C-2

Center for Research on Economic and Social Theory
Research Seminar in Quantitative Economics

Discussion Paper

[redacted]

DEPARTMENT OF ECONOMICS

University of Michigan

Ann Arbor, Michigan 48109

[redacted]

THE TOKEN ECONOMY REINFORCEMENT THEORY AND THE CONSUMER MODEL

John G. Cross*

C-2

February 1978

*Professor of Economics, The University of Michigan

THE TOKEN ECONOMY

REINFORCEMENT THEORY AND THE CONSUMER MODEL*

John G. Cross

University of Michigan
Ann Arbor, Michigan 48109

It seems probable that the contemporary microeconomic model of the consumer could be adapted to incorporate a more realistic description of ordinary human behavior. For many years economists have tolerated the amalgam of good sense and bad psychology which is their microeconomic theory because of the empirical plausibility of its predictions. There is growing pressure to expand the limits of this theory, however, and by now it has become common for empirical regularities to be rationalized in terms of variables which never have had a secure place in the theory, or which even may be in open conflict with it. This is exemplified by Houthakker's and Taylor's (1970) use of a consumer's "stock of habits" as a variable in their demand equations, and by Katona's (1960) appeal to low levels of "consumer confidence" in his rationalization of the observation that consumer saving rates tend to increase in times of unusual inflation. Of course, it is possible that these are merely cases in which economists have taken unconventional lines on situations in which orthodox models still might be made to work were one to introduce such factors as information costs and lags, expectations formation, transactions costs and the like. Nevertheless, it is intriguing to find

---

*The author would like to thank John Fitts, Robert S. Holbrook and Sidney G. Winter for their comments upon earlier drafts of this paper.

that behavior which does not correspond to the simplest traditional models is frequently explained, not with these modifications, but instead with propositions of a psychological or quasi-psychological character.

It seems worthwhile to investigate whether established psychological insights can be introduced into existing consumer theory in order to reduce our apparant dependence upon ad hoc hypothesizing. In this paper, the role of the learning process is addressed specifically, and a model is developed which incorporates a simple reinforcement view of behavior determination into the traditional consumption framework. The model is applied to three problem areas: (1) intertemporal choice and the role of time preference in consumer saving decisions, (2) the pervasive use of lagged variables in consumer demand studies, and (3) the failure of economists to introduce demonstration effects and the obviously significant influence of advertising explicitly into the model of the household.

The Token Economy

The work of James Duesenberry (1952) has already brought the relevance of learning processes to the attention of economists concerned with the nature of the consumption function, although his work antedates the development of most of the quantitative models of learning which now exist. An additional impetus for this research comes from compelling demonstrations of the practical usefulness of learning theory which have been provided by psychologists working in the area of simple reinforcement learning (i.e., "reward"-induced learning), an aspect of behavior which already has been subjected to extensive experimental study.[1]

It has long been recognized by psychologist and layman alike that particular actions or behavior patterns become "habitual" if they are regularly accompanied or closely followed by some kind of tangible payoff. Conversely, actions come to be avoided if they are regularly followed by some sort of "punishment". Beginning in 1961, Theodoro Ayllon and Nathan Azrin (1966), psychologists working in an Illinois State mental hospital, devised a sophisticated method for making practical use of this simple principle. Their procedure was to "reward" desired behavior on the part of patients with "tokens" which at a later time could be exchanged at specified rates for "reinforcers" (rewards) of the patients' own choosing.[2] This system was seen to have two virtues: First, even if the payoffs were necessarily deferred until later, the tokens provided an immediate and tangible reward which, by proxy, could have an equivalent effect. Second, the patients were able to exchange the tokens for any of a wide variety of goods, thus choosing for themselves the most powerful "reinforcer" which was available.

The token device has since generated enormous interest among psychologists, and it has been applied to a broad spectrum of problems in behavior change, ranging from classroom situations in which children are awarded tokens for good performance (especially in mathematics) to prisons in which inmates are rewarded for "socially acceptable" behavior.[3]

One of the remarkable features of these experiments is that despite the obvious analogy to ordinary market economies, the effectiveness of the tokens is rationalized by their users entirely in terms of psychological learning theory; there is no acknowledgment of the consumer as a rational, maximizing decision-maker, although the results of token use correspond directly to the predictions which can be made on the basis of economic theory. For example, in carefully controlled sets of circumstances, Ayllon and Azrin were able to obtain downward-sloping demand curves, upward-sloping supply curves (including the suggestion of a backward-bending supply curve), supply-demand equilibrium prices, and even the purchase of franchises (the payment of tokens in exchange for the guarantee of a particular job).[4]

To an extent, such results may be interpreted as empirical support for the positivistic proposition that individuals behave "as if" they were utility maximizers. As it happens, however, learning experiments turn up a few phenomena which are quite difficult to reconcile with the maximization model. One such example occurs in Ayllon's and Azrin's work on the impact of "advertising" on consumption decisions.[5] The experimental procedure was to provide "free samples": small quantities of those goods or services which were to be "advertised." For a wide variety of items, the provision of these "free samples" was found to increase purchases (expecially on the part of those who had not purchased the commodity before). Such a result, of course, is consistent with an optimization model with imperfect information; the free sample may have showed consumers that they "liked" the commodity. However, when the "free sample" programs were terminated, purchases made by the same individuals dropped abruptly--not to levels as low as before the advertising was instituted, but lower than during the advertising program.[6] Although results of this kind are readily explained in psychological terms, they do not seem to be consistent with microeconomic consumer theory as it presently stands, even allowing for the amendments which in other circumstances seem to be so useful. This example will be interpreted in terms of reinforcement theory in a later section of this paper.

Intertemporal Consumption Decisions and Saving Behavior

Before proceeding with any formal model, we must recognize that with respect to one aspect of individual behavior, the conflict between economists' and psychologists' views appears to be irreconcilable. Most economists seem to subscribe to the general intertemporal decision model in which consumers are presumed to formulate income and consumption plans, in some cases extending over entire lifetimes, which maximize the utility, properly discounted, of all future consumption. As many economists realize, however, this view of consumption as the outcome of forward-looking planning is incompatible with current beliefs about the psychological determinants of behavior. Learning experiments have demonstrated repeatedly that unless some means (such as Ayllon and Azrin's "tokens") are available for bridging time gaps, even moderate lags between an action and its "reward" virtually eliminate the reinforcing potential of that reward.[7] Thus, from the point of view of learning theory, the supposition implicit in time preference models that the present action of saving is encouraged merely by the anticipation of consumption several years in the future is simply insupportable.

Learning models are essentially backward-looking, reflecting the proposition that behavior is determined by past experiences rather than future possibilities. This is not to say that these experiences must apply only to oneself. Within the learning context there still exists the possibility of observing and reacting to the experiences of others. The young person choosing a professional career will certainly be influenced (reinforced) by the apparent wellbeing of those who are already in that career, and the savings behavior of the thirty-year-old consumer will reflect impressions gained from observation of those already in retirement. Nevertheless, the behavior predicted by this theory is significantly different from that which would be produced by life-cycle planning. For example, cobweb cycles are much more likely to occur if markets are composed of learning consumers because individuals will draw their behavioral signals from the wrong generations. Rather than using currently available data to forecast the likely consequences of their own actions, they will react to the wellbeing of others whose choices may have been made under quite different circumstances. This will produce consumption and investment lags even in markets for which the data necessary for rational planning are readily available.

To formulate a model of saving behavior which is consistent with the perspective of learning, we will postulate that the reinforcer in the consumption-deferral process is the accumulation of savings itself rather that the uses to which the savings ultimately are put, or the dates at which those uses come up. Thus if the consumer suffers (or observes someone else suffering) the consequences of insufficient resources in the face of sudden financial necessity, the response will be not to plan more carefully or thoroughly for that particular contingency, but simply to save more thereafter.[8]

The sacrifice which this modification entails is one of theoretical elegance rather than of empirical substance, since the qualitative implications of the theory are very little affected. So long as the consumer's valuation of savings is to some extent influenced by the interest income they earn, a traditional comparative static model with "savings" and "current consumption" has implications similar to those produced from a model with "current consumption" and "future consumption" as arguments.

Although the theory of investment is harder to retain in these terms, we will argue in a later section that the same mechanism which makes advertising effective also implies that the demonstration of return which investment has made possible for others will itself tend to reinforce that investment in oneself, without one's having to make any explicit calculations of the return to one's own investment *per se*.[9]

A Learning Model

In a previous paper (Cross, 1973), I have outlined the possibility of constructing a model of the firm from the perspective of learning theory, using a modification of the specific learning model of Bush and Mosteller (1955). We may apply a similar learning paradigm to consumers if we adopt the following definitions and notation:

X_i refers to an "action" on the part of the consumer; it is a vector whose components, $x_{i1},\ldots, x_{is}$, describe quantities consumed of each of s different commodities (including saving). The set of feasible commodity bundles is confined to the consumer's budget: if I represents the consumer's income at the beginning of a budget period, and $Q = q_1,\ldots, q_s$ represents the vector of commodity prices, then X_i must satisfy the condition $Q \cdot X_i = I$. Although dissaving is certainly possible, we presume that there are limits to lenders' willingness to support the overcommitment of current income. This, together with the natural non-negativity of real commodities, places a lower bound on each of the X_{ij}s. The number of feasible (integral) commodity bundles is therefore finite for all non-zero prices, and, for convenience, we shall denumerate these with the index $i = 1,\ldots, n$. Naturally n is much larger than s.

It is a property of virtually all psychological learning theory that choices are considered not to be determinate but to be governed by probability distributions. Experimental studies of behavior reinforce this attitude very strongly, suggesting that even under carefully controlled circumstances, one cannot predict a particular behavior pattern at a particular time with certainty. This may of course be only a reflection of our limited understanding of the mechanisms of behavior, but it seems to be safe to infer from the empirical evidence which exists already that ordinary economic variables such as price and income are not adequate for the determination of choice behavior, and that the psychologists' distribution functions are the best that we presently can do. It is these probability distributions which are considered to be subject to the influence of learning. Since we have chosen to identify a finite set of alternative consumption bundles, we shall describe this distribution with a vector of probabilities $P_{i,t}$ for $i = 1, \ldots, n$, where $P_{i,t}$ refers to the probability that the bundle X_i will be the one chosen at time t.[10]

We will use a variable U_i to refer to the payoff which the consumer receives from the bundle X_i. From a psychological standpoint, the definition and units of U_i are quite controversial; indeed, if economists were to follow up on their frequently expressed intention to "go to the psychologists" to find the appropriate characterizations for their utility functions, they would come away severely disappointed. Nevertheless, for the purposes of this paper, we must avoid this very interesting detour and treat U_i as though it were obtained from a familiar neoclassical cardinal utility index: thus U_i may be interpreted as the "utility" of the consumption bundle acquired by the action X_i. In keeping with the usual properties of such a function, all the U_is will be considered to be positive. The problem of satiation or the possibility of unfortunate consumption experiences (such as the acquisition of a defective automobile) can be introduced into a learning model, but the benefits of doing so are slight. Moreover, there is to be no notion of "opportunity cost" in this model. Our consumers react to how well they do, but notions as to how well they could have done if they had chosen differently, if introduced at all, belong in the category of advertising and related external influences over behavior.

According to the principles of learning, if a choice X_i is made at time t and is positively reinforced ($U_i > 0$) then the likelihood that that choice will be made at time t + 1 is increased. Necessarily, the likelihoods of alternative choices are reduced correspondingly. Using our notation, $P_{i,t+1} > P_{i,t}$ and $\sum_{k \neq i} P_{k,t+1} < \sum_{k \neq i} P_{k,t}$ whenever X_i is chosen and U_i is positive. Naturally, we also require $0 \leq P_{k,t} \leq 1$ and $\sum_{k=i}^{n} P_{k,t} = 1$ at times t and t+1. There are many models of learning which reflect these basic features and only because it is the easiest to use, we will here adopt the linear version proposed by Bush and Mosteller. If X_i is chosen at time t, and U_i is positive, then we will write:

$$P_{i,t+1} = P_{i,t} + \alpha(U_i)(1-P_{i,t}) \tag{1}$$

$$P_{k,t+1} = P_{k,t}\,[1-\alpha(U_i)] \text{ for } k \neq i \tag{2}$$

Equations (1) and (2) are a direct statement of the Bush-Mosteller theory except for the insertion of the function $\alpha(U_i)$. In their original model, when any action is positively reinforced, the likelihood of its occurrence upon a subsequent occasion is increased by an amount which is determined by a constant, independent of the nature of the reward.

Equations (1) and (2) replace that constant by the function $\alpha(U_i)$ which can take into account the differential impact of differential payoff (utility) magnitudes. It is postulated that larger payoffs have a greater impact of learning; hence $\alpha'(U_i) > 0$. Since by the nature of the probability calculus, we must have $0 < \alpha(U_i) < 1$, then for increasing values of U_i we must eventually have $\alpha''(U_i) < 0$. This is a sort of diminishing marginal effectiveness of payoffs on learning.

The actual value of $P_{i,t+1}$ depends upon the particular choice which is made in time t. If that choice is X_i, then $P_{i,t+1}$ is determined by equation (1). If the choice is some X_k where $k \neq i$, then $P_{i,t+1}$ is determined by an equation analogous to (2). The choices at time t are governed by the vector $P_{1,t},\ldots,P_{n,t}$. Thus we can determine an expected value for $P_{i,t+1}$ from equations (1) and (2) using $P_{i,t},\ldots,P_{n,t}$ as weights. After some manipulation, this produces equation (3):

$$E[P_{i,t+1}] = P_{i,t}[1+\alpha(U_i) - \sum_{k=1}^{n} P_{k,t}\alpha(U_k)] \quad i = 1,\ldots,n \qquad (3)$$

Suppose that X_i happens to be a unique utility maximizing bundle ($U_i > U_k$ for all $k \neq i$). Then we must have $\alpha(U_i) > \alpha(U_k)$ for $k \neq i$, and whenever $P_{i,t} < 1$, it must be the case that $\alpha(U_i) > \sum_{k=1}^{n} P_{k,t}\alpha(U_k)$. From equation (3), this implies that $E[P_{i,t+1}] > P_{i,t}$ for every value of $P_{i,t}$ (except for the extreme case of $P_{i,t} = 1$, which then gives $P_{i,t+1} = 1$ as well). It follows that $E[P_{i,t+2}] > E[P_{i,t+1}]$, because for every element in the weighted average $E[P_{i,t+1}]$, the corresponding expectation for the period t + 2 exceeds the value of that element. Moreover, the strict inequality will apply for every value of $E[P_{i,t+r}]$, implying that $\lim_{r\to\infty} E[P_{i,t+r}] = 1$. That is, we expect the likelihood of the utility-maximizing choice X_i to approach unity, and we may conclude from this that the long run behavior which is predicted by this learning theory is not different from the utility-maximizing behavior which is specified by the traditional model.

Short Run Price and Income Variations

We have shown that the introduction of learning need not affect the standard comparative-static model of the consumer: The long run equilibrium theory implies the same behavior in either model. By its very nature, however, the comparative-static theory never has provided a viable short run model. Economists concerned with short run consumer demand and the short run consumption function have been forced either to apply the comparative static model directly, thereby implying instantaneous adjustment on the part of consumers, or to invent some *ad hoc* dynamic adjustment process ("information lags", "uncertainty," etc.) which will produce more plausible implications but often at the expense of appearing to have been introduced on the spur of the moment solely for the purpose of plugging an embarrassing gap in the model. A useful property of any learning theory is that it begins with a short run dynamic process and derives from that the properties of long run equilibrium without any need for additional assumptions.

According to the learning theory, moreover, short run adjustment cannot be inferred from the properties of equilibrium, because in the general case of a dynamic, changing economy, the consumer is usually out of equilibrium, continually reacting and adjusting to new conditions and to past successes and mistakes. In addition, in a short run adjustment model, unlike an equilibrium model, it is not possible to use the budget in the usual way. Since purchases during the budget period are subject to a stochastic influence, the extent to which the consumer's income is under- or over-committed is also determined by stochastic processes. Necessarily, then, some consumption items must play the role of residuals in the budget, decreasing or increasing inversely with earlier expenditures. Cash accumulations and other forms of saving are natural candidates for this role, and it seems proper to include many durable goods purchases in this category as well. If food and other current "necessities" have absorbed all of a household's income, the tendency may be to tolerate or repair old durable equipment (at least for this period) until enough money is left over for a larger expenditure upon replacement.[11] This is not to say that saving or durable goods purchases play a passive role in consumer spending decisions; they are part of the overall consumption bundle, and they contribute to the reinforcing effect of that bundle. Alternative buffers are of course possible, for virtually any consumption good could fill this role so long as the household can survive through a budget period without it.

We separate the commodity bundle, X_i, into two components: the vector of "ordinary" commodities

$$Y_i \equiv x_{i1},\ldots,x_{i\ s-1}$$

and a single "buffer" commodity, x_s. This procedure preserves the meaning of of the index i with respect to the consumption of ordinary commodities when there is a price or income change. Moreover, if we take saving to be represented by x_s, then, since the price of saving equals 1, we have $x_s = I - Q\cdot Y_i$ where Q is now the vector of prices $q_1,\ldots,q_{s-1}$. The probability $P_{i,t}$ now refers to the bundle Y_i.

The utility derived from the choice Y_i is $U_i = f(Y_i, x_s)$ where we may consider the form of the function f to incorporate any possible influence of the consumer's wealth holdings.

We do not know with certainty the particular Y_i which the consumer will choose at time t; however, we do have a vector associating a given probability with each alternative. Therefore, an expected consumption bundle at time t may be defined as the vector of expected goods purchases:

$\bar{x}_{1,t}, \ldots, \bar{x}_{s-1,t}$ where each of these elements is given by

$$\bar{x}_{j,t} = \sum_{i=1}^{n} P_{i,t} x_{ij} \qquad j = 1, \ldots, s-1$$

In the same spirit, an expected consumption bundle at time t+1 can be defined as the vector $\bar{x}_{1,t+1}, \ldots, \bar{x}_{s-1,t+1}$ where each of these elements is defined using expected consumption probabilities at t+1:

$$\bar{x}_{j,t+1} \equiv \sum_{i=1}^{n} E[P_{i,t+1}]\, x_{ij} \qquad j = 1, \ldots, s-1$$

Then, using (3) and simplifying, we have:

$$\bar{x}_{j,t+1} = \bar{x}_{j,t} + \sum_{i=1}^{n} (x_{ij} - \bar{x}_{j,t}) P_{i,t} \alpha(U_i) \tag{4}$$

Since in this model the consumer is normally out of equilibrium, $\bar{x}_{j,t+1} \neq \bar{x}_{j,t}$. Nevertheless, we can address the usual questions regarding the effect of changes in prices and incomes in period t upon the consumption which is expected for the following period. For example, to determine the effect of a change in income in period t upon expected consumption in period t+1, we differerentiate (4) with respect to I, recognizing that the vector $P_{1,t},\dots,P_{n,t}$ is not affected by any small price or income change during t (and therefore that $\frac{\partial x_{s,t}}{\partial I} = 1$):

$$\frac{d\bar{x}_{j,t+1}}{dI} = \sum_{i=1}^{n} (x_{ij} - \bar{x}_{j,t})P_{i,t}\alpha'(U_i)\frac{dU_i}{dx_s} \text{ for } j=1,\dots,s-1 \qquad (5)$$

As assumption used to simplify consumer models in a wide variety of applications states that "the marginal utility of money (income) is constant."[12] This assumption, referring as it does to a comparison between long run equilibrium states, is not readily applicable to this short run disequilibrium model. We can make the more restrictive assumption that the marginal utility of one particular good is constant over some "relevant" range, however, and recognize that the usual assumption follows as a consequence. Suppose that good to be x_s. In the context of the learning model, the assumption would then be: "the marginal reinforcing power of small additions to x_s is constant." This proposition can be written: $\alpha'(U_i)\frac{\partial U_i}{\partial x_s} = b$ where b is a constant. Substituting this into equation (5) yields: $\frac{d\bar{x}_{j,t+1}}{dI} = b\Sigma(x_{ij} - \bar{x}_{j,t})\, P_{i,t}$.

Applying the definition that $\bar{x}_{j,t} = \Sigma P_{i,t}x_{i,j}$ this gives $\frac{d\bar{x}_{j,t+1}}{dI} = 0$ for $j = 1,\dots,s-1$, and this is the same result which constant marginal utility of money implies for the static consumer maximization model.

Of course, $\bar{x}_{j,t+1}$ is only an expectation. Although in large populations, we might expect the change in market demand to be negligible, the purchases of individual consumers could vary markedly from one another and from one time period to the next, depending upon the values of the choice probabilities. It is an interesting contrast to the maximization theory that the demand function of any particular learning consumer is not a microcosm of the market but may actually contain actions which are contrary to the tendency of consumers as a whole, even if everyone has the same "tastes." Moreover, it is clear from equation (5) that, except in our special case of constant marginal reinforcing power of cash, the impact of a change of income upon expected consumption depends heavily upon the probability vector in time t. Even the direction of the impact could be altered by altering the values of the probabilities, *ceteris paribus*.

An example of the dynamic workings of this model is given in Figure 1. A utility function with "consumption" and "saving" as arguments was used to generate a hypothetical consumption function whose long run marginal propensity to consume is 0.8.[13] An approximation of the series $\bar{x}_t$, $\bar{x}_{t+1}$, $\bar{x}_{t+2}, \ldots$, is generated by substituting $E[P_{i,t+r}]$ for $P_{i,t}$ in equation (3) to obtain an estimate of $E[P_{i,t+r+1}]$ for $r = 1, \ldots$.[14] In this manner, a series of expected consumption levels is obtained as a function of income, as well as of the time lag following a once-and-for-all change in income. As can be seen from the figure, the very short run consumption function (labelled t = 1) displays very little responsiveness to income changes, whereas the longer run functions (t = 15, t = 50, t = 100) approach more and more closely to the "equilibrium" relationship of C = 0.8I. (Note, however, that since we have not specified *a priori* the length of one time period as used in the model, the lags t = 1, t = 15, t = 50 and t = 100 can be used for comparison purposes only, and do not refer to specific values of real time.)

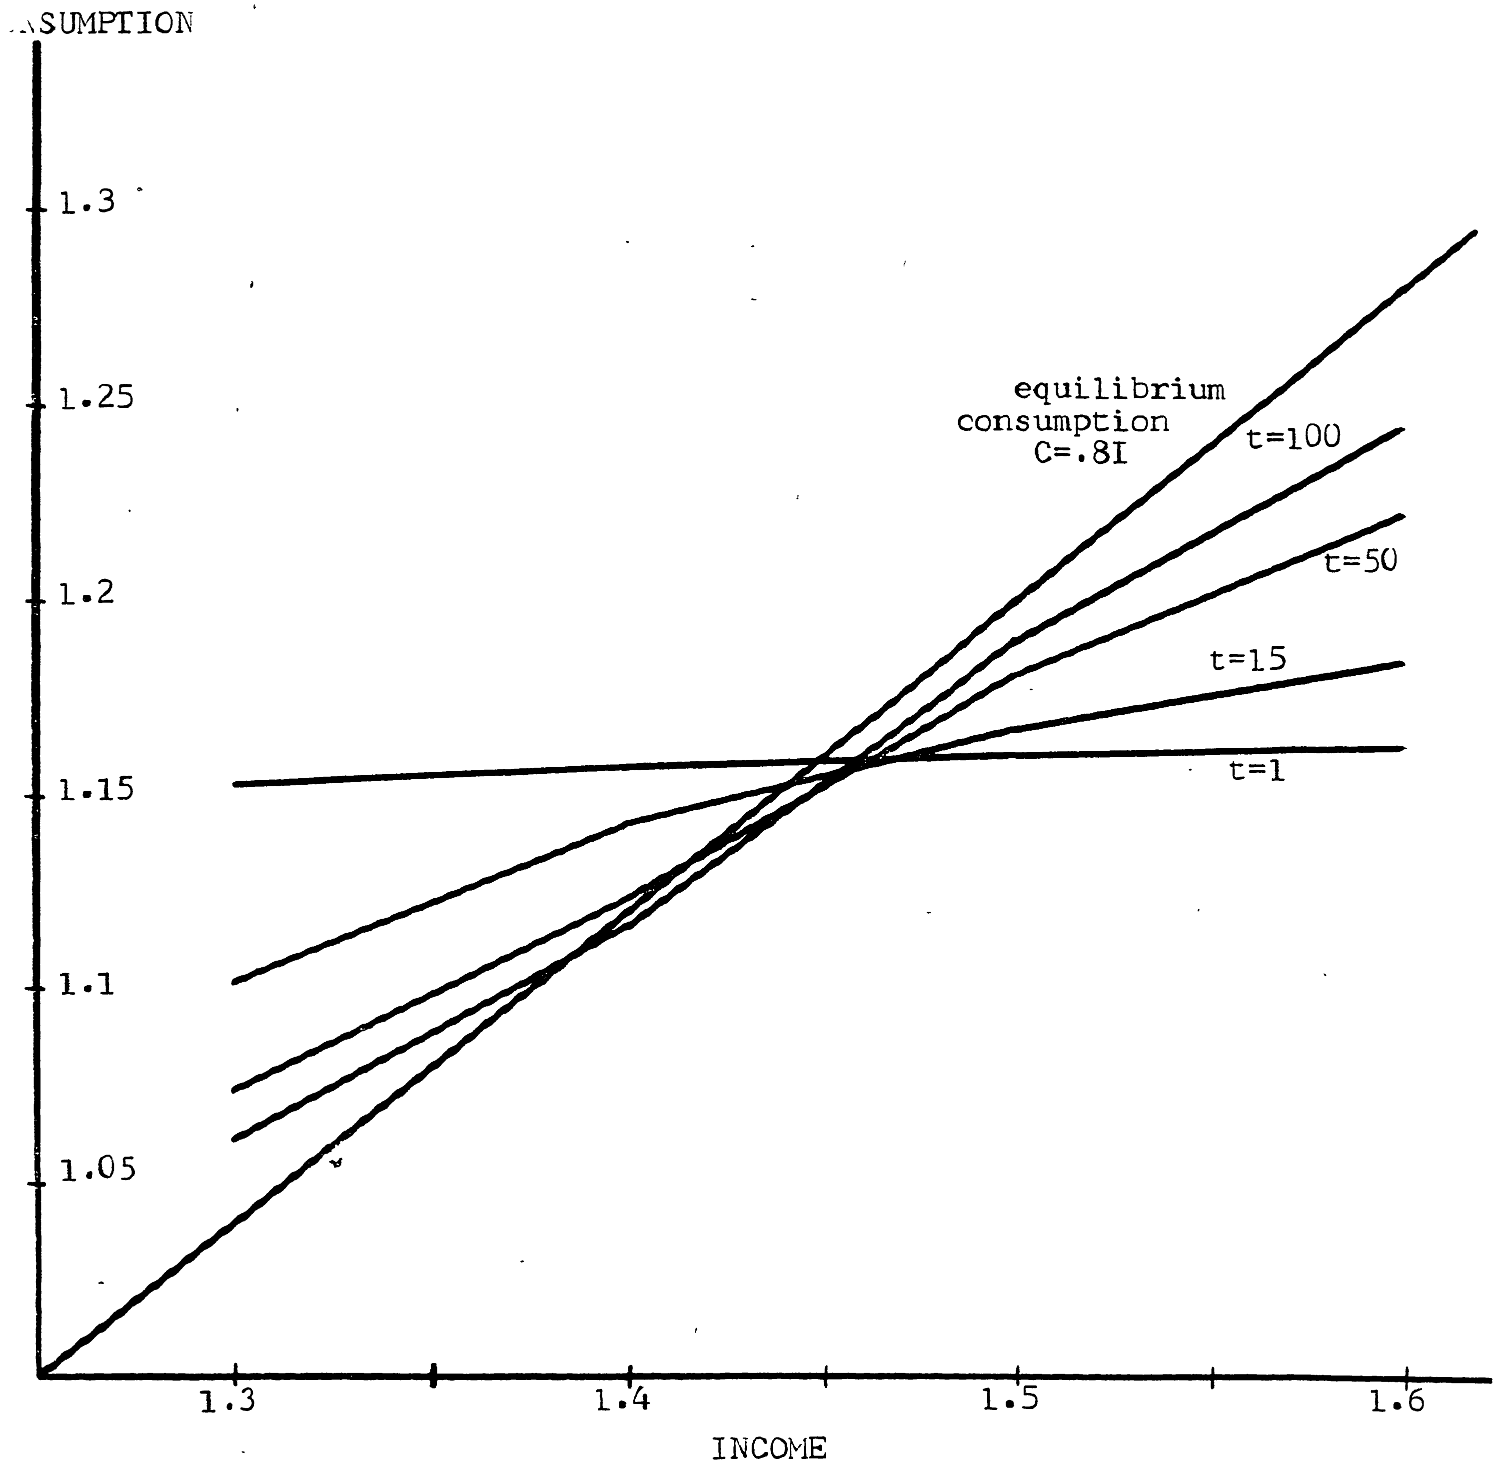
CONSUMPTION
1.3
1.25
1.2
1.15
1.1
1.05
equilibrium
consumption
C=.8I
t=100
t=50
t=15
t=1
1.3
1.4
1.5
1.6
INCOME

Figure 1

The short run dependence of the model upon the stochastic elements of the theory is equally evident in the case of a change in price. Suppose we change the price of commodity k and use equation (4) to determine its impact upon the expected consumption of commodity j:

$$\frac{\partial \bar{x}_{j,t+1}}{\partial q_k} = -\sum_{i=1}^{n} (x_{ij} - \bar{x}_{j,t}) P_{i,t} x_{ik} \alpha'(U_i) \frac{\partial U_i}{\partial x_s} \quad j=1,\ldots,s-1 \qquad (6)$$

Consider the consequence of again assuming "constant marginal reinforcing power of x_s"; then (6) becomes:

$$\frac{\partial \bar{x}_{j,t+1}}{\partial q_k} = -b\sum_{i=1}^{n} (x_{ij} - \bar{x}_{j,t}) P_{i,t} x_{ik}$$

which in turn reduces to:

$$\frac{\partial \bar{x}_{j,t+1}}{\partial q_k} = -b\sigma_{kj}^{2}$$

where $\sigma_{kj}^{2} = \Sigma(x_{ij} - \bar{x}_{j,t})(x_{ik} - \bar{x}_{k,t}) P_{i,t}$ which is the covariance between the consumption of x_j and x_k at time t (or simply the variance in the consumption of x_j if we are considering an own demand curve, k = j).

The sign of σ_{kj}^{2} is opposite to the sign of the cross-elasticity of demand between commodities i and j. It is interesting to note that even in this disequilibrium model, the symmetry found in the formal definition of substitutes and complements is preserved. Since $\sigma_{kj}^{2} = \sigma_{jk}^{2}$, we must have $\frac{\partial \bar{x}_{j,t+1}}{\partial q_k} = \frac{\partial \bar{x}_{k,t+1}}{\partial q_j}$. The introduction of income effects (b not constant) would destroy this symmetry just as it does in the case of the ordinary optimization theory.

σ_{kj}^2 is also an index of the extent to which the consumer has settled down on an equilibrium commodity bundle: the more stable the consumption choice period after period, the smaller will be the values of σ_{kj}^2 and the smaller will be the short run impact of a change in prices. For example, if the consumption of a particular quantity of x_j is firmly established as a habit, then $x_{ij} = \overline{x}_j$ for all i with $P_{i,t} \neq 0$ and, of course, that quantity of x_j will continue to be purchased (in the short run) despite the change in q_k. If it is the quantity of x_k which is firmly established as a habit ($x_{ik} = \overline{x}_k$ for all i with $P_{i,t} \neq 0$), then the change in q_k amounts simply to a change in income equal to $\overline{x}_k \Delta q_k$, and such a change in income already has been shown to have no immediate effect upon the short run consumption of any of the commodities. Over the long run, if course, adjustments in consumption will take place following the predictions of maximization theory.

Figure 2 provides an example of the long run and short run demand curves which are generated by this model. Using the same procedure as was applied in Figure 1, values for expected consumption of a commodity were generated as a function of price and length of the lag after a once-and-for-all price change. Again, the very short run demand curve displays very little responsiveness to price, whereas the longer run functions (t = 15, t = 50, t = 100) approach more and more closely to the static equilibrium demand curve.[15]

It is surely no surprise to find that short run elasticities are smaller than long run elasticities. Unlike optimization theory, however, this property is imbedded in this consumer model itself and does not depend

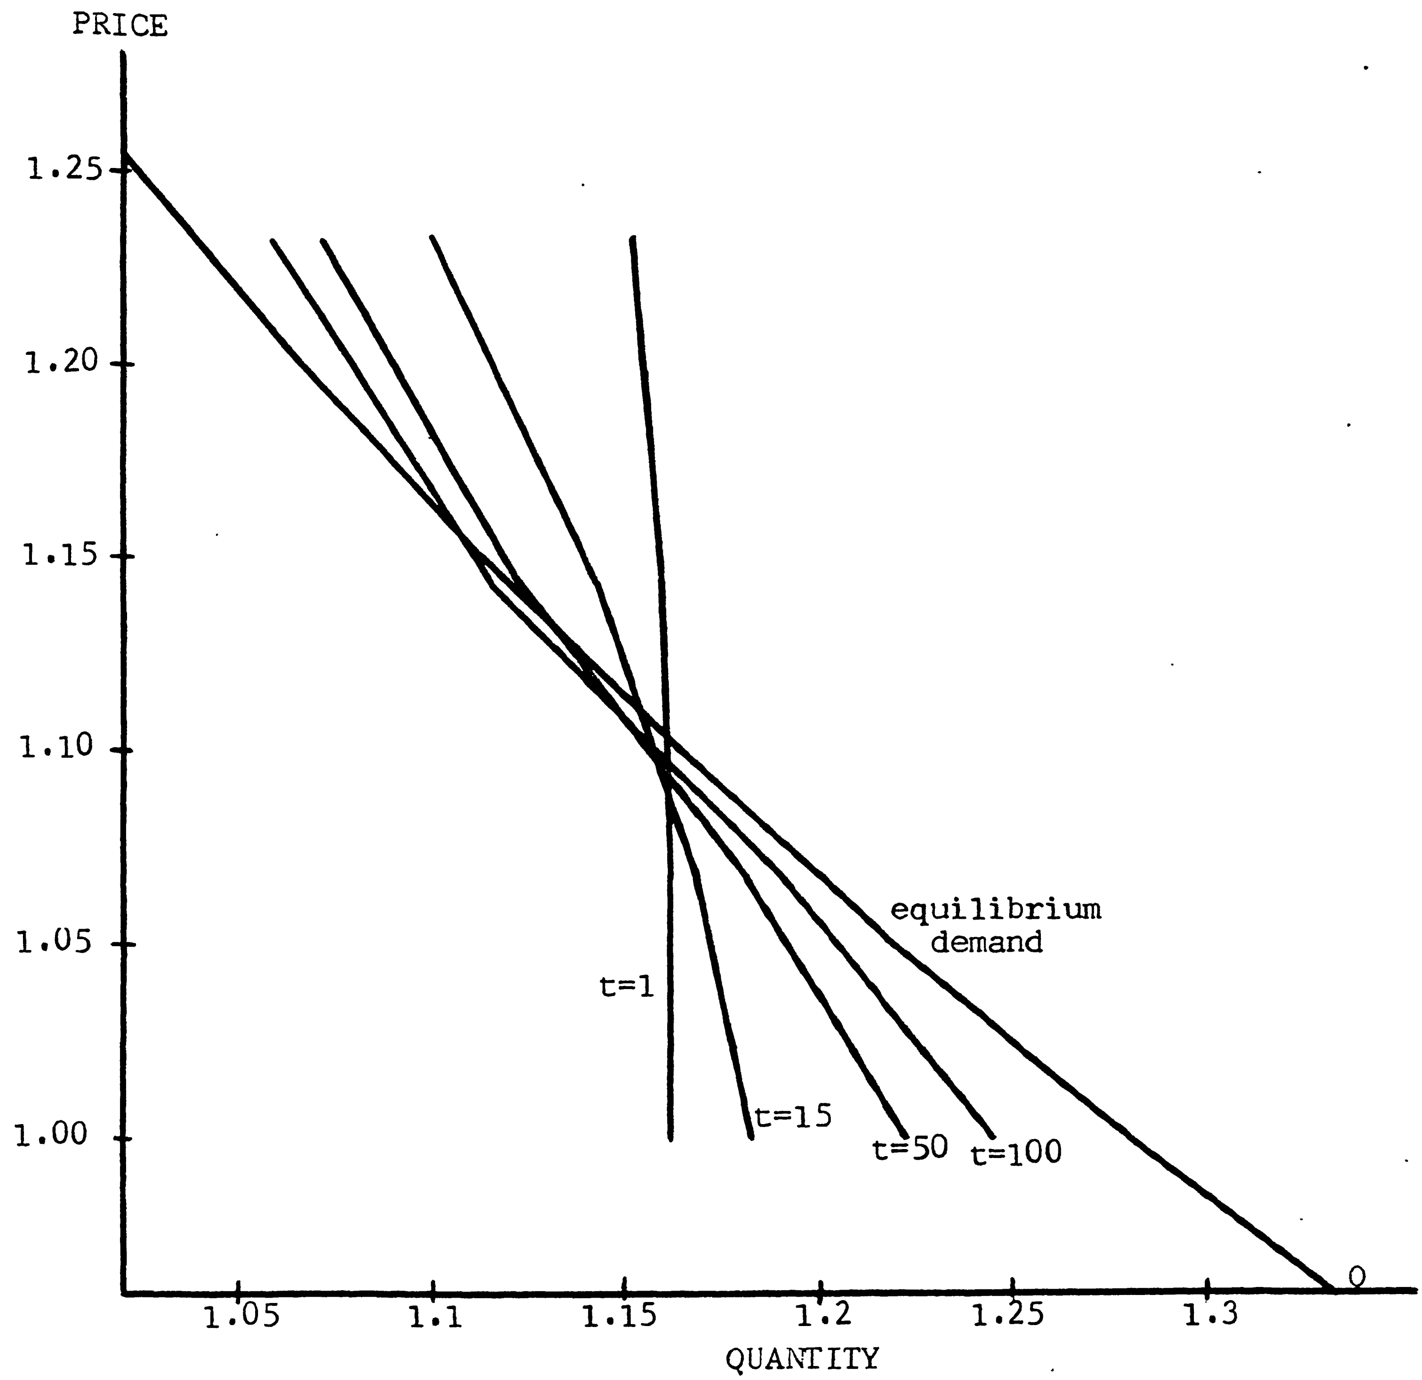
PRICE
1.25
1.20
1.15
1.10
1.05
1.00
equilibrium
demand
t=1
t=15
t=50
t=100
0
1.05
1.1
1.15
1.2
1.25
1.3
QUANTITY

FIGURE 2

for its existence upon any additional hypotheses. We need suggest only that consumers acquire purchasing "habits" which will not be altered immediately by changes in income or price, but instead will adjust only gradually to new circumstances, the speed of adjustment itself depending upon the degree of entrenchment of those habits. Such a proposition is compelling on intuitive grounds alone, and in addition it has a close similarity to hypotheses which already have been advanced in defense of observed consumer demand behavior. Houthakker and Taylor (1970), in their extensive demand study, not only associate lags with "habits," but they even go so far as to develop variables reflecting the consumer's "psychological stock" of "habits." Moreover, their empirical work provides substantial support for this approach, with the psychological stock variable often playing a major role in their equations. However, introduction of such variables into a traditional economic framework is not only entirely foreign to the existing theory found in micro-economics, but in many respects it is an outright contradiction to it. Indeed, it is a little surprising that Houthakker and Taylor do not place more stress upon the fact that their results, interpreted in the fashion they have chosen, provide important disconfirming evidence for the application of orthodox consumer maximization theory to any short run problems.[16]

Suppose that the learning consumer's income were increased for a short time and then reduced to its former level. Such income would act in the short run only to increase stocks of the "buffer" goods (which we have taken to be composed essentially of cash and durables), and since the higher income level is not sustained, the consumer would continue to purchase essentially the same bundle as he had before. Thus we have a marginal propensity to consume out of transitory income which would be small or even zero. Moreover, we have reason to expect purchases of durables to absorb a substantial part of transitory income fluctuations--a proposition which also enjoys significant empirical support.[17]

Finally, we may note that even in the case of consumption function models which are derived explicitly from future-oriented theories, the actual testing has of necessity relied upon historical income and consumption data. Potentially, then, the empirical models are equally consistent with an explicit backward-looking model of consumer learning.

Demonstration Effects

The conventional belief that advertising can induce fully informed consumers to buy what they otherwise would not buy has been a continuing source of embarrassment to those economists who adhere to traditional utility-maximizing models. Some deny the possibility entirely, and stress instead the information content of advertising, and (sometimes) the possibility that advertising actually may contribute to consumption utility itself. Other economists acknowledge the potential importance of advertising, as well as "bandwagon" effects, Snobbism, and Veblen effects, but although Leibenstein (1950) has shown how some of these influences can be taken into account in the construction of market demand curves, and Duesenberry (1952) has stressed their role in aggregate consumption, these phenomena do not seem to have been considered explicitly in any microeconomic model of the household.

It is easy to introduce these effects into learning models, however. For example, if it should happen that advertising or some other demonstration were directed at one particular commodity bundle, then we could describe this effect with equations precisely analogous to (1) and (2):[18]

$$P_{i,t} = P_{i,t}^* \, [1-\beta_i] + \beta_i \tag{7}$$

$$P_{k,t} = P_{k,t}^* \, [1-\beta_i] \text{ for all } k \neq i \tag{8}$$

Here $P_{i,t}^*$ is defined as the probability that the bundle X_i will be selected at time t in the absence of the demonstration. β_i, the effect itself, may be a function of several variables: advertising expenditures, E_i; the fraction of other households consuming that bundle (in the case of bandwagon effects), N_i; the fraction of other households not consuming that bundle (in the case of snob effects), $1-N_i$; or the prices of its component commodities (in the case of Veblen effects). Since we must always satisfy $0 \leq P_{i,t} \leq 1$, it is evident from (7) that we must have $0 \leq \beta_i \leq 1$. In addition, we generally expect stronger demonstrations to have larger effects upon consumption: In the case of advertising, for example,

$$\frac{\partial \beta_i}{\partial E_i} > 0, \text{ and for bandwagon effects, } \frac{\partial \beta_i}{\partial N_i} > 0.$$

The upper bound on β_i, however, implies eventually diminishing marginal impacts; advertising expenditure is subject to diminishing marginal returns, and the bandwagon effect is subject to Leibenstein's "diminishing marginal external consumption effect."

When advertising expenditures on individual products have effects upon many commodity bundles containing that product, we may turn to a more general formulation of the same set of relationships:[19]

$$P_{i,t} = P_{i,t}^*(1-B) + \beta_i \qquad i=1,\ldots,n \tag{9}$$

Here $B \equiv \sum_{k=1}^{n} \beta_k$ and $0 \leq B \leq 1$. Equation (9) reduces to (7) and (8) whenever all $\beta_k = 0$ for $k \neq i$.

Suppose that the consumer happens to select the bundle X_i at time t. Then from equations (1) and (2):[20]

$$P_{i,t}^* = P_{i,t} = \alpha_i(1-P_{i,t}) \tag{10}$$

$$P_{k,t+1}^* = P_{k,t}(1-\alpha_i) \text{ for } k \neq i \tag{11}$$

where α_i is the value of the function $\alpha(U_i)$, and $\overline{\alpha}_t$ is its expected value at time t. X_i, however, is selected with a probability $P_{i,t}$; thus, combining (10) and (11) to obtain the expected value of $P_{i,t+1}^*$ (the same procedure used to obtain (3) from (1) and (2)):

$$E[P_{i,t+1}^*] = P_{i,t}\ (1 - \overline{\alpha}_t + \alpha_i) \tag{12}$$

Demonstration effects in a learning model have both a primary and a secondary impact upon consumption. The primary impact is that already described in (7)-(9): if a particular commodity bundle is advertised during period t, the likelihood of consumption of that bundle is increased. The secondary effect operates through the actual consumption experience. Selection of the bundle X_i leads to the reinforcement of X_i and this makes the choice of X_i more likely in all subsequent periods even if the advertising which was originally responsible for the increase in consumption is terminated. In effect, the advertising encourages the consumption of a particular set of goods, and the consumption habit which is thus formed makes more likely the future consumption of those same goods. This is reflected in the model by the positive value of the expression $(1 - \overline{\alpha}_t + \alpha_i)$ in equation (12), which indicates that an increase in $P_{i,t}$ will necessarily increase $P_{i,t+1}^*$ where

$P_{i,t+1}^*$ represents the likelihood of the consumption of X_i during period t+1 in the absence of any advertising during t+1. Naturally, this result will hold for longer periods as well; the increased likelihood of choice X_i in t+1 means an increased likelihood of the reinforcement of X_i which in turn would increase $P_{i,t+2}^*$ and so on. Thus the effects of advertising will persist for a period of time well beyond the termination of an advertising campaign.[21] This is the case even though the positive sign on $\frac{\partial P_{i,t}}{\partial \beta_i}$ from (9) indicates that a cessation of an advertising campaign will immediately reduce consumption. We can conclude that cessation of an advertising campaign will lead to reductions in demand, but not to demand levels as low as those which would have existed had the advertising never been instituted. This result is supported by the empirical observation of the same phenomonon which is described by Ayllon and Azrin in their Token Economy study, which was referred to earlier in this paper.

We may use (12) also to evaluate the effect upon the choice X_i when it is the alternative bundle X_j which is advertised. Differentiating (12) respect to E_j using (9) to obtain $\frac{\partial P_{i,t}}{\partial E_j}$ and simplifying, we obtain:

$$\frac{\partial E[P_{i,t+1}^*]}{\partial E_j} = - \frac{\partial \beta_j}{\partial E_j} \{P_{i,t}^* (1 + \alpha_i - \bar{\alpha}_t) + P_{i,t}[\alpha_j - \Sigma \alpha_k P_{k,t}^*]\} \quad (15)$$

The sign of this expression is ambiguous. The term $(1 + \alpha_i - \bar{\alpha}_t)$ is always positive, but the sign of the term $[\alpha_j - \Sigma\alpha_k P_{k,t}^*]$ depends upon the relative magnitude of α_j. If α_j is large (or if it is maximal), then this term is also positive and therefore the entire expression is negative. This is the result expected intuitively; as competing choices are advertised, the selection of X_i becomes less likely. Suppose, on the other hand, that α_i

and α_j are both small; then the second term may become sufficiently negative to make the entire expression positive, and we must conclude that it is possible for a demonstration in support of bundle X_j to increase the probability of X_i. In fact, the mechanism for this result is not implausible. If α_i is very small, then it is alternative bundles whose consumption likelihood is expected to increase most rapidly over time. The advertising for X_j increases the likelihood that X_j will be chosen instead. Alternative bundles, however, would tend to reduce $P_{i,t}{}^*$ much more strongly than does the occurrence of X_j itself, and hence $P_{i,t+1}{}^*$ is expected to be higher than it otherwise would have been.

It seems most plausible to regard demonstration effects as temporary in nature. Eventually, experience with the commodities in question, and repeated discoveries that the satisfaction which they produce is not as great as that obtained from other goods, would be likely to reduce the effectiveness of the demonstrations themselves. If this is the case, then the variable β_i is a function of time, with $\frac{\partial \beta_i}{\partial t} < 0$. The long run equilibrium of the model then corresponds to the long run equilibrium of the maximization model, and the only consequence of advertising or bandwagon effects is the (potentially very useful) one of accelerating the discovery and acceptance of new and superior products.

Should it be the case that a particular demonstration effect is persistent over time (or is periodically renewed through innov advertising campaigns), then we must use equation (9) to rewr obtain the expected actual selection probability for bundle) in time:

Revised
C-56
C-57

$$E[P_{i,t+1}] = P_{i,t}(1-B)(1 + \alpha_i - \overline{\alpha}_t) + \beta_i \qquad i=1,...$$

Let us define a long run "equilibrium" vector $P_1,\ldots,P_n$ with the conditions $E[P_{i,t+1}] = P_{i,t}$ for $i=1,\ldots,n$.[22] It is interesting to find that this equilibrium is unique. To see this, first suppose $\beta_i > 0$ for all i, and that there are two probability vectors $P_1,\ldots,P_n$ and $P_1',\ldots,P_n'$ which meet the equilibrium condition. Then using (16) we obtain:

$$\frac{B}{1-B} - \alpha_i + \bar{\alpha} = \frac{\beta_i}{1-B}\frac{1}{P_i}$$

$$\frac{B}{1-B} - \alpha_i + \bar{\alpha}' = \frac{\beta_i}{1-B}\frac{1}{P_i'} \tag{17}$$

where $\bar{\alpha} \equiv \sum_{k=1}^{n} P_k\alpha_k$ and $\bar{\alpha}' \equiv \sum_{k=1}^{n} P_k'\alpha_k$. Subtracting the second equation from the first we obtain: $$\bar{\alpha} - \bar{\alpha}' = \frac{\beta_i}{1-B}\left[\frac{1}{P_i} - \frac{1}{P_1'}\right] \tag{18}$$

This must hold for all i. But since $\frac{\beta_i}{1-B} > 0$, the expression $\frac{1}{P_i} - \frac{1}{P_1'}$ must have the same sign as $\bar{\alpha} - \bar{\alpha}'$ for all i, and this is impossible unless $\bar{\alpha} = \bar{\alpha}'$ since the probabilities summed over i must always equal 1. If $\bar{\alpha} = \bar{\alpha}'$, then we must have $P_i = P_i'$ for every i, and there is only one equilibrium vector. Now let there be some set S of bundles with $\beta_i = 0$ for every $X_i \in S$. In this case, (17) requires $\bar{\alpha} = \bar{\alpha}'$ and then (18) requires $P_i = P_i'$ for any i with $B_i > 0$. For the remaining bundles, (16) may be rewritten: $\sum_{k=1}^{n} P_k\alpha_k = \alpha_i - \frac{B}{1-B}$ for i such that $X_i \in S$.

Since for every X_i not in S, P_i is unique, this system can have multiple solutions only in the trivial case in which two or more bundles, say X_j and X_k, have $\beta_j = \beta_k = 0$ and yield identical utilities: $\alpha_j = \alpha_k$. In this case, P_j and P_k can vary although they must always sum to a unique value (which itself is zero unless α_j is maximal).

Suppose, for a simple example, that the bundle X_i is the obtimal one from the standpoint of the consumer's own utility (α_i is maximal) but that it is unadvertised ($\beta_i = 0$). Equation (16) indicates that for any unadvertised bundle, S_j, which is inferior to X_i, ($\beta_j = 0$, $\alpha_j < \alpha_i$), equilibrium can only occur at $P_j = 0$.[23] Finally, suppose that only one bundle, X_k, is advertised. Then, using (16) and the long run equilibrium conditions $[(P_{k,t+1}] = P_{k,t}$ and $E[P_{i,t+1}] = P_{i,t}$, we obtain:

$$
\begin{aligned}
P_k &= \frac{\beta_k}{1-\beta_k}\,\frac{1}{\alpha_i-\alpha_k} \\
P_i &= 1 - P_k
\end{aligned}
\tag{19}
$$

From (19) we note that if X_i is not greatly superior to X_k, and if X_k is advertised heavily enough, then the likelihood of X_i may be reduced even to zero. Thus advertising or other demonstration effects can eliminate selection of a superior commodity bundle entirely. Under less extreme circumstances, for which $\frac{\beta_k}{1-\beta_k} < \alpha_i - \alpha_k$, the consumer is left vacillating, sometimes choosing the superior bundle, but sometimes induced through outside influence to choose an inferior one instead.

Two Further Areas of Application

Before concluding this discussion, it might be useful to indicate two further areas in which the introduction of learning theory into the consumer model could provide useful hypotheses. The suggestions given here are largely speculative in nature, although they are derived from well-established properties of learning behavior.

Large Price Changes and Inflation

Learned behavior patterns, even thoroughly habitual ones, generally do not occur spontaneously, but are triggered by specific sets of external events, generally referred to as "stimulus" events. In specific experimental situations, these events usually are carefully controlled, and are confined to very simple phenomena such as a ringing bell or a flashing light. In applications to ordinary day-to-day situations, a stimulus event is less well defined and is often said to be composed of a complex combination of circumstances, which might include "need," the time of day and the day of week, opportunity, and prices. Of course, there may be many different situations which generate the same behavior, and moreover, through a process known as "stimulus generalization," novel situations which are similar but not identical to those which have occurred before may also give rise to the same behavior. Thus, for example, if laundry detergent for sale at $.57 a box previously has cost $.54, purchases will not be significantly affected in the short run because the two prices are so similar that the individual does not distinguish between them. This is the case to which the model in previous sections of this paper applies.

On the other hand, without the necessary stimulus, learned behavior may fail to occur. From the standpoint of psychological learning theory, a dramatic change in the price of the soap may prevent habitual purchases from taking place because the usual stimulus event (which includes the usual price) no longer arises. According to the model which has been developed in this paper, this would lead in the short run to an increase in the acquisition of the "buffer" commodities. This provides a possible explanation for the discovery by Katona and others that consumers may increase their rates of saving in times of higher than usual inflation.[24] If many prices undergo unusual increases, the consumer is placed in many unfamiliar situations, and not having learned to buy the commodities at such prices, at the end of the budgetary period the residual commodities, especially saving, will have risen by substantial amounts.[25]

The Influence of Windfall Income

Consumer spending from windfall income has long been regarded as a source of potentially valuable evidence for testing life-cycle or permanent income models. According to the theory, the proceeds from a windfall should be distributed throughout the consumer's planning period and only a fraction of this income should be used for consumption in the current period. According to the same theory, consumption should increase even if the windfall is only an anticipation: an increase in expected future income should increase consumption in the present so long as the consumer expects to receive the income within the planning period.

Some observations by Michael Landsberger (1966) concerning German reparations payments to families in Israel in 1957-1959 provide an interesting challenge to these predictions. In fact, the reparations were negotiated before they were received, yet Landsberger was able to detect no difference in consumption expenditures in 1957-1958 between families which could expect to receive payments and families which could not. Second, Landsberger found a declining propensity to spend out of windfall income as the size of the windfall rises.[26] [Bodkin (1960) has countered this second result using data from the U.S. National Service Life Insurance Dividend of 1950, but the sizes of the windfall payments in this case are all near the small end of Landsgerber's scale.] Both these results are consistent with the view that learning plays an important role in consumption behavior. Thus, for example, the reinforcement model makes current consumption a function only of current and past income, so that even expectations of future windfalls would not have an effect upon current spending behavior. Moreover, we have observed already that it is unlikely that the consumer will choose to purchase items which have not been reinforced previously. If our "buffer" commodities include expensive durable goods as well as saving, relaxations of the budget may lead to the acquisition of some durables, but in the case of large windfalls, the set of desired durable purchases is quickly exhausted; the consumer simply has nothing left other than saving on which he has learned to spend. In addition to Landsberger's data, casual observation of the winners of lotteries, sweepstakes, and similar gambles seems to support this observation: small winnings are spent, but large winnings are banked. Eventually, of course, if income is permanently increased, the occasional consumption "events" which occur will lead to reinforcement of a wider variety of goods (or more expensive ones) and thus the consumer's expenditure will adjust gradually to a new level.

Summary

Strictly speaking, most of the implications in this paper are not new to economists. For many years we have been aware of the importance of learning and habit formation in consumer behavior, and this awareness has been reflected by the freedom with which these concepts are incorporated into econometric studies of consumption. We also have been conscious of a variety of conflicts between empirical data and a strict interpretation of the maximization model, and consequently we have been willing to provide *ad hoc* adjustments to the theory which will suit it to whatever our particular purpose may be at the time. This process has continued for so long, however, that the time surely has come to attempt to introduce some of these modifications into the formal theory. We have attempted to demonstrate that a formal learning model would comprise at least an initial step in this direction while preserving what is most useful in traditional analysis.

Footnotes

[1]A general review of learning theory and extensive references may be found in Hilgard and Bower (1966).

[2]Specifically, the tokens were given for a variety of activities ranging from personal hygiene to jobs working in the hospital kitchen and laundry, and they could be exchanged later for such things as grounds passes, candy, movies, and cigarettes.

[3]It goes without saying that the ethics of these procedures are at least suspect, and public opposition to their use is growing. Our concern here, however, is not how psychologists should employ what is proving to be an astonishingly powerful tool, but whether the demonstrated effectiveness of this device can shed any light upon consumer behavior in ordinary markets.

[4]Ayllon and Arzin (1966), especially Ch. 4, 5, and 9.

[5]Ayllon and Azrin (1966) pp. 103-113.

[6]Unfortunately, the experiments were not continued long enough after the free samples were terminated to enable us to determine whether or not purchases would eventually decline all the way back to their original levels.

[7]Hilgard and Bower (1966) pp. 166-618.

[8]Strictly speaking, psychologists would call saving a "secondary" reinforcer. Payoffs which are known _a priori_ to induce learning are generally described as "primary" reinforcers: if some other event is frequently associated with a primary reinforcer, it is well known that this other event will take on reinforcing powers of its own, despite the fact that it may have no obvious "utility" producing potential at all. It is in this sense that the savings may be viewed as reinforcing, since high savings levels are associated with interest income, parental and cultural approval, reduction in anxiety regarding possible financial crises, and so forth.

[9]This view of saving may be more consistent with various empirical studies which have produced impressively high implied rates of time discounting, ranging from 33% to 100% or more, implying that potential events taking place five or more years in the future would have such an insignificant effect upon present consumption behavior as to be negligable. [Friedman (1957) pp. 142-152; and Holbrook (1967) pp. 750-754.] In one extreme case, Landsberger (1966) found that discovery of large future windfalls had no discernable effect upon current consumption at all. In fact, this personal rate of discount has some other distressing features. One may rationalize the high observed rates by suggesting that real borrowing rates are in fact enormously high, but it seems to be much more common, in life-cycle models particularly, to introduce into the theory a property variously known as "subjective discounting" or "pure" time preference: a skewness which is frequently alleged to characterize consumer preference mappings and which would lead to higher consumption levels in the present than are planned for the future, even in the presence of a positive interest rate. [See, for example, the use of an explicit discount function by Yaari (1964.] The support for this proposition is generally derived from the observation of individuals who do not seem to save "enough:" who do not manage to accumulate sufficient assets to maintain reasonable consumption levels after retirement, who drop out of school at the expense of future income, or who go into debt at high rates of interest to finance unnecessarily high levels of current consumption. Even Irving Fisher (1961, pp. 80-90), despite his rigorous analysis of decision-making over time, tends to describe these phenomena as consequences of lack of "foresight," lack of "self-control," or even of "habits." A case can be made, I think, that Fisher's verbal description of

"impatience" frequently corresponds more closely to the view of savings as a weak reinforcer than it does to the contemporary view that some people's tastes, as an empirical matter, just happen to be constructed so as to give greater weight to current consumption than to future consumption. Certainly Fisher's own description provides something of a contradiction to many contemporary analyses: for example, one can hardly be satisfied with a life-cycle model of a rational maximizing consumer when that model incorporates a phenomenon which according to him, is itself attributable to a "weak will" in conjunction with a "weak intellect."

[10]It is possible to think of $P_{i,t}$ as compounded from the probabilities with which the components of X_i are selected. Thus if $P_{ij,t}$ represents the probability with which X_{ij} is to be chosen at time t, we have $P_{i,t} = \prod_{j=1}^{s} P_{ij,t}$. Since it proves to be much more convenient mathematically, we will focus our attention upon the $P_{i,t}$'s in this paper and in face, $P_{i,t}, \ldots, P_{n,t}$ will be treated as the dependent variables in the model. Naturally, $\sum_{i=1}^{n} P_{i,t} = 1$.

[11]The treatment of savings as a residual is, of course, consistent with Duesenberry's (1952) treatment.

[12]See, for example, Hicks (1946) pp. 26-37.

[13]The function used was $\alpha = C^{.8}S^{.2}$ with values for I of 1.6, 1.5, 1.4, and 1.3. At these values, optimum C is 1.28, 1.20, 1.2, and 1.4 respectively, and for all combinations of these values of C and I, α remains below the necessary upper bound of 1. The initial probability vector used made all consumption levels equally likely $P_i = 1/4$.

[14]The iterative procedure of substitution $E[P_{i,t+r}]$ into equation (3) in place of $P_{i,t}$ in order to obtain an estimate of $E[P_{i,t+r+1}]$ is much

simpler than a direct calculation of $E[P_{i,t+r+1}]$ from the initial vector of $P_{i,t}$'s, but it can introduce a substantial error. For example, in determining $E[P_{i,t+2}]$, this procedure provides an estimate which deviates from the correct value by an amount equal to

$$P_{i,t}[\{\sum_{k=1}^{n} P_{k,t}(\alpha_k - \bar{\alpha})^2 - \alpha_i^{\,2}\}(\bar{\alpha} - \alpha_i) - \sum_{k=1}^{n} P_{k,t}\alpha_k^{\,2}(\alpha_k - \bar{\alpha})]$$

Here $\bar{\alpha} \equiv \sum_{k=1}^{n} P_{k,t}\alpha_k$. The absolute magnitude of this error can be shown never to exceed $P_{i,t}\alpha_j^{\,2}(\alpha_j - \bar{\alpha})$ where α_j is maximal. In the examples used in the text, this expression indicates that the error in any probability estimate would in no case exceed 6% of the probability estimate itself, and would usually be much smaller.

[15]For this example, the same model was used: $\alpha = C^{.8}S^{.2}$ except that income is held at 1.6 and the price of C is given the alternative values of 1, 1.067, 1.143, and 1.231 at which prices the optimal values of C are 1.28, 1.20, 1.21, and 1.04 as before. The initial probability vector again made the four alternatives equally likely.

[16]As an alternative to the explicit introduction of the learning process, one might obtain a model consistent with the findings of Houthakker and Taylor with a maximization theory incorporating changes in "tastes." This seems to be a much inferior approach, however, in that it uses variations in an unobservable independent variable to "explain" variations in an observable dependent variable. For formal models of taste changes, see Gorman (1967) or Pollack (1970).

[17]See Smith (1962), or Darby (1972).

[18]The learning model which we are using here, despite its empirical heritage, is simply a mathematical representation of a few observed regularities and does not in itself provide any explanation for why individuals

should behave in accordance with them. Indeed, psychologists who specialize in learning behavior seem to be much less prone than economists to invent hypotheses which purport to describe what goes on in the unobservable interiors of their subjects. Instead of speculating on what people "think" or how one might rationalize their actions, learning theorists generally develop their models directly from statements of empirical behavior. In accordance with this approach, we accept the proposition that demonstration effects (in which we mean to include advertising and bandwagon effects) do "work" and resist the temptation to try to explain how they influence consumer "thinking." Our main objective in any case is to show that demonstration effects are compatible with our learning model, and to do this, we need only purpose that just as in token economy experiments, the tokens are presumed to act as proxys for future substantive rewards, so the observation of an allegedly satisfactory consumption experience in others operates as a proxy for a satisfactory experience of one's own. Thus the behavioral effect on the consumer is similar to that which would occur if there had been a tangible reinforcer present. This view permits us to construct a model using equations such as (1) and (2) virtually unchanged from their original form.

[19]Since there are many consumption bundles which contain positive quantities of any given commodity, demonstrations in favor of commodity j tend to increase the probabilities of consumption of several different X_i's, encouraging the consumption of bundles containing large quantities of the commodity in place of those which contain none at all. Presumably, the distribution of the impact over these X_k's is in part a function of the specific situation: the main effect may be to induce non-users of a good

to purchase it, or instead it may be to induce moderate users to purchase more. In either case, the probability associated with each X_i containing the good in question might be increased, but the extent of the increase is not determinable on the strength of any a priori criterion.

[20]An alternative formulation for equations (10) and (11) could be:

$$P_{i,t+1}^* = P_{i,t}^* + \alpha_i(1-P_{i,t}^*) \qquad (10')$$

$$P_{k,t+1}^* = P_{k,t}^*(1-\alpha_i) \qquad (11')$$

One might use this version of the model if one considered advertising to have only a very limited direct effect: if the advertising for a product is terminated before any of that product happens to be purchased, then this formulation would make the net impact of the advertising equal to zero. In fact the qualitative conclusions drawn in this paper from equation (10) and (11) are identical with those which could be derived from (10') and (11').

[21]This result may also be demonstrated using the formulation described in footnote 16, even though, under those assumptions, advertising has no lasting direct effect at all. In that case, equation (14) would be:

$$\frac{\partial\, E[P_{i,t+1}^*]}{\partial\, E_i} = \frac{\partial\, \beta_i}{\partial\, E_i}\, \alpha_i\, (1-P_{i,t}^*)^2 + P_{i,t}^* \frac{\partial\, \beta_i}{\partial\, E_i} \sum_{k\neq 1} \alpha_k P_{k,t}^* \qquad (14')$$

and this expression is also positive.

[22]The "equilibrium" defined by this condition does represent a central tendency of the model; nevertheless, even if $P_{i,t} = P_i$ for all i, only one choice will actually be made at time t and the reinforcement of that

particular bundle will naturally lead to $P_{i,t+1} \neq P_{i,t}$. Thus whenever $E[P_{i,t+1}] \neq 1$, we know that the actual $P_{i,t+1}$ will not equal the expectation. We are led to describe our condition as an "equilibrium" only in the sense that there is a systematic tendency toward $P_1,\ldots, P_n$, and hence we can identify those X_i's whose corresponding P_i's are positive, as bundles which will continue to be selected over time, roughly with those probabilities. Obviously, we must have $P_i > 0$ for $i = 1,\ldots, n$ and therefore if our condition implied a negative value for P_i, that probability would in fact fall only to 0.

[23]If we are in equilibrium with $P_j > 0$, then from (16): $P_j = P_j(1-B)(1+\alpha_j-\bar{\alpha})$ and this requires $(1-B)(1+\alpha_j-\bar{\alpha}) = 1$. Since $\alpha_i > \alpha_j$, it follows that $(1-B)(1+\alpha_i-\bar{\alpha}) > 1$, and from (16) $E[P_{i,t+1}] > P_{i,t}$ which contradicts the assertion of equilibrium.

[24]The observation that inflation tends to lead to increased saving was made some years ago by George Katona (1960). Further empirical support has recently been contributed by Juster and Watchtel (1972) and by Taylor, L.D.: "Savings of U.S. Households: Evidence from the Quarterly Flow of Funds," forthcoming.

[25]Two qualifications must be made to this conclusion. First, for the sake of simplicity in exposition and analysis, this paper embodies a very simple model in which the actual purchases of specific commodities are the objects of the learning process. It would be desirable ultimately to substitute simple learned rules of thumb for many types of purchases. Thus if the price of steak rises, for example, the household does not simply go hungry, but instead substitutes via a learned rule of thumb some adequate alternative. Second, our example assumes prices to have been stable over a period of time before the inflation. If prices had been rising rapidly

throughout recent history, the next rise would not present an unfamiliar situation and purchases would go on.

[26] A similar result with reference to transitory income in general has been reported by Darby (1972).

References

Ando, A. and Modigliani, F. "The 'Life Cycle' Hypothesis of Saving: Aggregate Implications and Tests," American Economic Review, LIII (March 1963), pp. 55-84.

Ayllon, T. and Azrin, N. The Token Economy. Appleton-Century-Crofts, New York, 1966.

Bodkin, R. "Windfall Income and Consumption," with comment by Milton Friedman, in Friend, I. and Jones, R., ed. Consumption and Saving, Vol. II, University of Pennsylvania Press, 1960, pp. 175-212.

Bush, R. R. and Mosteller, F. Stochastic Models for Learning. John Wiley & Sons, Inc., New York, 1955.

Cross, J. G. "A Stochastic Learning Model of Economic Behavior," Quarterly Journal of Economics, May 1973.

Darby, M. R. "The Allocation of Transitory Income Among Consumers' Assets," American Economic Review LXII (December 1972), pp. 928-941.

Duesenberry, J. S. Income, Saving, and the Theory of Consumer Behavior. Harvard University Press, Cambridge, 1952.

Fisher, I. The Theory of Interest. Augustus M. Kelly reprint, 1961, pp. 80-90.

Friedman, A. A Theory of the Consumption Function, Princeton University Press, Princeton, New Jersey, 1957.

Gorman, W. M. "Tastes, Habits, and Choice," International Economic Review, Vol. 8, 2 (June 1967).

Hicks, J. R. Value and Capital. Oxford University Press, London, 1946, pp. 26-37.

Hilgard, E. R. and Bower, G. H. Theories of Learning, Appleton-Century-Crofts, New York, 1968.

Holbrook, R. S. "The Three Year Horizon: An Analysis of the Evidence," Journal of Political Economy, 75 (October 1967), pp. 750-754.

Houthakker, H. S. and Taylor, L. D. Consumer Demand in the United States, Harvard University Press, Cambridge, 2nd ed., 1970.

Juster, T. F. and Wachtel, P. "Inflation and the Consumer," Brookings Papers on Economic Activity, ≠1, 1972.

Katona, G. The Powerful Consumer, McGraw-Hill Publishing Co., 1960.

Kreinin, M. E. "Windfall Income and Consumption--Additional Evidence," American Economic Review LI (June 1961), pp. 388-390.

Landsberger, M. "Windfall Income and Consumption: Comment" with "reply" by Bodkin, R., American Economic Review LVI (June 1966), pp. 534-545.

Leibenstein, H. "Bandwagon, Snob, and Veblin Effects in the Theory of Consumers' Demand," Quarterly Journal of Economics (May 1950), pp. 183-207.

Modigliani, F. and Brumberg, R. "Utility Analysis and the Consumption Function: an Interpretation of Cross-Section Data," in Kurikara, K. K. (ed.), Post-Keynesian Economics, Rutgers University Press, New Jersey, 1954.

Pollack, R. A. "Habit Formation and Dynamic Demand Functions," Journal of Political Economy, Vol. 78 (4, part 1), July-August 1970.

Smith, P. E. "The Demand for Durable Goods: Permanent or Transitory Income?" Journal of Political Economy LXX (October 1962), pp. 500-504.

Taylor, L. D. "Savings of U. S. Households: Evidence from the Quarterly Flow of Funds,"

Yaari, M. E. "On the Consumer's Lifetime Allocation Process," International Economic Review 5, (September 1964), pp. 304-317.