NOTES ON SOME LITERATURE IN EXPERIMENTAL ECONOMICS

Vernon L. Smith*

*Center for Advanced Study in the Behavioral
Sciences, Stanford, California

Vernon L. Smith*

Center for Advanced Study in the Behavioral Sciences
Stanford, California

1. Introduction

These notes provide a guide and some comments on a portion of the literature of experimental economics. They represent thoughts and criticisms developed in the course of a seminar in experimental economics at Purdue University. The references do not constitute an extensive bibliography. For an excellent summary of the experimental oligopoly literature, the interested reader is urged to consult Friedman [6], and for additional general discussion of the scope and significance of experimental economics and for further bibliographic references, see Naylor [11].

2. Utility and Decision Theory

The earliest contributions to experimental economics were concerned with the construction of empirical commodity-indifference curves and of Neumann-Morgenstern utility curves for individuals. So far as we have been able to determine, the work of Thurstone [21] on indifference curves represents the first attempt to validate an economic proposition by experimental means. Thurstone's study exhibits the usual complement of beginner's errors, but its lack of sophistication should not blur its significance in demonstrating the highly operational character of utility theory. It is tempting, if not already common practice, to teach utility and demand theory by assertion, by an appeal to the "mental experiment," by an appeal to "self-evident" propositions, or by taking refuge in the methodology of positive economics wherein consequences, rather than assumptions, are deemed to be the proper subjects of empirical testing in a theory. As a result students of economics are sometimes poorly motivated to accept, even provisionally, the most elementary and essential concepts of the science. Economic theory is easily considered to be more "unreal" or "abstract" than in fact need be the case. The study by MacCrimmon and Toda [10] provides a modern, experimentally sophisticated treatment of the much-neglected behavioral foundation of demand theory.

By contrast, much attention has been devoted to the experimental study of Neumann-Morgenstern utility theory for choices between uncertain prospects. Indeed, this work is sufficiently rich and sophisticated to pose difficult problems of representative selection, but Dolbear [4] is particularly useful in providing good summaries of the experimental literature, and the principal issues of measurement in utility and decision theory under uncertainty.

The concepts of utility are not only basic to microeconomic theory, they are also basic to experimental methodology in general. It is now well established experimentally that people prefer more money to less, that there are nonmonetary utilities and disutilities in any decision situation, and that subject's decisions can be interpreted in terms of balancing these subjective values. In this sense, it is important in all experimental design to take utility theory seriously. Consequently, one does not ask a subject if he prefers A to B. Rather one designs a choice situation in which it is in the interest of the subject to reveal that he prefers A to B if he does, and otherwise if he does not. Furthermore, choices should be offered and payoffs made so as to control on the effect of wealth and of previous outcomes on preferences. The common technique used by MacCrimmon-Toda and Dolbear, and suggested by Yaari, is to confront the subject with a series of pairwise preference choices with the understanding that when he has completed the task, one of the pairs will be chosen at random and the subject will receive (or engage in it, if it is a gamble) that member of the pair which he says he preferred. If he has a preference ordering and desires to make choices in accordance with those preferences, then he must state truthfully that member of each pair that he prefers. For example, if there are N equally likely pairs

(G_i, H_i), i = 1, 2, $\cdots$, N, and he is an expected utility maximizer, then

$$E(U) = \frac{1}{N} U\{G_1 \text{ or } H_1\} + \frac{1}{N} U\{G_2 \text{ or } H_2\} + \cdots + \frac{1}{N} U\{G_N \text{ or } H_N\}$$

and

$$\max E(U) = \frac{1}{N} \max\{U(G_1), U(H_1)\} + \frac{1}{N} \max\{U(G_2), U(H_2)\}$$
$$+ \cdots + \frac{1}{N} \max\{U(G_N), U(H_N)\}.$$

A subject can only maximize by revealing his preferred choice in each pair. By this device acceptance sets of objects or gambles can be separated from rejection sets to reveal indifference curves, or in the Dolbear study, permit utility indexes to be constructed.

Siegel's paper [14] should be considered as a case study in three closely related pitfalls of experimental economics:

(1) Caution should be exercised in interpreting experimental results as indicative of "irrational" behavior, especially where the reward structure is weak.

(2) Inherent in any decision situation (experimental or not) are certain subjective "costs" or disutilities associated with the process of decision making and reporting as well as the postulated utilities of the decisions. That is, there are "transactions costs" in any task situation which can be expected to influence decisions.

(3) The consequences of decision making in the task situation may have positive or negative commodity value whose effects are confounded with those of any explicit, controlled, reward structure.

In short take utility theory seriously, and interpret it broadly, in trying to understand subject responses.

In this much-studied experimental task, subjects must predict Bernoulli outcomes in a two-choice situation. They have two choices (a_1, a_2) on each trial, and, following each such choice, nature or chance chooses one of two alternatives (θ_1, θ_2). The consequence of (a_i, θ_j) is c_{ij}:

Probability State Act	π θ_1	$1 - \pi$ θ_2
a_1	c_{11}	c_{12}
a_2	c_{21}	c_{22}

The state θ_1 occurs with programmed probability π, θ_2 with probability $1 - \pi$. Typically in these experiments, the value of π is not revealed to the subjects. Let p be the asymptotic expected probability that the subject will choose a_1 (i.e., the proportion of trials in which he predicts θ_1). The Siegel models are concerned with the relation between p and such independent variables as π and the reward structure.

The significance of Siegel's contribution can only be fully appreciated against the background in which some behavioral scientists have interpreted as irrational the failure of subjects to maximize by repeated choice of the more frequent event. Yet experiments with this task, going back to 1939, had typically not reinforced subject responses with monetary rewards or penalties. Utility theory does not predict that people will make the "correct" decision when it is not in their interest to do so. What Siegel shows very clearly is that the alleged evidence for "irrational" behavior is the exception that proves the (utility) rule.

We will derive a somewhat generalized version of Siegel's model I, with a more explicit treatment of the reward structure and the underlying assumptions about utility. Siegel assumes that the subject's choices in the stationary (asymptotic) state can be regarded as consistent with the hypothesis that he maximizes a utility function $u(p)$. However, the subjective consequences of the choices associated with p involve considerably more than the subjective value of any explicit monetary payoffs. Siegel seeks to explain behavior in both payoff and nonpayoff experimental conditions. The consequences c_{ij} of (a_i, θ_j) are assumed to provide three sources of subjective value. There is value associated with just "winning," i.e., predicting correctly, independently of the amount won. Thus there are different utilities associated with consequences c_{11} and c_{22} (i.e., "win") than with c_{12} and c_{21} (i.e., "lose") quite apart from payoffs. Secondly, if in general there is a monetary reward for a correct prediction and a penalty for an incorrect prediction, these monetary consequences

have utility. Finally, the repeated choice of a_i generates utility losses, quite apart from either being correct, or receiving or losing money. This is due to boredom and monotony in not diversifying one's choices. Hence, values of p near 1/2 provide greater subjective value than those near 1 or 0.

Specifically, in the general case, we assume u(p) to be decomposable into three additive components:

(1) If x is the number of successful predictions in n stationary state trials, we associate a utility, $U_1(x)$ with these successes. If we assume that U_1 is proportional to the expected number of correct predictions in n trials, then

$$U_1(x) = aE(x) = anP$$

where a > 0 is a behavioral constant, and P = pπ + (1 - p)(1 - π).

(2) If μ ≥ 0 is the payoff for each successful prediction, and ρ ≥ 0 is the penalty charged for each failure, then the number of units of money received in n trials is y = μx - ρ(n - x), and the Neumann-Morgenstern utility function is

$$U_2(y) = E[U(y)] = E\{U[μx - ρ(n - x)]\}$$

where U(m) is the utility of m units of money. For illustration we postulate a quadratic utility function,

$$U_2(y) = AE(y) + BE(y^2)$$
$$= A[(μ+ρ)nP - ρn] + B(ρ + μ)^2 nP(1 - P)$$

where (A > 0, B < 0) are behavioral constants, and E(y) and E(y²) are computed with the binomial mass function.

(3) Finally, there is a utility $U_3(p)$ associated with variability of choice. For n trials, assume

$$U_3(p) = bnp(1 - p).$$

where b > 0 is a constant and $U_3(p)$ has the property that it is maximized with p = 1/2.

Hence, total utility for n trials is

$$u(p) = U_1(x) + U_2(y) + U_3(p)$$
$$= anP + An[(μ+ρ)P - ρ] + Bn(μ+ρ)^2 P(1 - P) + bnp(1-p).$$

This criterion function generalizes Siegel's model I, and explicitly distinguishes between reward-penalty parameters (μ,ρ) and behavioral parameters (a, b, A, B). Since u''(p) < 0, necessary and sufficient conditions for max u(p) are
0≤p≤1

$$p^o \begin{array}{c} > \\ < \end{array} \frac{b + (2π - 1)[a + A(μ+ρ) + B(μ+ρ)(2π - 1)]}{2b + 2B[(μ+ρ)(2π - 1)]^2}$$

where if > holds $p^o = 0$, if < holds $p^o = 1$, and if = holds we have $0 < p^o < 1$. In the absence of monetary payoffs, μ = ρ = 0,

$$p \begin{array}{c} > \\ < \end{array} \frac{1}{2} + (π - \frac{1}{2})(\frac{a}{b}),$$

and we get probability matching, P = π, if and only if a = b, as in Siegel's model I.

Siegel's model II could be generalized in the above spirit by assuming different utility functions for a correct prediction of the more frequent event and a correct prediction for the less frequent event. Thus, with π > $\frac{1}{2}$, we could write $U_1(x_1) = a_1 E[x_1] = a_1 npπ$, where x_1 is the number of correct predictions of $θ_1$, and $V_1(x_2) = a_2 E[x_2] = a_2 n(1 - p)(1 - π)$ with x_2 the number of correct predictions of $θ_2$. U_1 and V_1 provide richer means of accounting for the "commodity value" choice.

of gambling. Total utility would be $u(p) = U_1(x_1) + V_1(x_2) + U_2(y) + U_3(p)$.[1]

The above models, and those of Siegel, suffer by ignoring subject transient, or dynamic, behavior. Steady-state behavior should be a special case of a more general dynamic model. The natural model of optimal play-by-play behavior would be a Bayesian learning model that modifies subjective probabilities, and predictions, in the light of trial outcomes (θ_1, θ_2). Thus a consequence on a given trial not only yields immediate situation utility, but also information on the parameters of the Bernoulli process which influence future optimal decisions. Such a model has been derived by Emir H. Shurford[15], but it includes neither Siegel's consideration of the utilities of diversification nor gambling (being "correct").

Yaari [22] has proposed a return to the pre-Neumann-Morgenstern utility period by dropping the dominance, or independence, axiom, but adding the assumption of convexity which is, of course, fundamental to the theory of markets. The expected utility theorem which follows if we admit of dominance or independence is mathematically very convenient because it allows us to apply all of the expectation calculus. But does it adequately describe behavior? Yaari suggests it does not, and the experimental results bear him out. Furthermore, the experimental results strongly support convexity.

The importance of Yaari's contribution to experimental methodology is perhaps the emphasis on restricting theoretical constructs to those which are observable. What can be observed in utility theory are those gambles that an individual will reject or accept, i.e., his decisions, rather than probability numbers on which he is presumed to act. Furthermore, in Yaari's approach, it is still possible to interpret subject choices in terms of a concept of subjective probabilities provided that

the choices are consistent with such a concept. But the experimental procedure does not depend upon such a concept.

By way of criticism, Yaari's experimental procedure for constructing the offer curve by either the auction method or the method of letting subjects raise their own bids on an accepted gamble [22, page 285] is not satisfactory. The auction method does not measure the "highest" price a subject is willing to pay for a given gamble but rather the price he is willing to pay for a compound gamble yielding (i) the given gamble if his bid is higher than the competing bids, or (ii) nothing, if his bid is not higher. The auction method introduces the additional uncertainties analyzed in (sealed-bid) auction theory [19], and unnecessarily contaminates the experiment. The method of letting "the subjects themselves raise their own bids by 25 per cent and see if the resultant bid is 'definitely too high'" [22, p. 285] greatly complicates the decision task. Now we have to worry not only about the utility of a gamble but also the utility of giving false information on the value of the gamble. Such problems are likely to be avoided if the experimenter does a simple decision theoretic analysis of the experimental task from the point of view of the subject to determine if it is in the interest of the subject to reveal what the experimenter wants to know. The theory underlying the experiment (that subjects have a preference ordering over gambles) must be assumed also to apply to the subject's choices. If it does not, then there is nothing to measure in the first place. These criticisms should not be allowed to detract from the brilliance and ingenuity of Yaari's paper. They show, rather, how really difficult it is to achieve a clean experimental design.

Even aside from the above criticisms, it is difficult to assess Yaari's experimental technique and results due to incomplete reporting. Ideally, the report of experiments should include the instructions and the results in a form that permits other investigators to reproduce the

experiments and to use or evaluate the data. Yaari raises important issues with insight, and more rigorous tests of convexity would be valuable in increasing our understanding of behavior under uncertainty.

The motivation for the Yaari paper is, in part, to rationalize the Friedman-Savage "paradox," namely the co-existence of gambling and insurance for the same decision maker, but without the Friedman-Savage type of nonconvex utility function. An alternative means of reconciling gambling and insurance within the expected utility hypothesis is to redefine utility to include the commodity value of gambling. Certainly Yaari's proposal redefines the criterion of choice among gambles broadly enough to include any such commodity value. However, Yaari's definition of subjective probability, which corresponds to probabilities under the expected utility hypothesis, suggests an interpretation of behavior in terms of psychological distortions in "true" (or so-called "objective") probabilities. Under this interpretation Yaari's subjects tend to "overstate" low probabilities and "understate" high probabilities. This interpretation is entirely consistent with the data, and one cannot, on these grounds, argue with it as a behavioral hypothesis.

But many economists, such as Samuelson [13; pp. 136, 144], Hirshleifer [8, pp. 257-264], and Smith [20], to name only a few, have found it more natural to explain the discrepancy between actual expected utility-of-wealth behavior and postulated expected utility-of-wealth behavior, in terms of what we can call the commodity value (or cost, if values are negative) of gambling. "Beating the machine," winning on unlikely contingencies, or betting on a draw from an urn with unknown proportions of red and black balls (see Ellsberg [5]), may represent phenomena that yield special utilities or disutilities which are confounded with the subjective value of the monetary prizes won or lost. Thus, Dolbear's experimental procedure [4] carefully attempts to control on the

commodity value of gambling, by always confronting subjects with choices between gambles (as distinct from choices between gambles and cash prizes), while control over discrepancies between subjective and "objective" probabilities is attempted by using gambles with probability 1/2 associated with each payoff.
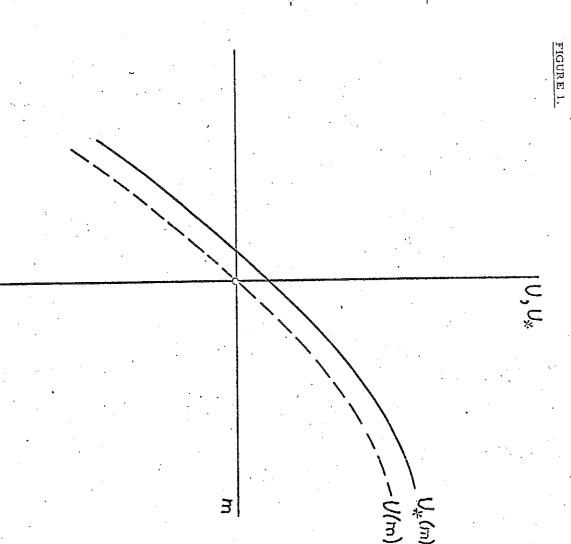
Although many economists have made reference to the importance of admitting of the commodity value of gambling it has not been systematically studied theoretically or empirically. Using a Friedman-Savage type diagram, the consistency of risk aversive behavior (concave utility of wealth) with gambling at mathematically unfavorable odds is easily illustrated. Let U(m) be the utility of wealth, monotone increasing and concave. This function is assumed to be given, quite uncontaminated by the commodity value of gambling. This implies that if we are to ever measure U by choice behavior it must be accomplished in a nongambling context, or by using a psychologically neutral gambling machine or "canonical experiment" (see Pratt, Raiffa and Schlaifer [12]). Now let $U_*(m)$ be the utility of wealth m, when one also engages in a gamble defined by specified unambiguous event contingencies. Think of utility as a function of two commodities, wealth, m, and "gambling." We can either gamble (denoted by 1) or not gamble (denoted by 0). Thus, in terms of a two-commodity utility function, ϕ, we have U(m) = ϕ(m, 0), $U_*(m) = \phi(m, 1)$. In general, we assume that gambling may interact with wealth in determining utility.

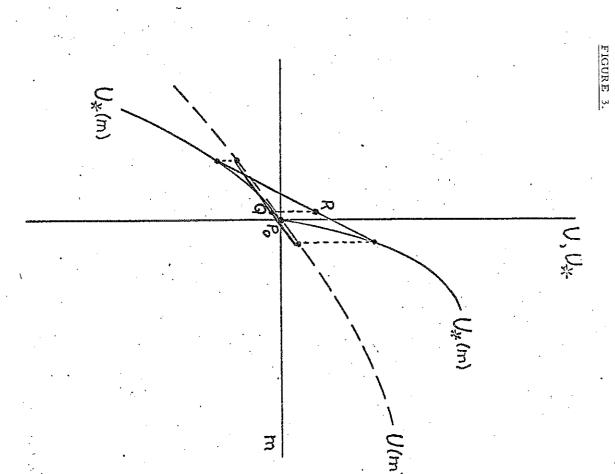Figures 1–4 illustrate various conceivable examples of $U_*(m)$ functions. In Figure 1 gambling yields utility at all wealth levels. In Figure 2 gambling is discommodious if the person loses wealth ("insult" is added to injury), but commodious if he gains wealth. In Figure 3, the person's utility or disutility of gambling is entirely symbolized by the amount won or lost so that a gambling game played without real money

yields no subjective value. In Figure 1 and, possibly, Figure 2, the individual would pay to gamble just to see if he can "win," even though there are no monetary payoffs (e. g., playing poker just for valueless chips). In Figure 4, gambling is discommodious for all wealth, but not unless real money is at stake.

In Figure 3, we illustrate the effect of the utility of gambling on an individual's acceptance of an unfavorable bet. If gambling had no commodity value, then $U(m)$ would be the appropriate utility function, and a mathematically unfavorable gamble with expected utility such as at Q would be rejected. But if $U_*(m)$ is the appropriate utility-with-gambling function, then this gamble would have expected utility given by R which exceeds the utility of not gambling at P_o.

Now, $U_*(m)$, viewed in terms of the commodity wealth, does indeed have a nonconvexity in a region near the origin, while interquadrant convexity is borne out by the Yaari experimental results. However, in order to pick up gambling utilities or disutilities included in $U_*(m)$, considerably more experimental control may be necessary over the conditions of subject entry. In taking the origin arbitrarily as a point on the offer curve, Yaari not only assumes zero cost of entry for all gambles but, more important to the issue of the acceptance of mathematically unfair gambles, he assumes zero <u>value</u> for entry. Paying subjects for participating in the experiment, as a means of reducing the "cost of entry," may have the effect of destroying the commodity value of gambling (as in Figure 3) by permitting it to occur with free or "house" money. Also if nonconvexities near the origin are to be identified, observations must be made in this region and it is not clear from Yaari's report on the test of interquadrant convexity that the experimental design allowed this.

FIGURE 1.

FIGURE 2.

FIGURE 3.

FIGURE 4.

FIGURE 4.

U(m)

U*(m)

m

3. Competitive Auction Markets

The first experiments in competitive auction markets reported in [16] were conducted against the background of considerable skepticism by Smith concerning the tentative hypothesis that supply and demand theory might have some relevance to the observed performance of such simulated markets. This bears testimony to the power and influence of the revolt against competition, contained largely in the concepts associated with the theory of monopolistic competition. Econometric studies were concerned with measurement, or testing hypotheses about coefficients, always under the matinained hypothesis that price data were somehow generated by supply and demand equilibria. Direct evidence was lacking, and, as Guy Orcutt once noted, the econometrician is in a position like that of an electrical engineer who must infer the laws of electicity by listening to a radio play.

References [16], [17], [18] and [7] constitute all the direct evidence that seems to exist as to the empirical relevance of supply and demand theory to competitive auctions. The research in these papers demonstrates the equilibrating power of the "public auction" mechanism, and the possibilities of experimental methods in exploring the elementary foundations of economics. Economists need not be content with "listening to a radio play."
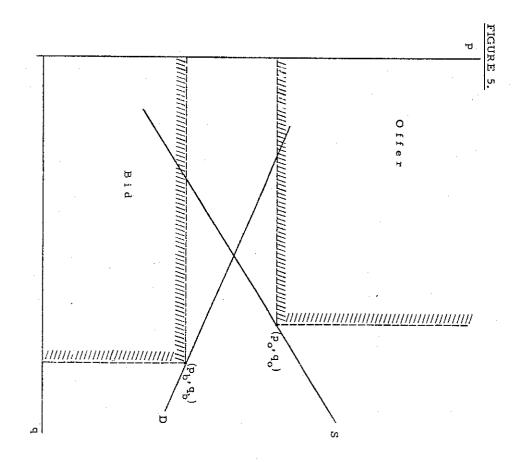
One of the important lessons of experimental economics is that the discipline of the laboratory is very demanding of economic theory. (We have already noted that experimental studies of utility cannot ignore "transaction" costs.) Experimental competitive market design requires the specification of rules of exchange which circumscribe some process whereby decision makers interact. The demand on economic theory is to provide models of the trading process which can be confronted with

experimental evidence. Traditional economic theory has not provided such a process. One macromarket adjustment hypothesis--the Walrasian hypothesis--has dominated supply and demand theory. Yet this hypothesis is essentially a logical argument as to why nonequilibrium states cannot persist and cannot be considered as a serious model of a trading process. Thus, so it is argued, if "price" is above equilibrium there is excess supply, and sellers will "competitively" lower prices in an attempt to sell the excess. Yet competitive market theory requires each participant to take price as given. Such a theory omits explicit consideration of a multilateral negotiation--bid, offer, acceptance-- process. There is no theoretical distinction between price quotations and actual transaction prices. Yet, all actual markets, including the over-the-counter securities markets as well as the organized commodity and stock exchanges, operate under formal or informal rules which govern the placing of bids and offers, and their acceptance to form contracts.

What is made plain by these experiments, and is the source of their greatest shortcoming, is the need for models of such a bid, offer, transaction process. In actual markets, as in the experimental markets, a bid is always a 2-tuple (p_b, q_b) specifying the maximum buying price and the maximum quantity legally acceptable to the buyer. That is, it is understood by the quotation that the bidder will accept at any lower price any quantity up to q_b. This interpretation, besides being legal practice, is formalized in the trading rules of the New York Stock Exchange. Thus, a "limit" bid (p_b, q_b) means that p_b is the upper limit of the bid price commitment by the buyer. That is, it is understood by the quotation that the bidder will accept at any lower price any quantity up to q_b. This interpretation, besides being legal practice, is formalized in the trading rules of the New York Stock Exchange. Thus, a "limit" bid (p_b, q_b) means that p_b is the upper limit of the bid price commitment by the buyer. Also the rules of the Exchange prohibit "all or none" bids (and offers) so that if $q_b = 500$, any part of the 500 may be taken at a price no greater than p_b.

Similar considerations govern an offer, which is a 2-tuple (p_o, q_o) specifying the minimum buying price and the maximum quantity acceptable

to the seller at that price.[2] The convex set of acceptable transactions represented by a bid (p_b, q_b) and by an offer (p_o, q_o) are each illustrated in Figure 5 for individual supply (S) and demand (D) curves. Figure 6 uses the Edgeworth box offer curve(O_1 and O_2) representation to illustrate an offer (p_I^o, q_{1I}^o) of commodity # 1 by trader I, and an offer $(1/p_{II}^o, q_{2II}^o)$ of commodity # 2 by trader II.

FIGURE 5.

FIGURE 6.

P

q_{2II}^o

O_2

$(1/p_{II}, q_{2II}^o)$

$(1/p \geq 1/p_{II}, q_{2II} \leq q_{2II}^o)$

O_1

(p_I^o, q_{II}^o)

q_{II}^o

$(p \leq p_I^o, q_{II}^o \leq q_{II}^o)$

I

II

O

Referring to Figure 6, if the term "price" in the Walrasian tâton-nement process is interpreted as an offer, it is clear that it cannot be open-ended as to quantity. Thus if I, without knowledge of the offer curve O_2, makes the unlimited price offer OP shown in Figure 6, he may be exposing himself to the risk of an acceptance that exceeds his resource capacity, as for example at P.

On the other hand, the limit offer (p_I^o, q_{II}^o) assures I that he cannot be made worse off by the acceptance of an element in the offer set, and precludes a contract beyond his offer curve, O_1. Under the essential assumption of incomplete knowledge, there is no guarantee of contracts at the supply-demand equilibrium. However, the simple experiments in [16], [17] and [18] show strong convergence tendencies where the market is repeated under the same conditions of supply and demand.

Using the above concept of an offer (or bid) many models of price adjustment processes might be developed. One type of model would be to postulate a subjective probability density, for a given trader, which as-sociates a probability of acceptance with each possible offer. An offer by a trader could then be defined as one which maximized expected utility over this subjective density. Some appropriate assumption could then be introduced to modify (e.g., by Bayes' theorem) this density in response to information, i.e., offers by other traders and whether they were accepted. An acceptance by a trader would occur, if a given out-standing offer provided at least as good terms as that trader's expected utility-maximizing offer. Such a process provides a mechanism for generating both offers and transactions.

Carlson's [1] study of markets with a lagged supply response provides direct evidence on the relationship between past prices and expected future prices. The simple cobweb hypothesis, which has commanded so much attention in the literature, is not confirmed. Subjects are not so naive as to believe that next period's price will be

the same as last period's price. Two experimental supply and demand configurations, one divergent and one convergent, under the cobweb hypothesis, reveal stable and similar convergent behavior. Indeed, the results suggest that in a potentially divergent market, the greater are the initial price fluctuations, the less confidence a supplier will have in the last price as an indicator of the next price. In a subsequent paper Carlson [2] shows that if price expectations are based on an unweighted average of all past prices, if the excess demand function is decreasing, and if the supply and demand functions are linear, then an isolated market always converges to the equilibrium price.

Studies [1], [7], [16], [17] and [18] represent a small portion of the potential return from the application of experimental techniques to competitive market behavior. The problem of dynamic market adjustments over time to changing demand or supply conditions could be studied with a design which permits subjects to buy for inventory in anticipation of more favorable future sales.

General equilibrium models of exchange could be studied experimentally. For example, in the two-commodity case the objects traded might be red and blue poker chips. Utility value would be induced on holdings of the two objects by means of a table associating a value $V(q_r, q_b)$ in U.S. currency with each possible terminal holding (q_r, q_b) of red and blue chips. The table provided to each subject would then serve as his ordinal utility of red and blue chips provided only that the subject exhibit a monotone increasing utility function for U.S. currency, $U(V)$. Thus $U[V(q_r, q_b)]$ would induce a controlled subjective value on q_r, q_b, independent of each subject's actual utility for currency. That is, the marginal rate of substitution of red for blue poker chips,

$$\frac{\partial q_r}{\partial q_b} = - \frac{U'(\partial V/\partial q_b)}{U'(\partial V/\partial q_r)} \ ,$$ will depend only on $V(q_r, q_b)$ which is

predetermined by the table. Trading could be a matter of direct exchange of red chips against blue, or through a monetary medium with red chips first traded against "stage money," then "stage money" against blue chips to complete a full round of exchange. The postulate that demand is homogeneous in nominal prices and ("stage") money income could be tested, as well as the equivalence of monetary and direct exchange systems.

Production and a producers' market could be added by introducing production function tables and trading in claims on labor input endowments (for example, white chips). But note that in such a general equilibrium model one would not have to introduce profit payoff tables for producer subjects, as in partial equilibrium oligopoly experiments. The rewards of producers would be derived from their "production" of red and blue chips and sale to "consumers." The $U[V(q_r, q_b)]$ functions of "consumer" subjects would be the entire driving force of the economy, inducing value, through production, upon artificial labor input endowments.

All experimental economics depends upon the maintained hypothesis that the utility of money, $U(V)$, is monotone increasing. Different experiments differ only in the decision arguments, assigned by the experimental task, to the V function that one gives to the subject.

Footnotes

1. For a discussion and comparison of steady-state models of
   Bernoulli process predictions see R. Duncan Luce [9].

2. Special cases of the limit bid or offer are represented by so-
   called "market" orders which do not specify price. In these
   cases it is understood that $P_b = \infty$, or $P_o = 0$, i.e., there is
   no legal maximum price enforceable by a buyer, or legal
   minimum enforceable by a seller.

References

[1] Carlson, J.A., "The Stability of an Experimental Market with
    a Supply-Response Lag," Southern Economic Journal, Vol. 33,
    1967, 305-321.

[2] Carlson, J., "An Invariably Stable Cobweb Model," Review of
    Economic Studies, Vol. 35 (3), No. 103, July 1968, 360-362.

[3] Carlson, J.A. and T.B. O'Keefe, "Buffer Stock and Reaction Co-
    efficients: An Experiment with Decision Making Under Risk," Review
    of Econ. Studies, Vol. 36 (4), No. 108, October 1969, 467-484.

[4] Dolbear, F.T., "Individual Choice Under Uncertainty: An Experi-
    mental Study," Yale Economic Essays, Vol. 3, 1963, 419-470.

[5] Ellsberg, D., "Risk, Ambiguity, and the Savage Axioms,"
    Quarterly Journal of Economics, Vol. LXXV, No. 4, November
    1961, 643-669.

[6] Friedman, J.W., "On Experimental Research in Oligopoly," Review
    of Econ. Studies, Vol. 36 (4), No. 108, October 1969, 399-415.

[7] Hess, Alan C., "Experimental Evidence on Price Formation in
    Competitive Markets," Journal of Political Economy, Vol. 80,
    March/April, No. 2, 1972, 375-385.

[8] Hirshleifer, J., "Investment Decision Under Uncertainty: Appli-
    cations of the State-Preference Approach," Quarterly Journal of
    Economics, Vol. LXXX, No. 2, May 1966, 252-277.

[9] Luce, R. Duncan, "Learned Versus Optimizing Behavior in Simple Situations," Chapter 7 in M. W. Shelly, II and G. L. Bryan (editors), Human Judgements and Optimality, John Wiley and Sons (New York, 1964).

[10] MacCrimmon, K. R. and M. Toda, "The Experimental Determination of Indifference Curves," Review of Economic Studies, Vol. 36 (4), No. 108, October 1969, 433-451.

[11] Naylor, T. H., "Experimental Economics Revisited," Journal of Political Economy, Vol. 80, No. 2, March/April 1972, 56-84.

[12] Pratt, J., H. Raiffa, and R. Schlaifer, Introduction to Statistical Decision Theory, McGraw-Hill (New York, 1965), Chapters 2 and 3.

[13] Samuelson, P., The Collected Scientific Papers of Paul A. Samuelson, edited by J. Stiglitz, M.I.T. Press (Cambridge, 1966), Vol. 1, Chapters 13 and 14.

[14] Siegel, Sydney, "Decision Making and Learning Under Varying Conditions of Reinforcement," Annals of the New York Academy of Sciences, 1961, 89, 766-783.

[15] Shurford, E., "Some Bayesian Learning Processes," Chapter 9 in Human Judgements and Optimality, op. cit.

[16] Smith, V. L., "An Experimental Study of Competitive Market Behavior," Journal of Political Economy, April 1962, 111-137.

[17] Smith, V. L., "Effect of Market Organization on Competitive Equilibrium," Quarterly Journal of Economics, May 1964, 181-201.

[18] Smith, V. L., "Experimental Auction Markets and the Walrasian Hypothesis," Journal of Political Economy, Vol. 73, No. 4, August 1965, 387-393.

[19] Smith, V. L., "Experimental Studies of Discrimination Versus Competition in Sealed Bid Auction Markets," Journal of Business, Vol. 40, No. 1, January 1967.

[20] Smith, V. L., "Measuring Nonmonetary Utilities in Uncertain Choices: The Ellsberg Urn," Quarterly Journal of Economics, May 1969, 324-329.

[21] Thurstone, L., "The Indifference Function," Journal of Social Psychology, Vol. 2, 1931, 139-167.

[22] Yaari, M. E., "Convexity in the Theory of Choice Under Risk," Quarterly Journal of Economics, Vol. 79, May 1965, 278-290.