.97	2.5
Variance	.589	.384	.759	.96
4:†				
Mean	5.64	4.9	...	...
Variance	1.80	.96	...	...
5:†				
Mean	9.14	8.1	...	...
Variance	1.37	1.83	...	...
6:†				
Mean	13.22	12.1	11.21	12.1
Variance	4.31	3.0	8.20	6.4
9:‡				
Mean	31.02	28.9	27.02	28.9
Variance	4.91	8.38	18.66	18.85

SOURCE.—Cox, Roberson, and Smith (1982).
* $N = 70$.
† $N = 60$.
‡ $N = 30$.

11.21. The predicted variances are also close to those observed. As predicted by the model, people tend to bid their value when they participate in the second-price auction. Second, notice that the prediction about the market treatment variable is also correct. The average price for the second-price auctions is below the average price of the first-price auctions for every value of N. The first-price auction generates more revenue as predicted.

The risk-neutral model ($r = 1$) tends to develop inaccuracies when applied to the magnitude of first-price auction bids. Of course the risk-neutrality parameter was not controlled in these experiments. In any case, prices in the first-price auction are higher than those predicted by the model if we assume that $r = 1$. If the data are tested for every value of N against the risk-averse model, which predicts that observed prices will be above the risk-neutral prediction, the model cannot be rejected for $N > 3$.

The support for the Nash-equilibrium-based models has continued as research has expanded to a study of the multiple units case, although the model has encountered difficulties for some values of N. For the single-unit case, however, the full Nash equilibrium model with all its implicit and explicit rationality assumptions is the most accurate model that exists. To the extent that the model places restrictions on data it is consistent with the facts in an absolute sense.

C. *Signaling*

The third example is a demonstration that the equilibrium notions motivated by concepts of optimizing behavior can capture the essence of very complicated and interdependent phenomena. The model itself was originally motivated by a cynical view of education (Spence 1974). Imagine a world in which education has no intrinsic value but is very costly in terms of time and effort to all but the smartest people. By paying an appropriate premium for educated employees, employers can make education a profitable investment for smart people but not for others. Thus the employers can hire just the people they wish (smart) by paying a premium for an attribute they do not value (education). Theoretically, the employers can do this even though the intelligence level of the prospective employee prior to employment can be observed by no one other than the employee himself; and, when asked, a prospective employee has an incentive to lie.

The point of the exercise is not to explore the appropriateness of the reasoning when applied to investments in education. The purpose is to explore the nature of equilibrium when such asymmetric information exists in markets. We inquire about the appropriateness of equilibrating principles that are asserted to be operative and the ability of mathematical statements to capture them.

The example was intended only to help one understand the laboratory market that was created. In the laboratory market several sellers have 2 units each of a commodity that can be sold. The units have two characteristics: grade, which can be either Regular (R) or Super (S), and quality, which initially is zero but can be added by the seller. Grade is like a {dumb, smart} variable, and quality is like education that can be added at cost. A seller's 2 units are either both R or both S. Half the sellers have R's, and the other half have S's, as determined randomly and secretly before any trading begins. Before purchase, N buyers can observe quality added, but the underlying grade is discovered only after purchase and after the market period is closed.

Buyers like Supers better than Regulars, and buyers place some value on any additional quality added by sellers (i.e., education has some value). In particular, for each unit purchased buyers have the following value (determined by the experimenter by using the techniques of induced preference described in the introduction):

$$V(g, q) = G(g) + Q(q),$$

where

$$g \in \{R, S\} \equiv \{\text{Regular, Super}\};$$
$$q \in [0, \infty) = \text{quality added by seller};$$
$$G(g) = \begin{cases} \$2.50 & \text{if } g = S, \\ \$.50 & \text{if } g = R; \end{cases}$$

$$Q(q) = \begin{cases} \$.205q - \$.005q^2 & \text{if } q \leq 20, \\ \$[(.205)(20) - (.005)(20^2)] + \$.01q & \text{if } q > 20. \end{cases}$$

Sellers face costs of adding quality of \$.15q and \$.02q if the units are Regulars or Supers, respectively. It costs less to add q (get educated) if the unit is a Super (smart).

The most efficient signaling equilibrium is a fascinating concept when considered from a rationality perspective. The equilibrium is defined by the following equations:

all Regulars will be produced at the same quality, q_R, and will sell at the same price, P_R;
all Supers will be produced at the same quality, q_S, and will sell at the same price, P_S. (1)

The two conditions in (1) follow from an underlying axiom requiring that no arbitrage exists. If different prices and qualities existed within grades, then profit opportunities would exist, and rational agents would take advantage of them:

$$\begin{aligned} V(R, q_R) &= P_R; \\ V(S, q_S) &= P_S. \end{aligned} \tag{2}$$

The equations in (2) pick up two aspects of behavior. First, having observed the quality level, q_R or q_S, the buyer can infer the grade, R or S, with certainty. Quality and grade are perfectly correlated. Second, once this is known the demand and supply model under certainty becomes applicable. For any unit with characteristics (g, q) a horizontal demand exists. Recall that the values of consumers were defined per unit, so, without budget constraints and with prices below value and no uncertainty, the buyer would want an infinite quantity of all possible commodities. The limited supply (vertical supply curve) and horizontal demand curve drive prices to the maximum, that is, the demand price:

$$\begin{aligned} P_R - .15q_R &\geq P_S - .15q_S; \\ P_S - .02q_S &\geq P_R - .02q_R. \end{aligned} \tag{3}$$

The two conditions in (3) require that truthful revelation is incentive compatible. Regular sellers maximize profits by selling units at the quality level recognized by buyers as Regulars. Super sellers maximize profits by selling units at quality levels recognized by buyers as Supers:

$$\max\{[V(R, q_R) - .15q_R] + [V(S, q_S) - .02q_S]\}. \tag{4}$$

Condition (4) captures a type of "market rationality." It says that profits of the system will be maximized subject to the behavioral constraints defined in (1)–(3).

In less opaque terms, the final condition (4) can be interpreted as another type of demand and supply condition. The q_S and q_R will be adjusted to reflect gains from exchange. The maximization formulation captures the idea that this adjustment in the quality levels of the commodity will continue until further adjustments would negate the signaling value implicit in (1). The idea is explained geometrically in figure 2. The value functions for a single buyer are drawn from Regulars and Supers. The increases in value with additional quality are as shown. Equation (2) says that the price of an S will be along the top curve and that the price of an R will be along the bottom curve. The qualities, q_R and q_S, must be such that they are not equal and thereby signal to the buyer the underlying grade. The qualities should also be located such that sellers of R's have no interest in marketing their units at (P_S, q_S), and so forth, as demanded by (2). Finally, q_R should be located to maximize system profits, and q_S should be the minimum possible level consistent with (2). A check of the equations will demonstrate that $(q_R, q_S) = (6, 27)$, as shown in the figure, have the requisite properties.

Twelve markets with the above (and related) parameters were reported in Miller and Plott (1985). The results were mixed in the sense that other variations of the model outlined above were more accurate than was that particular model. However, the interesting thing from the perspective of this paper is that the model captured any of the market data at all; yet in two of the 11 markets this complicated model that is filled with rationality postulates is very accurate. The data points are shown in figure 2 near the predicted equilibrium. The quality of Regulars, q_R, is correct, and the quality of supers, q_S, is a little too high. Variances in qualities and prices are very low. Prices are slightly below the predicted level, reflecting a frequently observed property of markets that agents will not trade for zero reward. Behavior of the type described in this model is certainly not beyond human or market capabilities.

III. Ex Post Rationalization (Reparameterization)

When markets perform in unusual or unexpected ways, the rationality postulates suggest hypotheses to explain why. The econometrics and field studies literature are filled with ex post rationalization techniques, but very little has been said about them in laboratory economics papers.

The idea of reparameterization is important in a second way. Rationality at a market level of analysis can be separated from rationality at the individual level. Suppose that the market model works well given the individual agent's personal decision rules. So, from observed market behavior we can make some reasonable inferences about what actual individual decision rules must have been. Suppose further that from induced preference theory we obtain an independent theoretical idea about what a rational individual's decision rule would have been in

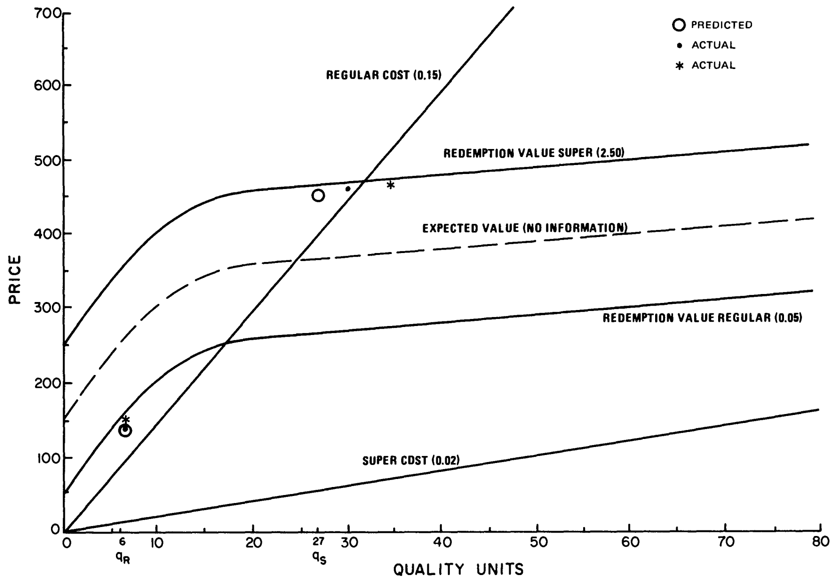

FIG. 2.—Model parameters and predictions displayed with actual experimental outcomes.

the experiment. By comparing the inferred actual with the theoretical rational we can perhaps develop a methodology for testing the latter as they are relevant for economics.

To demonstrate how rationality principles can be used in this capacity, the data from four experiments are analyzed. These are experiments that would have been discarded because of (allegedly) poor experimental control. These experiments were done in the mid-1960s and early 1970s before some of the experimental techniques currently used had been developed. They are all oral auction markets that differed in various ways from the oral double auction now in use. Each trader had 2 units and could tender an all-or-none offer. Offers remain open until canceled or taken. The instructions were not administered carefully. No tests on procedures or practice rounds were allowed. The accounting was not checked for confused or cheating participants, and so on. In essence the current operational procedures for making certain that subjects understand the reward medium and the market technology were not followed—or so we would like to believe.

The nature of the markets was to induce simple demand and supply functions different from those that had previously been examined. Also present were multiple units, which, at the time of the experiments, had not been studied. The question posed was whether the observed prices and volume would converge to the equilibrium predictions of the model.

The answer was a rather resounding "no." The initial parameters are shown as the solid-line demand and supply functions in figures 3–4. For the most part the data are well removed from the predictions of the model.

Whenever data are trashed, a danger exists that the problem is the principles that guided the models, not the lack of parametric control. However, when examining subjects' decisions, many seemed to violate the intuitive notions of rationality stemming from confusion or a willingness to violate the rules of the market.

An exercise was undertaken to "reparameterize" the experiments. We wished to provide a method of adjusting individual preference parameters in light of their choices and determine the extent to which the revised market model fits the data. The rules used were as follows. (*a*) If a buyer (seller) buys (sells) a unit for more (less) than the redemption value (cost) of the unit, then the limit price is adjusted to the transaction price. (*b*) If an agent never bid or traded during the entire experiment and passed up profitable opportunities (suitably defined), then the parameters are adjusted as if the agent were not present. (*c*) If an agent failed to trade for 2 consecutive periods and passed up profitable opportunities (suitably defined), then the limit prices are revised to equal the highest (lowest) bid to buy (offer to sell) that the agent tendered or accepted for that unit in any period throughout the experiment. (*d*) If an agent transacts for more units than the maximum permitted, then the units are adjusted to the maximum number of such extra units traded in any period, and the limit prices are the highest (lowest) price paid (received) for those units during the entire experiment.

The revised demand and supplies are the dotted curves in figures 3–4. The price predictions of the revised model fit much better than they do in the original in three of the four cases, and in the fourth case the price predictions are identical. The volume figures are worse after reparameterization because in all cases the actual volume was low relative to the original model and because the revised parameters predicted even lower volume.

The exercise demonstrates two properties of rationality-based theories. First, the adjustment of parameters need not induce circularity in the reasoning. Ex post theories based on rationality can certainly be rejected. For example, the observed volume can be used to reject the revised model. Second, in view of subsequent experimentation, the decision to discard the data was probably correct. If these subjects are "equivalent" to those used in subsequent experiments, and if the market organization had no special effects, then the actual preferences used by the subjects were not those the experimenter attempted to induce. If preferences are stable, we know now that under the double oral auction prices converge to the competitive equilibrium. Thus sub-

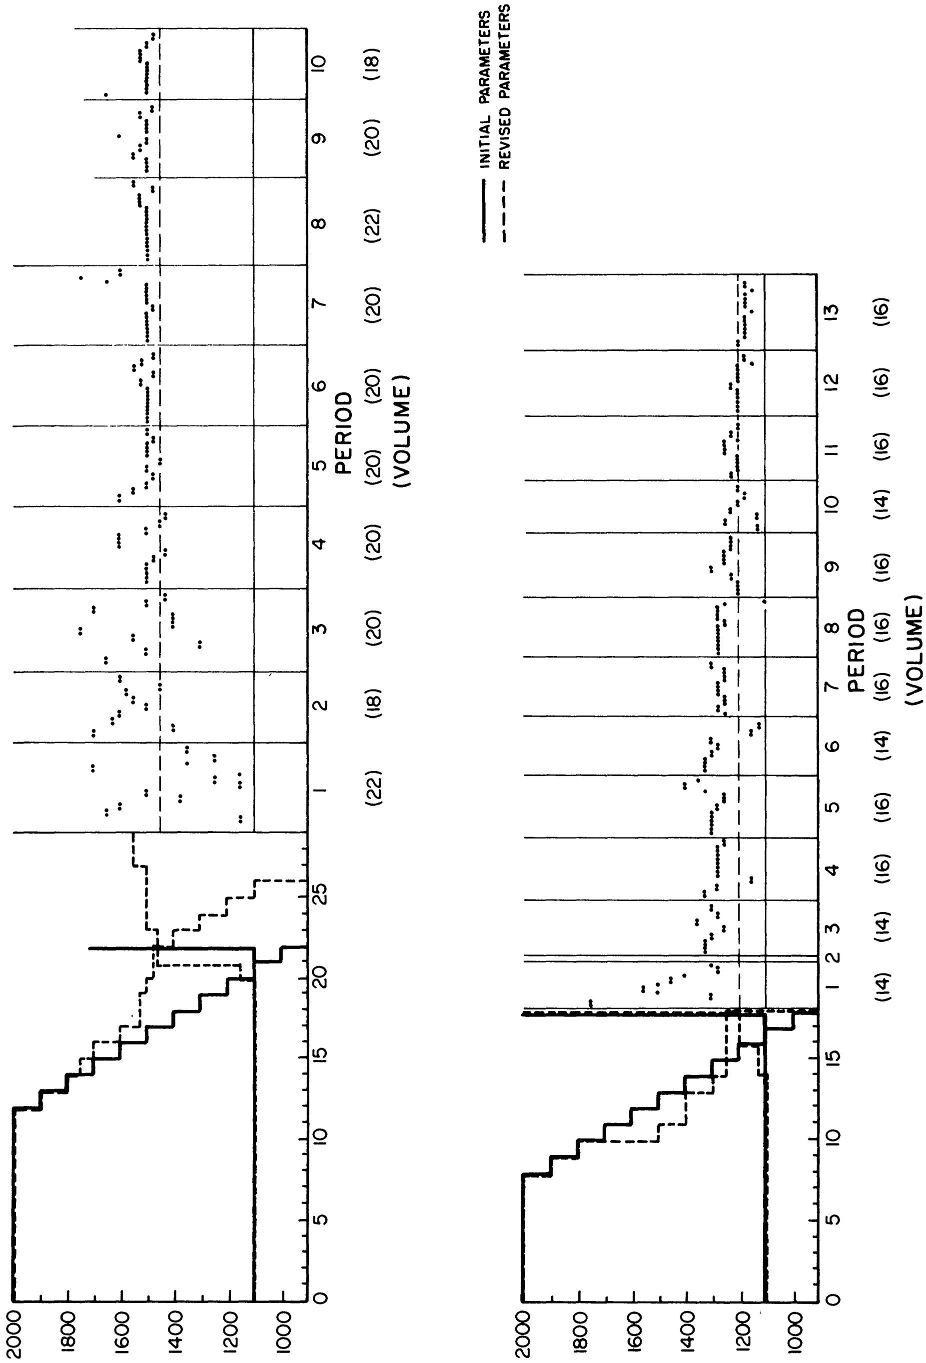

FIG. 3.—Parameters of initial and revised models, predictions, and price time series of two experimental markets

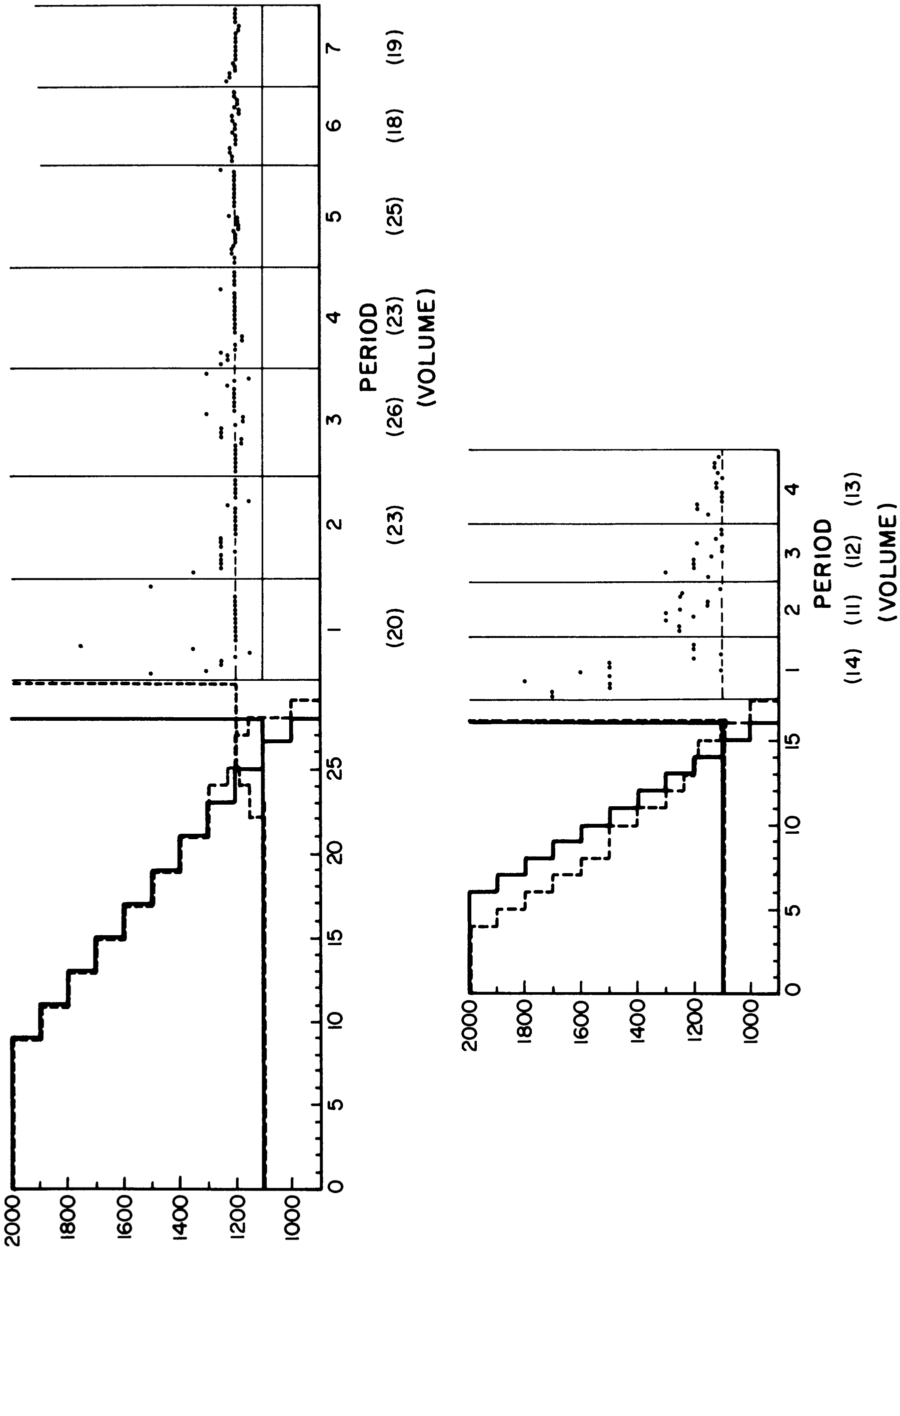

FIG. 4.—Parameters of initial and revised models, predictions, and price time series of two experimental markets. Solid line represents initial parameters and broken line revised parameters.

sequent experiments tell us that the markets in the figures above had adjusted to the actual preferences. Even if participants are confused and "irrational" from certain perspectives, the market model can still be applied.

The only other attempt to revise parameters of a market model based on decisions made during the experiment is the ongoing work by Knez, Smith, and Williams (1985). They have attempted to measure individual attitudes during a market, and to use those parameters for prediction they tested a market model based on measured parameters against the model with parameters as specified a priori by the experimenter. The markets themselves are for lotteries in which subjects stated their limit prices—maximum (minimum) willingness to pay (sell)—prior to the opening of each market period. Their conclusions are (*a*) that the act of measurement does not appear to affect the market; (*b*) that the market model drawn from the measured parameters is more accurate than is the model constructed from induced preferences; and (*c*) that many subjects (in the 40% range) exhibited a willingness to violate their own stated limit prices. Knez et al. suggest that the elicited parameters are analogous to guesses about how subjects will trade or, perhaps, are similar to wishes as opposed to true limit prices. Nevertheless, the measured parameters improve predictions about market prices.

IV. New Institutions

The rationality postulates have been useful in suggesting new institutional arrangements that have never before existed. The research on public goods provision mechanisms is a good example. Other examples include the work by Grether, Isaac, and Plott (1981) on the allocation of landing rights by auction or the work by Rassanti, Smith, and Bulfin (1982) on a combinatorial auction to solve the same problem. Experimental methods have been the only source of data about how these new institutions might perform.

An interesting example, with possibly limited social usefulness, is the unstable dollar auction.[4] The idea is to create processes that applications of rationality theory suggest will have bizarre properties. In this case the objective is to attempt to sell a dollar to perfectly informed people for much more than a dollar. Intuitively, it seems that rational consumers would never do such a thing, but intuition is not always a reliable scientific tool.

Subjects, after having attended an economics experiment, are fre-

4. This auction process first appears in print in Shubik (1971). In conversation Shubik tells me that he hesitates to take full credit for having invented the process because many unusual processes were proposed in conversation among game theorists at Princeton in the early 1950s. The theorists were using game theory to invent processes in which rational behavior by individuals would lead to surprising behavior.

quently in a room calculating their earnings. Having calculated their earnings and having not yet been paid, a dollar auction is announced. Subjects are carefully told that an English auction will be used. The market will stop if 45 seconds elapse after a bid with no intervening bid. The dollar will be given to the highest bidder, but the second-highest bidder must pay the amount of his own bid; that is, high bidder gets the dollar, but the second-highest bidder pays for it. Bids cannot exceed the amount earned in the previous experiment, and no talking is allowed.

The game is not well understood from a game-theoretic perspective. The version with unlimited budgets and unlimited time has no solution except infinite bids. With limitations on endowments, under no circumstances can nonparticipation by everyone be a Nash equilibrium. Models of the situation exist in which a solution involves participation from everyone and in which everyone should be prepared to bid their endowment.[5] The point is that models based on concepts of rationality suggest that rational people might produce intuitively impossible, or perhaps irrational, results (i.e., selling a dollar for much more than a dollar).

The data from five such auctions are in figure 5. The dots are the actual bids in dollars as they occurred in sequence. As can be seen, the dollar always sold for much more than a dollar. In auction 1, for example, the dollar went to a bid of $27, and the price actually paid by the second-highest bidder was $20. Some of the relevant data are in table 2. Participants are indexed according to the size of their budget, with the person with the largest budget called person number 1. In auction 1 the

5. A complete game-theoretic treatment of the auction is not available. Kim Border and Joel Sobel (private correspondence) have produced the following model. The insight of the model is to treat the auction like a sealed-bid auction. The sealed bid is interpreted as a reservation price above which the subject will not go during the actual English auction bidding process. Consider only the two-person case for exposition purposes with the following rules: (i) high bidder receives $1.00 and pays nothing; (ii) second-highest bidder receives zero and pays his bid; (iii) bids must be nonnegative and no more than wealth; and (iv) common knowledge is that wealth is independently and identically distributed from cumulative density function $F(\cdot)$, is supported on $[0, A]$, and has continuous density $f(\cdot)$. Let $V(x) = (1 + x)F(x) - x$; $M(W) = \max\{x \leq W: x \text{ max's } V \text{ on } [0, W]\}$; and $b(W) = W - \{V[M(W)] - V(W)\}/[1 - F(W)]$, using the convention that, if $F(W) = 1$, then $\{V[M(W)] - V(W)\}/[1 - F(W)] = 0$. The bidding function, $b(W)$, is the equilibrium strategy of a symmetric Bayesian Nash equilibrium with risk-neutral players. Generalization to N bidders is straightforward. As an example consider the two-person case in which $F(\cdot)$ is uniform over $[0, A]$. The optimum bidding function is

$$b(W) = \begin{cases} W, & \text{if } W \geq A - 1; \\ \dfrac{W}{A - W}, & \text{if } W < A - 1. \end{cases}$$

In this case the equilibrium strategy is to be prepared to bid all your wealth if your wealth is one less than the maximum possible wealth. Border has also produced an example in which the optimal strategy is for all bidders to always bid all their wealth.

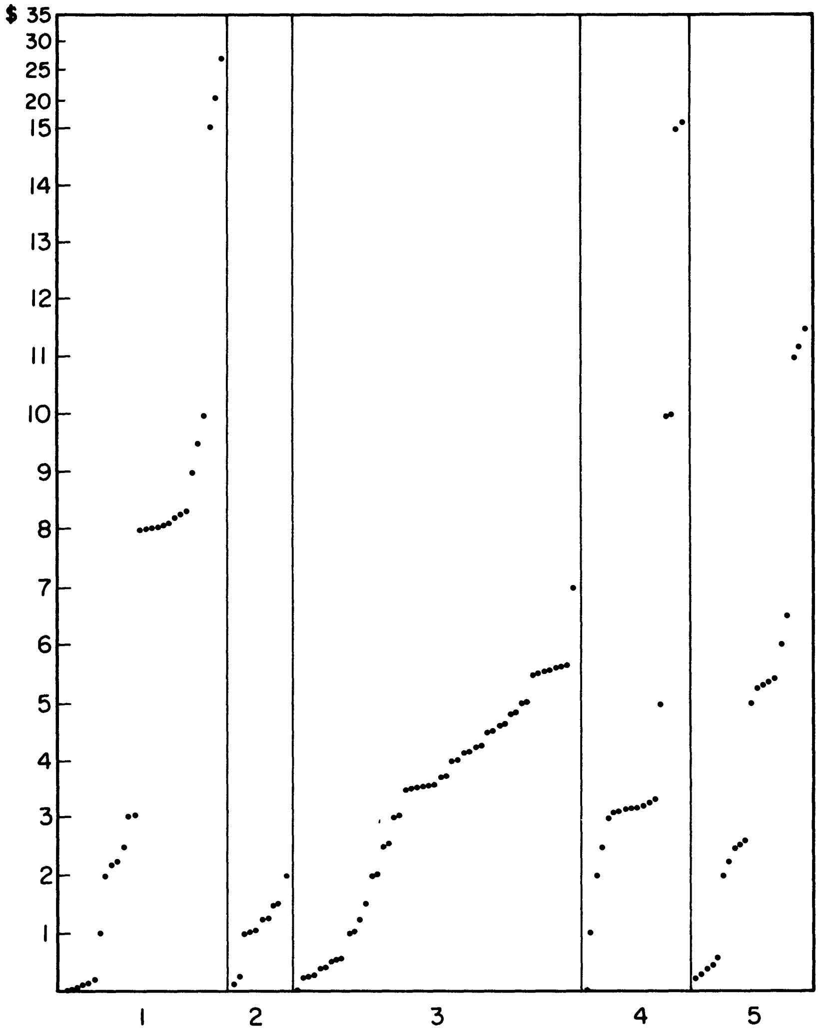

FIG. 5.—Bid time series from five dollar auctions (see table 2)

person with the largest endowment had $40. The auction winner had the sixth-largest endowment at $27.54. The person who paid $20 had an endowment of $20.70. These were the only two bidders after the fourteenth bid of $8.00. Frequently, the auction stopped only after a bidder hit a constraint. The individual who acquired the dollar tended to have an above-average endowment. On the average, people lost a great deal of money.

The phenomenon suggested by the models actually exists. The data contain three interesting lessons. First, models of rational choice help us look beyond the market organizations that have evolved through

TABLE 2 Parameters for Five Dollar Auctions, Bids That Initiate Two-Person Sequences, and the Number of People Bidding

Auction 1: PCC	Auction 2: PCC	Auction 3: PCC	Auction 4: CIT	Auction 5: CGS
Fourteen Participants	Fourteen Participants	Five Participants	Fourteen Participants	Eleven Participants
L_1 = \$40.00	L_1 = \$30.87	L_1 = \$31.68	L_1 = \$41.70	L_1 = \$14.70
$L_W = L_6$ = \$27.56	$L_W = L_{13}$ = \$23.07	$L_W = L_1$ = \$31.68	$L_W = L_5$ = \$21.30	$L_W = L_3$ = \$11.90
$L_L = L_{16}$ = \$20.70	$L_L = L_6$ = \$25.44	$L_L = L_5$ = \$11.28	$L_L = L_{11}$ = \$15.05	$L_L = L_6$ = \$11.20
B_W = \$27.00	B_W = \$2.00	B_W = \$7.00	B_W = \$16.00	B_W = \$11.50
B_L = \$20.00	B_L = \$1.51	B_L = \$5.63	B_L = \$15.00	B_L = \$11.20
Only traders 6 and 16, beginning with B = \$8.00 (fourteenth bid)	Only traders 6 and 13, beginning with B = \$1.02 (fifth bid)	Only traders 1 and 5, beginning with B = \$.52 (ninth bid)	Only traders 5 and 11, beginning with B = \$2.00 (third bid)	Only traders 3 and 6, beginning with B = \$11.20 (nineteenth bid)
Eight people bid at least once	Five people bid at least once	Three people bid at least once	Four people bid at least once	Five people bid at least once

NOTE.—PCC = Pasadena City College, CIT = California Institute of Technology, and CGS = Claremont Graduate School. L_i = limit price of the person with the ith highest limit price, L_W = limit price of the auction winner, B_W = bid of auction winner, and B_L = bid of auction loser—the person who paid.

history to find institutions that might be capable of performing some specific task. One might imagine nobler tasks than to sell a dollar for more than a dollar, but that is not the issue. Second, the existence of intuitively "irrational" market behavior is not conclusive evidence that models based on concepts of individual rationality are inappropriate or ill-equipped to be useful in applications. Finally, the example demonstrates that models based on optimization principles are filled with subtleties often unappreciated by critics of rationality. As far as I am aware, the game described above has never been solved with any degree of generality. We do not know some of the major properties of the Nash equilibrium strategies should they exist. Given the current development of theory, the data cannot be used either to confirm or to reject a theory.

V. Pending Problems

The review above contains no examples of the failure of an economic model to confront the data successfully. I do not want to leave the reader with the impression that such examples do not exist. This section is intended to disabuse anyone of the notion that our models are in perfect shape and that the rationality foundation needs neither examination nor modification. Many problems and paradoxes exist. This paper was not organized around the failure of the models because the reasons for the failures are not clear. Arguments like those in Section III that show differences of procedures and incentives as explanations for unexpected market behavior are very much in contention with arguments that would change entirely the way we think about economics.

The potential problems with rational choice models that have been identified by psychologists and that might be manifest in market behavior have not been systematically explored. This lack of study reflects a resource constraint and not a lack of interest or enthusiasm. Two exceptions to the general rule exist currently, and I understand that more attempts to study markets for evidence of "heuristics" are under way.

In an experimental study by Plott and Wilde (1982) the "representativeness" heuristic (Tversky and Kahneman 1974) was given an opportunity to work. Subjects had valuations for commodity units that were contingent on an underlying state of nature. Prior probabilities were generated by a bingo cage. Once a state of nature was chosen (one for each buyer), a clue to the state was generated by a draw from a second bingo cage. The distribution governing the draws from the second cage was contingent on the state determined by the first draw. After receiving their personal clue, buyers would participate in a market in which the units were being sold. After this process was repeated for several periods (during which the market equilibrated in the usual fashion), the market supply was shifted. The representativeness heuristic predicted no change in volume due to the lowered price. In reality, slight increases occurred as would be anticipated from risk-averse expected utility behavior with Bayesian agents.

A more direct examination of the base-rate fallacy has been conducted by Duh and Sunder (1985). The experiment was similar to the Plott and Wilde experiment, but the supply was completed inelastic, thereby letting price serve as a measure of valuation, and the markets were organized differently. The experiments also varied the base rate to see if the markets responded appropriately. A model based on the principle that base rates would be ignored was rejected in favor of a model based on the principle that people would follow Bayes's law.

The present lack of support in experimental markets for the psychology-based ideas is not going to be the end of the story. Many properties of markets have been observed that are not explicable in terms of current models. Posted prices have an independent effect on market prices (Plott and Smith 1978); nonbinding price ceilings affect market convergence (Isaac and Plott 1981; Smith and Williams 1981); and bubbles can be observed in asset markets (Plott and Sunder 1983; Camerer 1984; Smith, Suchanek, and Williams 1986). The dynamics of the convergence process in equilibrating markets is not theoretically understood at all (Davis and Williams, in press). In fact we have only begun to develop a theory based on individual strategic decisions about why equilibrium is attained in any experimental markets where equilibrium has been observed (Wilson 1982; Friedman 1984; Easley and Ledyard 1986). The Dutch auction behaves differently from the first-price auction even though they are supposed to be behaviorally isomorphic (Coppinger et al. 1980). The signaling experiments discussed above contain events that suggest that some of the markets studied failed to incorporate information that was clearly present in a statistical sense. The markets appeared to adjust appropriately only after a change in experimental procedure drew attention to the statistical regularity. The questions that now exist about the need for economists to consider the decision process used by individuals in addition to observed choices are likely to occur with increasing frequency.

The role of morality, altruism, and ethical predispositions in forming choice is another area of potential discoveries. Needless to say, there has been no way of separating theories of altruistically based behavior and moralistic behavior from preference theory or rational choice theories. Furthermore, since preference theory requires no theory about the source of preferences, no overriding need for a separate theory of moral behavior has been solidly demonstrated. The fact that preferences might include or reflect moral considerations does not, on the surface, contradict a theory of rational choice or maximizing behavior. Moral considerations might influence the shape and form of preferences, but that does not contradict the existence of preferences or choices based on them. One can argue that the existence of morally based behavior provides evidence of rational choice. Experimental markets with externalities, public goods decision processes, and related commons dilemma experiments have not shown the domination of moral considerations over financial motivation (Dawes 1980). Thus no review of procedures and theories has been forced on experimentalists. Nevertheless, evidence of morally based decisions does exist (Palfrey and Rosenthal 1985). In committee experiments the evidence is pronounced especially when as few as three participants are involved (Isaac and Plott 1978; Eavey and Miller 1984). Furthermore, a methodology for investigating experimentally the phenomena and related theories of moral choices is being explored. Hoffman and Spitzer (1985) formulate a strong case that it is possible to formulate in operational terms competing theories about moral attitudes and that it is possible to use experimental techniques to assess their relative accuracy. How our models of rational choice become modified to include the technical features of moral attitudes (consistent? myopic? stable? sensitivity and responsiveness of choice to evidence?), if such exist, remains an open question.

VI. Closing Remarks

The tone of this paper is defensive. Claims about the irrelevance of models of rational choice and the consequent irrelevance of economics are not uncommon topics of conversation. Even economists sometimes engage in disparaging remarks about the discipline because of doubts about either the testability or the validity of the optimization hypothesis. If one looks at experimental markets for evidence, the pessimism is not justified. Market models based on rational choice principles (including the subspecies of satisficing) do a pretty good job of capturing the essence of very complicated phenomena.

On the other hand, the evidence presented here should provide no one with a feeling of overconfidence. Referees who summarize experimental papers by saying, The results are obvious because they follow immediately from rational choice, have not looked very closely at the theory and the data. While the theory of rational choice provides a very useful set of general principles, it is a mistake to elevate the theory to the status of irrefutable law that always reliably operates and need not be challenged.

References

Aizerman, M. A. 1985. New problems in the general choice theory: Review of a research trend. *Social Choice and Welfare* 2 (December): 235–82.

Camerer, Colin. 1984. Speculative price bubbles in asset markets: A theoretical survey and experimental design. Unpublished manuscript. Philadelphia: University of Pennsylvania, Wharton School.

Cohen, Linda; Levine, Michael E.; and Plott, Charles R. 1978. Communication and agenda influence: The chocolate pizza design. In H. Sauermann (ed.), *Coalition Forming Behavior: Contributions to Experimental Economics*. Vol. 8. Tübingen: Mohr (Siebeck).

Coppinger, Vicki M.; Smith, Vernon L.; and Titus, Jon A. 1980. Incentives and behavior in English, Dutch and sealed-bid auctions. *Economic Inquiry* 18 (January): 1–22.

Cox, James C.; Roberson, Bruce; and Smith, Vernon L. 1982. Theory and behavior of single object auctions. In V. L. Smith (ed.), *Research in Experimental Economics*. Vol. 2. Greenwich, Conn.: JAI.

Davis, Douglas D., and Williams, Arlington W. In press. The effects of rent asymmetries in posted offer markets. *Journal of Economic Behavior and Organization*.

Dawes, Robyn M. 1980. Social dilemmas. *Annual Review of Psychology* 31:169–93.

Duh, Rong Ruey, and Sunder, Shyam. 1985. Incentives, learning and processing of information in a market environment: An examination of the base rate fallacy. Department of Accounting Working Paper 1985–5. Minneapolis: University of Minnesota at Minneapolis St. Paul.

Easley, David, and Ledyard, John. 1986. Theories of price formation and exchange in double oral auctions. Social Science Working Paper no. 611. Pasadena: California Institute of Technology.

Eavey, Cheryl, and Miller, Gary J. 1984. Fairness in majority rule games with a core. *American Journal of Political Science* 28 (August): 570–86.

Fiorina, Morris P., and Plott, Charles R. 1978. Committee decisions under majority rule. *American Political Science Review* 72 (June): 575–98.

Friedman, Dan. 1984. On the efficiency of experimental double auction markets. *American Economic Review* 74 (March): 60–72.

Grether, David M. 1981. Financial incentive effects and individual decision making. Social Science Working Paper no. 401. Pasadena: California Institute of Technology.

Grether, David M.; Isaac, R. Mark; and Plott, Charles R. 1981. The allocation of landing rights by unanimity among competitors. *American Economic Review* 71 (May): 166–71.

Grether, David M., and Plott, Charles R. 1979. Economic theory of choice and the preference reversal phenomenon. *American Economic Review* 69 (September): 623–38.

Hoffman, Elizabeth, and Spitzer, Matthew L. 1985. Entitlements, rights, and fairness: Some experimental results. *Journal of Legal Studies* 14 (June): 259–98.

Isaac, R. Mark, and Plott, Charles R. 1978. Cooperative game models of the influence of the closed rule in three person, majority rule committees: Theory and experiments. In P. C. Ordeshook (ed.), *Game Theory and Political Science*. New York: New York University Press.

Isaac, R. Mark, and Plott, Charles R. 1981. Price controls and the behavior of auction markets: An experimental examination. *American Economic Review* 71 (June): 448–59.

Knez, P.; Smith, Vernon L.; and Williams, A. 1985. Individual rationality, market rationality and value estimation. *American Economic Review* 75 (May): 397–402.

Ledyard, John O. In press. The scope of the hypothesis of Bayesian equilibrium. *Journal of Economic Theory*.

Miller, Ross M.; and Plott, Charles R. 1985. Product quality signaling in experimental markets. *Econometrica* 53 (July): 837–71.

Palfrey, Thomas R., and Rosenthal, Howard. 1985. Altruism and participation in social dilemmas. Graduate School of Industrial Administration Working Paper no. 35-84-85. Pittsburgh: Carnegie-Mellon University.

Plott, Charles R. 1979. The application of laboratory experimental methods to public choice. In C. S. Russell (ed.), *Collective Decision Making: Applications from Public Choice Theory*. Baltimore: Johns Hopkins University Press.

Plott, Charles R. 1982. Industrial organization theory and experimental economics. *Journal of Economic Literature* 20 (December): 1485–1527.

Plott, Charles R., and Smith, Vernon L. 1978. An experimental examination of two exchange institutions. *Review of Economic Studies* 45 (February): 133–53.

Plott, Charles R., and Sunder, Shyam. 1983. Rational expectations and the aggregation of diverse information in laboratory security markets. Social Science Working Paper no. 363. Pasadena: California Institute of Technology.

Plott, Charles R., and Uhl, Jonathan T. 1981. Competitive equilibrium with middlemen: An empirical study. *Southern Economic Journal* 47 (April): 1063–71.

Plott, Charles R., and Wilde, Louis L. 1982. Professional diagnosis versus self diagnosis: An experimental examination of some special features of markets with uncertainty. In V. L. Smith (ed.), *Research in Experimental Economics*. Vol. 2. Greenwich, Conn.: JAI.

Rassenti, S.; Smith, Vernon; and Bulfin, R. 1982. A combinatorial auction for airport time slot allocation. *Bell Journal of Economics* 13 (Autumn): 402–17.

Shubik, Martin. 1971. The dollar auction game. *Journal of Conflict Resolution* 15 (Fall): 109–11.

Simon, H. A. 1979. Rational decision making in business organizations. *American Economic Review* 69 (September): 493–514.

Smith, Vernon L. 1976. Experimental economics: Induced value theory. *American Economic Review* 66 (May): 274–79.

Smith, Vernon L.; Suchanek, Gerry L.; and Williams, Arlington W. 1986. Bubbles, crashes and endogenous expectations in experimental asset markets. Working Paper no. 86-2. Tucson: University of Arizona, Department of Economics.

Smith, Vernon L., and Williams, Arlington W. 1981. On nonbinding price controls in a competitive market. *American Economic Review* 71 (June): 467–74.

Spence, M. A. 1974. *Market Signaling: Informational Transfer in Hiring and Related Screening Processes*. Cambridge, Mass.: Harvard University Press.

Tversky, Amos. 1969. Intransitivity of preferences. *Psychological Review* 76 (January): 31–48.

Tversky, A., and Kahneman, D. 1974. Judgment under uncertainty: Heuristics and biases. *Science* 185:1124–31.

Wilson, Robert. 1982. Double auctions. Technical Report 391. Stanford, Calif.: Stanford University, Institute for Mathematical Studies in the Social Sciences, October. Reprinted in T. Groves, R. Radner, and S. Reiter (eds.), *Information Incentives and Economic Mechanisms: Essays in Honor of Leonid Hurwicz*. Minneapolis: University of Minnesota Press (in press).