Journal of Economic Behavior and Organization 1 (1980) 197–221.

LEARNING TO SEARCH

John C. CROSS*

University of Michigan, Ann Arbor, MI 48109, USA

Received November 1979, final version received January 1980

Consumer search behavior in a market characterized by differentiated products is investigated in terms of a psychological reinforcement learning model. It is demonstrated that under these assumptions, consumers will vary in their propensities toward search despite similarities in search costs. There is no 'equilibrium' search behavior for an individual household, but there does exist a unique equilibrium distribution of search propensities. Although the behavioral hypothesis which underlies the model is quite different from the usual Bayesian formulations, the mean of the equilibrium distribution has properties which are very similar to those of normative Bayesian search models.

1. Introduction

In response to Professor Stigler's 1961 paper on the Economics of Information, a substantial amount of work has been done on the theory of competitive markets with imperfect information. Research has centered around (1) whether in the presence of poorly informed buyers, competitive markets are capable of sustaining a variety of different prices for the same good [Diamond (1971), Rothschild (1974b), Butters (1977), Reinganum (1979), Salop and Stiglitz (1977)], and (2) given that such price dispersions might exist, what sort of search behavior can be expected from the typical consumer [Kohn and Shavell (1974), McCall (1970), Phelps (1970), Rothschild (1974a), Whipple (1973)]. The purpose of this paper is to investigate these questions by means of a learning paradigm of the sort which psychologists have long employed in elementary models of human behavior. Recent work by Cross (1973) and Himmelweit (1976) has demonstrated how such a psychological theory can usefully be applied to economic problems, particularly in the context of market disequilibria, and the model developed in this paper is designed to include the possibility that buyer search and price dispersions are potential consequences of market disequilibria as well as of imperfect information. Furthermore, we will argue that learning models are in some respects simpler than optimization models and that they can always yield results through simulation, however complicated the market situation may be. Finally, for some examples in

*The author would like to thank W.J. Adams, W.H.L. Anderson, T. Bergstrom, J. Laitner, C. Simon and S. Winter for comments on earlier drafts of this paper.

which analytic solutions can be obtained from ordinary search models, simulations of the learning model are shown to tend toward similar aggregate results, even though it turns out that individual consumers with similar search costs may be found to vary widely in their search propensities.

In general, we will suppose that a consumer investigates n different dealers, and incurs a search cost, C, at each one. We will assume that the order in which the n dealers are encountered is irrelevant; this amounts to saying that a favorable offer from a dealer who was investigated early in the search process can be retrieved costlessly. This implies in turn that if dealers were to differ only in the prices which they charged for a given homogeneous product, then each firm i would find its expected sales, Q_i, to be a declining function of its price, P_i; some consumers who found P_i to be minimal would accept alternative price offers if P_i were raised.[1]

2. Differentiated products

Since we wish to compare the implications of a learning model to those of more conventional theories, we must begin by establishing a reasonably general analytic framework which will justify search behavior even under equilibrium conditions. If we were to confine ourselves to price distributions alone, this would be quite difficult under the conditions which we have specified. An 'equilibrium' price distribution would be characterized by non-negative profits and by the condition that no firm i earn profits which are so high as to induce entry:

$$P_iQ_i - C(Q_i) \leqq K \quad \text{for any } i, \tag{1}$$

where $C(Q)$ is the cost function of any firm, and we presume the number of consumers to be large enough to justify the use of expected sales in place of their realization. K is the smallest profit which will bring about entry, and we allow it to reflect any expectations which any firm might have regarding the effects of its entry upon the price distribution. As Stigler observed in his paper, this condition is likely to require declining average costs: in the limited case of free entry and a large number of competitors, K becomes zero, and since $\mathrm{d}Q_i/\mathrm{d}P_i < 0$, condition (1) can only be satisfied if $(\mathrm{d}/\mathrm{d}Q)(C(Q)/Q) < 0$.

Condition (1) characterizes equilibrium in terms of entry and demands invariant (or bounded) profits across the price distribution. In addition, it is necessary that marginal changes in a firm's price cannot increase that firm's individual profits. We note, however, that just as firm i loses sales as its price

[1] If there is a continuous price distribution $F(P)$ and if each buyer investigates exactly n dealers and then purchases at the lowest price, each firm i will find expected sales Q_i equal to $kn[1-F(P_i)]^{n-1}$, where k is the number of potential buyers per seller in the market.

rises, other firms may experience sales increases as P_i rises, because their prices are now more 'competitive' compared to P_i. Customers who sample both firms i and j will choose to buy from i if $P_i < P_j$, but if P_i rises so that the inequality no longer holds, firm j will enjoy increased sales (even if P_i only rises to $P_i = P_j$ in which case customers may select randomly). Suppose that $K=0$ and firm i were to raise its price to P_j exactly. Then firm j would experience an increase in profits due to the fact of declining average cost. In the context of homogeneous products, however, firm i is now identical to firm j, and it, too, must enjoy the same profit level. Since the profits of firms i and j were identical before the increase in P_i (because $K=0$), firm i must have benefited from its price increase and we conclude that the price dispersion could not have been an equilibrium phenomenon.[2]

One simple way to retain search behavior under equilibrium conditions is to allow for the presence of product differentiation. That is, although firms are concerned only with unit cost, price, and sales, and therefore do not see any significant differences between their own and their competitors' products, consumers may respond differently to product characteristics (such as color, style, or perceived quality) and therefore find significant variations in the desirability of the products.

We characterize a household in two ways: it has a set of tastes for the characteristics of the product, and it has an inclination to investigate more than one supplier before making a final purchasing decision. We will compress the appeal of the product characteristics into a single variable, v, which itself is comparable to price, so that the household can compare two alternative suppliers, i and j, in terms of the net values $v_i - P_i$ and $v_j - P_j$. Naturally, no purchase will be made at all unless these quantities are positive. For the moment, we will also characterize the household with the index n, representing the number of different suppliers the household can be expected to consult before making a decision. Finally, we assume that the v's have an upper bound.

We use a cumulative function $V(v)$ to describe the distribution of the v's which is faced by our n-searching household, and assume $V(v)$ to be continuously differentiable. A similar function $W_i(v)$ describes how the product of firm i is distributed in appeal among its potential customers. In the context of differentiated products, it is, of course, possible for firms to charge different prices simply because of differences in their W-distributions. A firm with a wide distribution may choose to exploit its very-high-v customers with high prices (becoming an 'exclusive' shop) while firms with narrow W-distributions must be content with lower prices. However, we are

[2]This argument is quite different from the one given by Butters (1977) or Diamond (1971) to the effect that if search costs, C, are positive, any firm could raise its price by an amount up to C without losing customers, because that argument does not take into account the possibility that even small changes in price can divert customers to competitors which have already been sampled.

not interested here in these conventional sources of price differences, and we will therefore assume that each firm faces the same distribution $W(v)$ and that $W(v)$ is fixed. For similar reasons, we will make a strong assumption that households are identical in that each faces the same distribution $V(v)$.[3] Since firms as a whole and consumers as a whole are sampling the same population of v's, it follows from these two assumptions that $V(v)$ and $W(v)$ are identical.

Since all firms and consumers face the same v-distributions, P and v must be distributed independently. Let $F(P)$ be the price distribution and $\Gamma(\Delta)$ be the distribution of the net values $v-P$ where $\Delta \equiv v-P$ and the corresponding density function $\gamma(\Delta)$ is given by

$$\gamma(\Delta)=\int_0^\infty V'(P+\Delta)F'(P)\,\mathrm{d}P. \tag{2}$$

The n-searching consumer who visits firm i encounters a net value Δ_i. If $\Delta_i \geq 0$ he buys from i with a probability

$$[\Gamma(\Delta_i)]^{n-1}=\left[\int_0^\infty V(P+\Delta_i)F'(P)\,\mathrm{d}P\right]^{n-1}. \tag{3}$$

The P_i for seller i is given, but the v is distributed according to $V(v)$. If there are k customers per seller in the market, then the expected sales of firm i are given by

$$Q_i=kn\int_{P_i}^\infty\left(\int_0^\infty V(P+v-P_i)F'(P)\,\mathrm{d}P\right)^{n-1}V'(v)\,\mathrm{d}v, \tag{4}$$

where k is the number of potential buyers per seller in the market, and the lower limit to the integration reflects the fact that the consumer will buy only if v exceeds P_i.

It is clear from (4) that regardless of the nature of the distribution F, firm i faces a downward sloping demand curve for any finite value of n. As P_i rises, the cumulative distribution V falls for every value of v, reflecting the loss of potential customers to competitors. Moreover, the demand curve itself has non-zero elasticity, (imperfectly) reflected in this model by the range of the integration.

For the reasons already given, it is impossible to have an equilibrium price dispersion in this model. In order to derive the (Nash) equilibrium market price, we suppose that all firms other than firm i charge a common price P^*,

[3] These are quite restrictive assumptions: if automobile dealers sell products ranging from 'very small' to 'very large', the customer who wants a 'very small' car is likely to see quite a different distribution of values than is the customer who wants an intermediate.

find an optimal price for the remaining firm i, and solve for the P^* which leads to $P_i = P^*$. When all other firms charge P^*, the distribution $F(P)$ collapses to a point, and (4) becomes

$$Q_i = kn \int_{P_i}^{\alpha} [V(P^* + v - P_i)]^{n-1} V'(v)\,dv. \tag{4'}$$

In a case of constant marginal costs, the net profit of firm i would be

$$(P_i - C_u)Q_i - nkC_0 - FC, \tag{5}$$

where C_u is unit production cost, C_0 is the cost of space and personnel necessary to display or demonstrate the product both to those who buy and to those who do not, and FC is fixed cost. If we find the first order conditions for the maximization of (5)[4] and impose our condition that each firm find its profits maximized at the same value of $P_i = P^*$, we find (after some manipulation) that P^* is the solution to (6),

$$1 - [V(P)]^n - n(P - C_u)[V(P)]^{n-1} V'(P) - n(n-1)(P - C_u) \int_P^{\alpha} [V(v)]^{n-2}[V'(v)]^2\,dv = 0. \tag{6}$$

It is a curious fact that this differentiated product equilibrium may reflect a positive association between search activity and equilibrium market price. Suppose, for example, that the distribution $V(v)$ is given by $(v-a)^k$, where $0 \leqq a \leqq v \leqq (a+1)$ and $k > 0$. Then if $C_u = 0$ and $nk \neq 1$, eq. (6) can be solved to give[5]

$$1 = \frac{n(n-1)k^2}{kn-1} P \quad \text{if} \quad a \geqq \frac{kn-1}{n(n-1)k^2}, \tag{7a}$$

$$1 = \frac{n(n-1)k}{kn-1} P + (P-a)^{nk} + P(P-a)^{nk-1} \frac{nk(k-1)}{nk-1} \tag{7b}$$

if

$$a < \frac{kn-1}{n(n-1)k^2}.$$

Suppose that (7a) applies (this implies pricing in a range where all

[4]Since the derivative of (5) is positive for $P_i = 0$ whenever $P^* \geqq 0$ and is negative for sufficiently large P_i, a non-negative optimal P_i is always available.

[5]Eq. (7a) applies whenever $P \leqq a$, (7b) whenever $P > a$. Since $a > 0$, eq. (7b) can be solved only for a $P > 0$. Descartes' rule of signs guarantees that this positive value of P is unique

consumers eventually buy from some dealer), and define P' as the solution to (7a) with $n+1$ substituted for n. Using (7a),

$$P'-P=\frac{2-k(n+1)}{n(n^2-1)k^2}. \tag{8}$$

If we meet the conditions $n>1$, and $2>k(n+1)$ then eq. (8) implies that P rises as n is increased to $n+1$. (Since $n\geq 2$, we note that this result is only possible if $k<2/3$.)[6]

If all households were assumed to search exactly n times, then eq. (6) could be used directly to determine the market price. As we develop the theory further, however, we will find that even households with identical search costs can be expected to display a variety of different search behaviors. If, in fact, the inclination to investigate different sources of supply is distributed across the population according to a vector of probabilities, q_n, and the firm is unable to distinguish the search propensities of various customers, then eq. (6) should be replaced by

$$1-\sum_{n=1}^{\infty}\{[V(P)]^n+n(P-C_n)[V(P)]^{n-1}V'(P)$$
$$+n(n-1)(P-C_n)\int_P^{\infty}[V(v)]^{n-2}V'(v)]^2\,\mathrm{d}v\}q_n=0, \tag{9}$$

where the qualitative properties of eq. (6) will be retained by (9).

3. Consumer search

Optimal search models vary widely in the amount of prior information which is assumed to be available. At one extreme, Kohn and Shavell (1974) assume the distribution of values to be completely known. In this case a sequential search procedure is optimal: defining Δ' to be the best 'offer' which the household has encountered so far in its search, the household goes on searching so long as the cost of search is less than the expected benefit to be obtained from one further sample; that is, whenever

$$C<\int_{\Delta'}^{\infty}(\Delta-\Delta')\gamma(\Delta)\,\mathrm{d}\Delta. \tag{10}$$

[6]The possibility that increased search might have a perverse impact upon equilibrium price corresponds to a general property of models of monopolistic competition which has already been observed by Archibald (1961–1962). Interdependence among the demand functions which face individual firms with differentiated products makes it possible for entry of new firms into a market to shift the demand curves of existing firms. If the nature of the interdependence happens to lead to a general decrease in demand elasticities after entry, a new market equilibrium may come about at higher prices than before. In our model, an increase in n, from the point of view of firm i, is equivalent to an increase in the number of its competitors, and the distribution $V(v)$ produces the perverse elasticity consequence

This procedure suffers from the stringency of its assumption that buyers are fully informed as to the distribution of values, but have no idea where any particular dealer falls in that distribution. A rather more attractive approach is embodied in the Bayesian procedure, suggested by Rothschild (1974a) in which the consumer is assumed to be ignorant of the distribution, γ, so that the search process must be used both to estimate the parameters of the distribution and to obtain favorable outcomes. Rothschild does pay a price for this approach, but it is a price which economists are habitually willing to pay: we must assume that consumers actually do behave 'as if they were solving general stochastic optimization problems which are susceptible to analytical solution by professional statisticians in only a few special cases. Apart from this, the plausibility of the Bayesian view is reduced by the fact that the amount of search is likely to be too limited to provide estimates of distributions in which one could have any confidence. Except for very expensive durables such as housing, search costs are likely to be quite significant compared to the dispersion in perceived product quality. For example, if the density function $\gamma(\Delta)$ is known to be rectangular over [0, 1], the expected number of samples taken during sequential search, n_s, varies with search costs according to

$$n_s = 1/\sqrt{2C}. \tag{11}$$

A value of C equal to 10% of the standard deviation of the distribution would seem to be a modest cost; nevertheless in this case that would amount to $C=0.0289$, and this value substituted into (11) gives $n_s = 4.16$. Since the median number of searches is even smaller than this, this means that for most sequentially-searching households, the question on the margin is essentially whether or not to investigate a third, fourth, or fifth supplier. Since in principle, the form of $\gamma(\Delta)$ is unknown, the rational Bayesian must use his three or four data points to determine both the nature of the distribution and its parameters, and the confidence intervals which would have to be applied to the resulting estimates would be so broad as hardly to be worth the effort. The problem is reduced to the extent that the consumer enters the market with a prior estimate of that distribution, but if this is based on repeated buying experiences in the past, the additional bits of data will have only an updating effect, and to that extent we are back to the theories of Kohn and Shavell.

4. Learning to search

The attractive feature of the Bayesian procedure is its emphasis upon learning. Formal Bayesian models, however, are subject to the two complaints just noted: search costs are too high to permit the acquisition of reliable estimates of the parameters of any individual market distribution, and the image of the household as a mathematical statistician places an extraordinary (although not unprecedented) strain upon our credulity. We may deal with the first of these objections by directing the learning process toward a less ambitious goal: that is, we will assume that the object of the learning process is the value of search as an activity, and that this learning extends across markets rather than being confined to only one. We deal with the second objection through an appeal to the psychologists: rather than simply inventing a Bayesian procedure which is plausibly 'rational', we will employ the reinforcement learning paradigm which most psychologists already accept as a fundamental building block of all human behavior.

For the sake of a simple model, we will make three restrictions: first, we are concerned only with relatively expensive, differentiated products. The motivation for using differentiated goods is already explained. The reason for using goods with high prices is to ensure: (1) that price and quality variations are not overwhelmed by search cost, and (2) that no price encountered is ever actually lower than search costs, making further search unnecessary in any case. Second, we will assume that the buying experience is repeated from time to time in several markets and that the distributions found in those markets are similar (although mean price and quality levels may be entirely different); so that search behavior which may be learned from one market may reasonably be carried over into another. Third, any income effects which arise in consequence of a large number of expensive searches do not affect the valuations, v_i.

The household approaches firm i and is shown a product with appeal v_i and price P_i, giving a net value Δ_i. The household suffers a search cost, C, whenever it seeks out a dealer, and we will use Δ' to represent the most favorable offer which the household has encountered so far in its search.

When the consumer investigates a new source, three outcomes are possible:

S_1 – The search was successful in the sense that the new Δ is larger than Δ' by an amount greater than C. The household has benefited to the extent $(\Delta-\Delta'-C)$.

F_1 – The search was a failure in that despite an improvement in Δ, that improvement does not cover C, and the household has lost $-(\Delta-\Delta'-C)$.

F_2 – The search was a disaster in that $\Delta<\Delta'$, and the household has wasted its search cost, losing C.

Our main hypothesis is that households learn to search on the basis of experience with these three cases. If case S is frequently encountered, search will be encouraged (in a sense, the household is convinced that shopping around does pay off). On the other hand, cases F_1 and F_2 discourage search and at the same time encourage purchase from the best of the dealers already sampled (so long as $\Delta' > 0$).

The term 'learning' will be used in this paper to refer to the psychologist's sense of reinforcement of behavior through experiences of success rather than the sense of 'finding out' as it is most commonly used in the economic search literature. [For a general review see Hilgard and Bower (1968).] The consumer is not assumed to use the results of his searching to estimate statistical properties of the market, but only to react more or less passively to the encouragement and discouragement which market experiences provide. Moreover, in keeping with all psychological models of learning behavior, the dependent variable is a choice probability rather than a determinate decision. In this case, the household is characterized by a propensity to stop searching and buy from the source which is currently most favorable, where this propensity is given by $1-\phi$. Search continues with a likelihood ϕ, and the experience gained through searching modifies ϕ until the process stops (the likelihood of stopping being itself governed by ϕ). At the market level, one could reinterpret ϕ in a way which is more conventional from the point of view of economists: ϕ could be defined to be the expected proportion of households which, given identical experiences, would go on searching. Such a model would arise, for example, if consumers were Bayesians, but with differing priors. For many purposes, this interpretation is consistent with the formal material which follows, although the equilibrium properties which apply to the psychological learning theory would not hold in a Bayesian context.

Suppose the household has reached stage t. and has gone on to investigate a new dealer, $t+1$. The probability ϕ_t is modified to ϕ_{t+1} according to a series of transition rules. Formal learning models usually present these transition rules in some specific mathematical form, either as an operator on ϕ_t or as a Markov matrix representing transitions among a finite set of values of ϕ. Indeed, much of the literature in mathematical learning theory is devoted to comparisons among quite restrictive analytic representations [see Coombs, Dawes and Tversky (1970), or Norman (1972)]. This tendency to adhere to a particular functional form also characterizes the applications of learning theory to economics which are found in Cross (1973) and Himmelweit (1976). Nevertheless, the transition rules are in principle quite general and are restricted only by the requirements that $0<\phi_t<1 \Rightarrow 0<\phi_{t+1}<1$ and that $\phi_{t+1}-\phi_t$ be a positive monotonic function of the degree of success which the associated behavior encounters. In this paper, we will adhere to the representation of learning as an operator on ϕ_t, but write our transition rules in the general form

$$\phi_{t+1}-\phi_t = L(\phi_t, R), \tag{12}$$

where R is some index of 'reward' magnitude.

We attribute the following properties to the function $L(\phi_t, R)$:

(a) $L(\phi_{i,t}, R_i) > 0$ for all $R_i > 0$, $\phi_{i,t} < 1$,

(b) $L(\phi_{i,t}, R_i) < 1 - \phi_t$ for all $\phi_t < 1$,

(c) $L(1, R_i) = 0$ for all $R_i > 0$.

[Assumptions (b) and (c) preserve the upper bound of 1 on ϕ. The strict inequality in (b) corresponds to most formulations of learning functions; it has the effect of preventing the model from 'locking in' to some behavior on the basis of one or two chance events.]

(d) $L(\phi_t, 0) = 0$ (A zero reward leaves behavior unaffected.)

(e) $L(\phi_t, R_i)$ is continuously differentiable.

(f) $\dfrac{\partial L(\phi_t, R)}{\partial R} > 0$ for any $\phi_t < 1$. (Larger rewards increase the likelihood that an action will be repeated.)[7]

(g) $\dfrac{\partial L(\phi_t, R)}{\partial \phi_t} > -1$. This condition ensures that given a value of R, ϕ_{t+1} varies positively with ϕ_t.

In binary choice situations such as we are describing ('search' and 'stop searching'), learning functions are applied symmetrically. Thus if a search is successful, the degree of success would be used to reinforce further search. If the search is a failure, the degree of failure would be used to reinforce *stopping*. We will therefore apply the learning function to either ϕ or to $1-\phi$ depending upon the circumstances, and by definition then, our values for R will always be non-negative. Using this formulation, the transition rules are described below:

Case S. Search is successful, so that the likelihood of further search is increased,

$$\phi_{t+1} = \phi_t + L(\phi_t, \Delta_{t+1} - \Delta' - C). \tag{12a}$$

Case F_1. Search is a failure because search costs are not fully recovered, so that the likelihood of stopping is increased,

$$1 - \phi_{t+1} - (1 - \phi_t) = L(1 - \phi_t, C + \Delta'_t - \Delta_{t+1}).$$

[7] It is worth noting that the upper bound of 1 on φ and the positive derivative of L with respect to R impose a condition of 'eventually diminishing marginal effectiveness' of R.

This may be rewritten,

$$\phi_{t+1}=\phi_t-L(1-\phi_t, C+\Delta_t'-\Delta_{t+1}). \tag{13b}$$

Case F_2. Search is a disaster, also increasing $1-\phi$,

$$\phi_{t+1}=\phi_t-L(1-\phi_t, C). \tag{13c}$$

Define ϕ_T to be the value which ϕ happens to have when the household terminates the search process. From (13), we can show that ϕ_T has the following properties:

Property A. $C>0 \Rightarrow \phi_t<1$ for large t. Hence search costs prevent indefinite search, and the household will eventually terminate. Thus $C>0$ implies existence of a $\phi_T<1$.

Proof. Suppose we have a $\phi_t=1$, then from (13) we have

Case S. $\phi_{t+1}=1$ from assumption (c).

Case F_1. $\phi_{t+1}=1-L(0, C+\Delta_t'-\Delta_{t+1})<1$ from assumption (a).

Case F_2. $\phi_{t+1}=1-L(0, C)<1$ from assumption (a).

We would fail to encounter a $\phi_{t+1}<1$ only if the household encountered an infinite sequence of successes, which is a sequence of Δ_t's for which $\Delta_{t+1} \geqq \Delta_t + C$ for all t. If the values of Δ are all bounded, this is impossible.

Property B. $C=0 \Rightarrow \phi_T \geqq \phi_0$ for any $\phi_0<1$. This implies that if search costs are zero, search probabilities increase (with probability one) toward a limit of 1 (which would have the household investigate entire markets).

Proof. Suppose we have a $\phi_t<1$. Then from (13)

Case S. $\phi_{t+1}=\phi_t+L(\phi_t, \Delta_{t+1}-\Delta_t'-C)>\phi_t$.

Case F_1. Does not exist.

Case F_2. $\phi_{t+1}=\phi_t-L(\phi_t, 0)=\phi_t$ by assumption (d).

Therefore since $L(\cdot)$ approaches 0 only at $\phi_t=1$, $\phi_{t+1} \geqq \phi_t$ always, and since in repeated trials, successful searches will arise, search probabilities will rise monotonically toward 1.

Property C. Suppose the values of Δ have an upper bound Δ_{max}. $C>\Delta_{max} \Rightarrow \phi_T \leqq \phi_0$ for any $\phi_0>0$.

Proof. Suppose $\phi_t>0$.

Case E. Never arises.

Case F_1. $\phi_{t+1}=\phi_t-L(1-\phi_t, C+\Delta'_t-\Delta_{t+1})<\phi_t$.

Case F_2. $\phi_{t+1}=\phi_t-L(1-\phi_t, C)<\phi_t$.

Thus for very large search costs, search probability is monotonically decreasing toward 0.

Using $\gamma(\Delta)$ to represent the density function for Δ_t and eqs. (13), we can determine an expected value for ϕ_{t+1} as a function of any given Δ'_t,

$$
\begin{aligned}
E[\phi_{t+1}\,|\,\Delta'_t]=\phi_t&-\int_0^{\Delta'_t} L(1-\phi_t, C)\gamma(\Delta)\,\mathrm{d}\Delta \\
&-\int_{\Delta'_t}^{\Delta'_t+C} L(1-\phi_t, C+\Delta'_t-\Delta)\gamma(\Delta)\,\mathrm{d}\Delta \\
&+\int_{\Delta'_t+C}^{\infty} L(\phi_t, \Delta-\Delta'_t-C)\gamma(\Delta)\,\mathrm{d}\Delta. \qquad (14)
\end{aligned}
$$

As an example of eq. (14) it is interesting to consider a linear form of the learning function: $L(\phi_t, R)=\alpha R(1-\phi_t)$, where α is small enough to guarantee $\alpha R<1$. In this case, eq. (14) reduces to

$$
\begin{aligned}
E[\phi_{t+1}\,|\,\Delta'_t]=\phi_t&-\alpha\phi_t\left(C-\int_{\Delta'_t}^{\infty}(\Delta-\Delta'_t)\gamma(\Delta)\,\mathrm{d}\Delta\right) \\
&-\alpha(2\phi_t-1)\int_{\Delta'_t+C}^{\infty}(\Delta-\Delta'_t-C)\gamma(\Delta)\,\mathrm{d}\Delta. \qquad (15)
\end{aligned}
$$

The expression in the brackets corresponds to condition (10) – that is, the sign of this expression reflects the optimal stopping rule of sequential search models. The integral in the third term of (15) is always non-negative, and thus the sign of the third term of (15) depends upon whether ϕ_t is larger or smaller than 0.5. According to (15), then, the expected transition from ϕ_t to ϕ_{t+1} given Δ'_t depends upon an optimal stopping rule subject to a central tendency toward $\phi=1/2$. This central tendency is common to psychological learning models. It comes about because any learning model which satisfies $0\leqq\phi\leqq 1$ and which incorporates a positive monotonic relation between payoff and ϕ must have a point of inflection in that relation. If an action is already habitual (ϕ is near to its upper bound of 1), further reinforcements can do little to encourage it still more whereas negative experiences can do a great deal to inhibit it (as we approach the inflection point ϕ becomes more sensitive to payoff). Conversely, negative experiences do little toward inhibiting actions which are already unlikely, but positive reinforcements may have a large impact on future behavior.

Eqs. (13) characterize a Markov process which begins with an initial state, ϕ_0, and produces a distribution of possible values for Δ'_1 depending upon the value encountered from the first dealer. If the household goes on searching, each of these possible values for Δ'_1 produces a joint distribution of possible values for ϕ_2 and Δ'_2 depending upon the value encountered from the second dealer. Corresponding to each pair (ϕ_2, Δ'_2) there is a joint distribution of ϕ_3 and Δ'_3 if a third dealer is sampled, and so on. Except in the case in which the household buys at the first available supplier, the household is virtually certain to complete its market adventure with a terminal value, ϕ_T, different from the one with which it began. We will define a density function $h(\phi_T, \phi_0)$ to describe the set of possible ϕ_T's at the end of a typical searching experience *given that search beyond the first dealer does take place.* (If there is no further search, then $\phi_T = \phi_0$ again.)

So long as $\gamma(\Delta)$ is continuously differentiable and $C>0$, it is clear from our characterization of the learning process that

Property D. $h(0, \phi_0) = 0$ for any $\phi_0 > 0$ [from assumption (b) and the fact that $R_t \geqq 0$ always].

Property E. $h(\phi_T, \phi_0)$ is continuously differentiable in both variables over the intervals $0 > \phi_0 \geqq 1$; $0 \geqq \phi_T \geqq 1$.

Property F. $h(1, \phi_0) = 0$ for any $C > 0$ (from Property A).

Property G. $h(\phi_T, \phi_0) > 0$ for all ϕ_T which fall in an interval $0 < \phi_T \leqq \phi_0 + a < 1$ for some $a > 0$. (Some successes are always possible; hence $a > 0$ and a converges to 0 only at $\phi_0 = 1$. Strings of failures of any magnitude between 0 and C, and of any length are always possible: hence we can obtain any positive $\phi_T < \phi_0$.)

Although we were able to show that in the linear case, the learning model is related to an ordinary sequential stopping rule, the Markov process nevertheless differs from most search models in that it indicates that a positive (although perhaps very small) amount of search is consistent with relatively large values of C. A sequential search rule, for example, would permit no search at all if $\gamma(\Delta)$ were a rectangular distribution over [0, 1] and $C > 0.5$. In contrast, the learning process is compatible with positive search at much higher cost levels.

Define $\xi(\phi_0) \equiv \int_0^1 \phi_T h(\phi_T, \phi_0)\, d\phi_T$ (a function describing the expected value of the terminal ϕ given ϕ_0); then

Property H. For any $C < \Delta_{\max}$ there exists a positive number ε such that $\xi(\phi_0) > \varepsilon$ for any $0 < \phi_0 \leqq 1$.

Proof. At the tth stage of the search process, Δ_t' is distributed according to some density function $\psi(\Delta_t')$, and

$$E[\phi_{t+1}]=\int_0^1 E[\phi_{t+1}|\Delta_t']\psi(\Delta_t')\,\mathrm{d}\Delta_t'.$$

From eq. (14) and assumption (c) of the learning function, we have

$$\frac{\mathrm{d}E[\phi_{t+1}]}{\mathrm{d}\phi_t}>0,$$

and

$$\lim_{\phi_t\to 0} E[\phi_{t+1}]=0-0+\int_0^1\int_{\Delta_t'+C}^{\infty} L(0,\Delta-\Delta_t'-C)\gamma(\Delta)\psi(\Delta_t')\,\mathrm{d}\Delta\,\mathrm{d}\Delta_t'.$$

By assumption (a), the condition $C<\Delta_{\max}$ guarantees that the integral expression is strictly positive, and therefore $E[\phi_{t+1}]$ has some minimal value $\varepsilon_t>0$. Since at every stage, the search process is terminated with a probability $(1-\phi_t)$, the value of $\zeta(\phi_0)$ is composed of a series of values of $E[\varphi_{t+1}]$'s weighted by the convergent series $1-\phi_1$, $\phi_1(1-\phi_2)$, $\phi_1\phi_2(1-\phi_3),\ldots$, and therefore there is some value $\varepsilon>0$ with $\zeta(\phi_0)>\varepsilon$ for any $0<\phi_0\leqq 1$.

Using a proof similar to that of Property H, it is easy to show that

Property I

$$\lim_{\phi_T\to 0}\frac{\mathrm{d}}{\mathrm{d}\phi_T}h(\phi_T,\phi)=0 \quad\text{and}\quad \lim_{\phi_T\to 1}\frac{\mathrm{d}}{\mathrm{d}\phi_T}h(\phi_T,\phi)=0.$$

Since $\zeta(\phi_0)>\varepsilon$ for $\phi_0<1$ and $\zeta(\phi_0)<\phi_0$ for $\phi_0=1$ (from Property F) we have $\varepsilon\leqq\zeta(\phi_0)\leqq 1$ for all $\varepsilon\leqq\phi_0\leqq 1$. Since the Markov process is continuous, the Brouwer Fixed Point Theorem guarantees the existence of a ϕ^* such that $\zeta(\phi^*)=\phi^*$ with $\varepsilon<\phi^*<1$, and one might be tempted to use ϕ^* as an 'equilibrium' point for the model. Quite apart from the obvious difficulty in solving for explicit values of ϕ^*, however, it would not be an entirely satisfactory index of individual household behavior. Even if the household enters every new market with an initial ϕ_0 equal to the ϕ_t with which it terminated the last experience, and every market is characterized by the same value distribution $\gamma(\Delta)$, there can be no consumer 'equilibrium' in the sense that a rigid search habit is established (such as always investigating three dealers). Moreover, the dispersion in possible terminal ϕ-values may be quite large, so that ϕ^* occurs infrequently and in fact, given that initial ϕ-values are similarly distributed, there is no reason to expect ϕ^* actually to be the mean of the distribution. We know only a little more about the search behavior of one household at any point in time than we knew before.

The economy, composed of many households, has much more stable properties. Suppose that initial values ϕ_0 are distributed throughout the population according to some continuous density function $g(\phi_0)$. If each household proceeds to enter a market, search, and terminate search, then the economy will be characterized by a new density function $g'(\phi_T)$, where

$$g'(\phi_T)=(1-\phi_T)g(\phi_T)+\int_0^1 \phi_0 h(\phi_T,\phi_0)g(\phi_0)\,\mathrm{d}\phi_0. \tag{16}$$

Proposition 1. For every $C>0$, there exists a unique distribution $g^*(\phi)$ which satisfies (16) with $g(\phi)=g'(\phi)=g^*(\phi)$ for all ϕ and with $g^*(0)=g^*(1)=0$.

Intuitively, Proposition 1 arises from the fact that the learning operator is compact. Since the operator in (16) is not necessarily compact, a rigorous proof is difficult and is given in the appendix.

It is difficult to find any completely satisfactory criterion for stability of the distribution $g^*(\phi)$. Conceptually, it is simplest to look at the behavior of the mean of $g(\phi)$. Let us define ψ_0, ψ_T, and ψ^* as the means of $g(\phi_0)$, $g'(\phi_T)$, and $g^*(\phi)$, respectively. Since ψ^* is unique, we may define stability by the condition

$$\frac{\mathrm{d}\psi_T}{\mathrm{d}\psi_0}<1 \quad \text{at} \quad g(\cdot)=g^*(\cdot). \tag{17}$$

In general, this condition is not met for every possible perturbation in $g^*(\phi)$. Nevertheless, we can prove that for 'low' learning rates, the process is stable. Suppose that we introduce a new variable into the learning function,

$$L'(\phi_{i,t},R_i)=uL(\phi_{i,t},R_i), \qquad 0<u\leqq 1. \tag{18}$$

The new function $L'(\cdot)$ has all the properties of $L(\cdot)$, and the number u may be used to suggest a 'rate' of learning. We can now prove

Proposition 2. For every perturbation in $g^*(\phi)$ which is s standard deviations away from ψ^*, there exists a learning rate $u_0>0$ such that eq. (17) holds for all $u\leqq u_0$.

Proof. Using (16) we have

$$\psi_T=\psi_0+\int_0^1 \phi_0 g(\phi_0)[s(\phi_0)-\phi_0]\,\mathrm{d}\phi_0.$$

The expression $\phi_0[\zeta(\phi_0)-\phi_0]$ is a generally declining function of ϕ_0 because $\zeta(\phi_0)>\varepsilon$, and $\zeta(1)<1$, and so long as this function has a negative slope, distortions in $g(\phi_0)$ which put greater weight on larger values of ϕ_0 (thereby increasing ψ_0) will necessarily increase ψ_T by less, thus satisfying (17). Unfortunately, exceptions to the negative slope are found in the neighborhood of $\phi_0=0$. However, we have already noted that the function $\zeta(\phi)$ has a fixed point at $\phi=\phi^*$, and so long as ϕ^* is unique, the negative slope is guaranteed in the neighborhood of ϕ^*, and (17) holds. Now as the rate of learning $u\to 0$, $h(\phi_T,\phi^*)$ approaches a 'spike' around ϕ^*, $g^*(\phi)\to h(\phi_T,\phi^*)$, $\psi^*\to\phi^*$, and, by the uniqueness of ψ^*, ϕ^* is unique. For any distortion which is, say, s standard deviations away from ϕ^*, we can find a u_0 small enough to put that distortion in the vicinity of ϕ^*, where (17) must hold.

Finally, we can use (17) to show that increases in search cost will have the expected effect of inhibiting equilibrium search activity:

Proposition 3. For stable distributions, high search costs reduce search activity: $d\psi^*/dC<0$.

Proof. Consider any indefinite sequences of quotes $\Delta_1,\Delta_2,\Delta_3,\ldots$ Corresponding to this a sequence of search probabilities $\phi_0,\phi_1,\phi_2,\ldots$ The expected terminal value of P from this sequence is $\phi^e=\sum_{t=1}^{\infty}\phi_t^2\prod_{i=1}^{t-1}(1-\phi_i)$ which is the sum of the stopping probabilities times the associated values of ϕ_t at the time of stopping. From the learning function, we have $\partial\phi_t/\partial C=\partial\phi_{t-1}/\partial C+\partial L(\phi_t,R)/\partial C$ and since every $\partial L(\phi_t,R)/\partial C<0$ we must have $\partial\phi_t/\partial C<\sum_{i=1}^{t-1}(\partial\phi_i/\partial C)\leq 0$ for all $t\geq 2$. This condition is readily shown to be sufficient for $\partial\phi^e/\partial C<0$, and since this applies to all values of ϕ_0 and to all sequences of quotes, it follows that $d\psi_T/dC<0$ as well. If we begin with an equilibrium distribution $g_1^*(\phi)$ with ψ_1^*, then increasing C will lead to a $\psi_T<\psi_1^*$ and (17) implies that the new equilibrium distribution $g_2^*(\phi)$ must have $\psi_2^*<\psi_1^*$.

In summary, Propositions 1–3 imply that even though consumers all behave differently, and have their behaviors modified differently, the market will eventually come to be characterized by a stable equilibrium distribution of propensities toward search. This distribution, moreover, can be characterized by a mean which has the properties which are usually attributed to individual search processes.

5. Simulations

Analytic derivations of $g^*(\phi)$, even for very simple distributions $\gamma(\Delta)$, are not readily obtainable. The form of $g^*(\phi)$ is generally not even symmetric.

Nevertheless, simulation of the learning process is quite simple, and a computer can produce stable distributions involving thousands of 'households' very cheaply (in our experiments, half a second proved to be sufficient to produce stationary distributions for markets of 10,000 consumers). Many simulations of this sort have been run using various distributions $\gamma(\Delta)$ and alternative functional forms for $L(\cdot)$. Fig. 1 describes a typical set of six simulations, in which the variable under consideration was search cost. In these, $\gamma(\Delta)$ is a simple rectangular distribution, and the learning function is linear with $L(\phi_t, R) = 0.2(1-\phi_t)R$. Examples are given for $C=0.01$, $C=0.05$, $C=0.1$, $C=0.2$, $C=0.3$ and $C=0.5$. Fig. 1 reflects the property that higher search costs reduce the quantity of search, and in addition describes the effect of search costs on the dispersion of the ϕ-distribution. Relatively high and relatively low values of C decrease the width of the distribution, and also skew it toward the extreme values of $\phi=0$ and $\phi=1$.

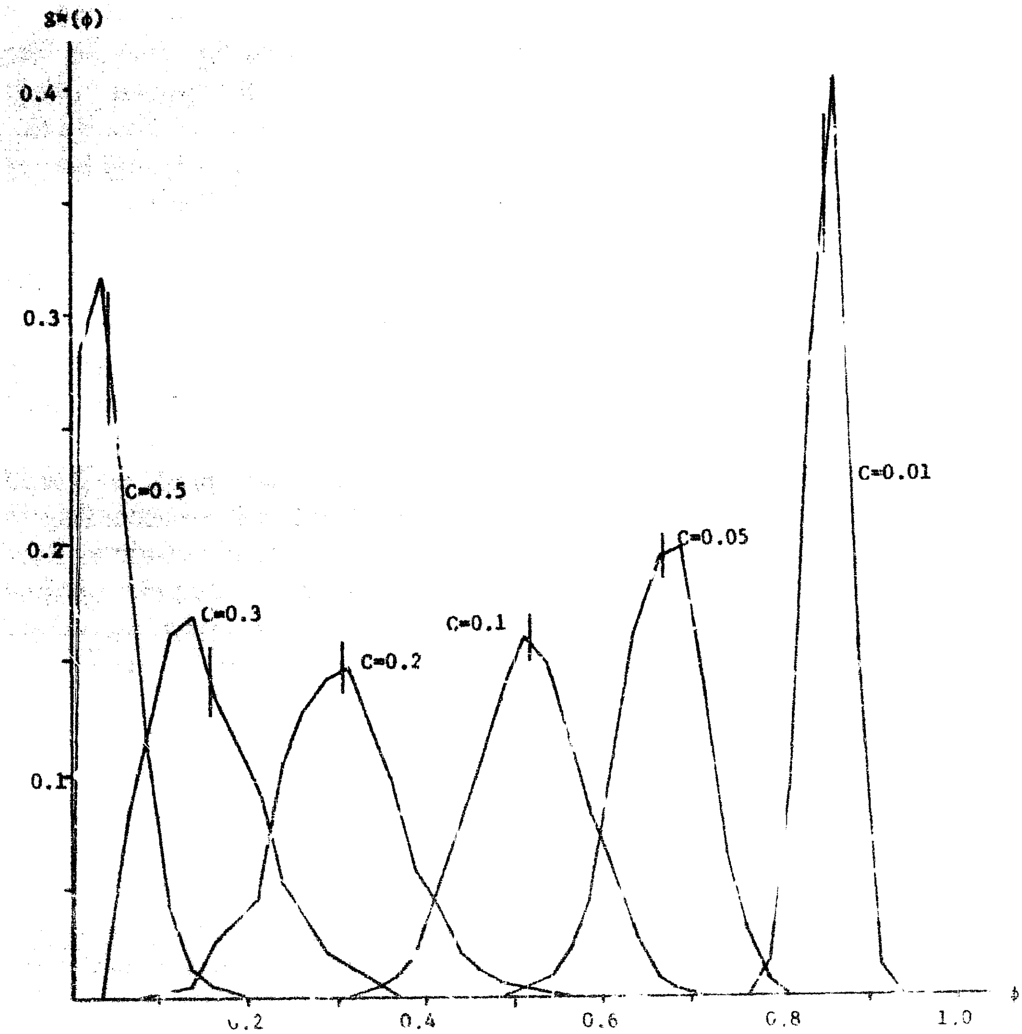

Fig. 1. Simulated distributions $g^*(\phi)$. $L(\phi_t, R) = 0.2R(1-\phi_t)$, $\gamma(\Delta)$ rectangular.

Corresponding to $g^*(\phi)$ is a vector of searching frequencies, Q^*, each of whose elements, q_n^*, describes the expected proportion of the population which will search exactly n times. Table 1 describes the simulation estimates for q_n^* for each of the five cases used to construct fig. 1.

ψ^* is the mean of $g^*(\phi)$, and we will call n^* the corresponding expected number of searches: $n^* = \sum nq_n^*$. Note that n^* may be close to, but is not precisely equal to the reciprocal of $1-\psi^*$. Values of ψ^* and n^* obtained from the simulation experiments are reported as columns 2 and 3 of table 2.

Simulations using different learning functions and different value distributions have produced similar results, although we can hardly claim to have investigated more than a small fraction of all the possible functions. Variations in the learning model, for example, appear to have a considerable effect upon the dispersion of $g^*(\phi)$, but relatively little influence over its mean: in one case, increasing learning rates by a factor of 4 altered the simulated values of ψ^* by less than 15% of the values in table 2 (and for most values of C, much less).

The particular examples used for fig. 1 are interesting because the rectangular distribution $\gamma(\Delta)$ permits comparisons with optimal search models. Suppose that this distribution was fully known to households. Optimal sequential search would then produce an expected number of searches, n_s, given by eq. (11), while the variance in n, σ_s^2, is given by

$$\sigma_s^2 = 1/2C - 1/\sqrt{2C}. \tag{19}$$

Values for n_s and σ_s^2 are listed as columns 5 and 6 of table 2, where they may be compared to n^* and s^2 (the variance of n) from the simulated distributions. The similarities among these numbers are striking, and this comparison might well encourage one in the view that markets indeed operate 'as if' they were composed of households which behave according to principles of optimization. The learning model presumes that consumers have neither the capacity nor the sophistication necessary to employ optimal search rules, yet it produces a very similar relation between search propensity and search costs. In some cases, it produces values of n^* which are indistinguishable from their 'optimal' counterparts.

It would be quite a different matter, however, to assert that the market is composed of households which individually behave as if they were maximizers. Although the total amount of search behavior in our model is similar to that which would be generated by optimal sequential search, our households actually are doing nothing of the kind. They do not have the information on $\gamma(\Delta)$ which would be prerequisite to optimal search behavior, nor do they use any analytic rules to govern their search decisions. Naturally, they do not do as well as fully informed optimizing households could do. The last two columns of table 2 list expected values for realized net payoffs (final Δ' minus total search costs expended) for our learning households, given as $\bar{v}$, and for optimal sequential search, given as v_s.[8]

Table 1
Estimates of q_n^* obtained from simulation.

C	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$	$n=6$	$n=7$
0.01	0.146	0.117	0.104	0.092	0.075	0.066	0.056
0.05	0.340	0.217	0.145	0.098	0.060	0.044	0.031
0.10	0.480	0.244	0.127	0.067	0.038	0.021	0.010
0.20	0.693	0.211	0.065	0.021	0.007	0.002	0.001
0.30	0.841	0.130	0.023	0.004	0.002	0.0	0.0
0.50	0.954	0.043	0.002	—	—	—	—

Table 2
Comparison of search propensities from learning and optimal search models.

Cost	ψ^*	n^*	s^2	n_s	σ_s^2	$\bar{v}$	v_s
0.01	0.85	7.0	40.1	7.1	42.9	0.72	0.86
0.05	0.67	3.0	6.2	3.2	6.8	0.52	0.68
0.10	0.52	2.1	2.4	2.2	2.8	0.41	0.55
0.20	0.31	1.5	0.68	1.6	0.92	0.27	0.37
0.30	0.16	1.2	0.25	1.3	0.38	0.17	0.23
0.50	0.05	1.1	0.06	1.0	0.0	−0.02	0.0

The dynamic responses to particular sequences of quotas are also quite different in the two kinds of model. For example, the sequentially searching household which happens to encounter a long string of low-value quotes will continue to search indefinitely, whereas the learning household may stop at any point and, in accordance with Property A, is certain to stop eventually (this may account for the fact that s^2 in table 2 is generally less than σ_s^2). As another example, suppose a household encounters first a low Δ and then a high one; this will encourage further search whereas a reversal in the order of the quotes would discourage further search. Under optimal search procedures, in contrast, the decision to investigate a third dealer would be unaffected by the ordering of the first two. In short, it is a distinguishing feature of learning models that the individual is not a microcosm of the market, and that the predominant features of a theory for one are not necessarily the predominate features of a theory for the other.

[8]Expected net payoff from optimal sequential search is given by $v_s = 1 - \sqrt{2C}$ if the distribution $f(\Delta)$ is uniform.

6. Dealers revisted

The theory of pricing with which we began this paper followed a simple profit-maximization procedure in which firms took advantage of the limited search propensities of their customers. Having introduced the learning household, however, it is appropriate to treat the firm in a similar fashion — that is, to make the price charged the object of a learning process as well. Since a model of this sort can be found in Cross (1973) (albeit in a restrictive linear functional form), and since we have already developed a more general model in detail in our treatment of consumers, there seems to be little to gain in introducing the notation and mathematics necessary to develop a full description of firm behavior here.[9] According to the linear model, firms will use profit experience to learn to maximize expected profits (or expected 'utility' of profits). This, of course, would imply convergence to a unique price, and market equilibrium would then be characterized by a price P^* which corresponds to the solution to (9) where the values of q_n are obtained from consumer search behavior, and the distribution $\gamma(\Delta)$ is simply a linear displacement of the underlying value distribution: $\gamma(\Delta)=V'(v+P^*)$. Alternative non-linear learning models of firm behavior may fail to produce a unique P^* but generate price distributions instead; nevertheless, the qualitative features of the equilibrium are unaffected, because consumers have an incentive to search in any case.

Even in the case of firm learning behavior which does converge to P^*, we have reason to expect price dispersions to persist for a considerable amount of time. If consumer search follows the pattern we have described, then $(1-\psi^*)$ is an estimate of the fraction of the population which will not proceed beyond the first dealer encountered. For this segment of the market, a dealer is a monopolist, and will do best at a correspondingly high price. Other customers will search more, and in a few cases, the dealer will be in what is equivalent to a highly competitive market. Over time, such a firm will receive a series of contradictory market signals: high prices sometimes will bring high profits and sometimes will drive customers away. The dispersion in these market signals will slow the learning of the firm even if the equilibrium learned behavior maximizes profits. In contrast to the usual competitive market model in which product homogeneity forces identical prices upon all firms at every instant of time, disequilibrium in this differentiated market is characterized by a variety of different prices, some too high and some too low, and this does not necessarily lead even to general shortages or surpluses at the market level. The main forces tending toward equilibrium are lost, and the disequilibrium will tend to reproduce itself.

An example of this sort of problem is presented in fig. 2. This distribution was obtained from a simulation involving 20 competing firms whose learning

[9] A full development of a general model of firm learning may be found in C.R.E.S.T. Discussion Paper c-3, Department of Economics, University of Michigan.

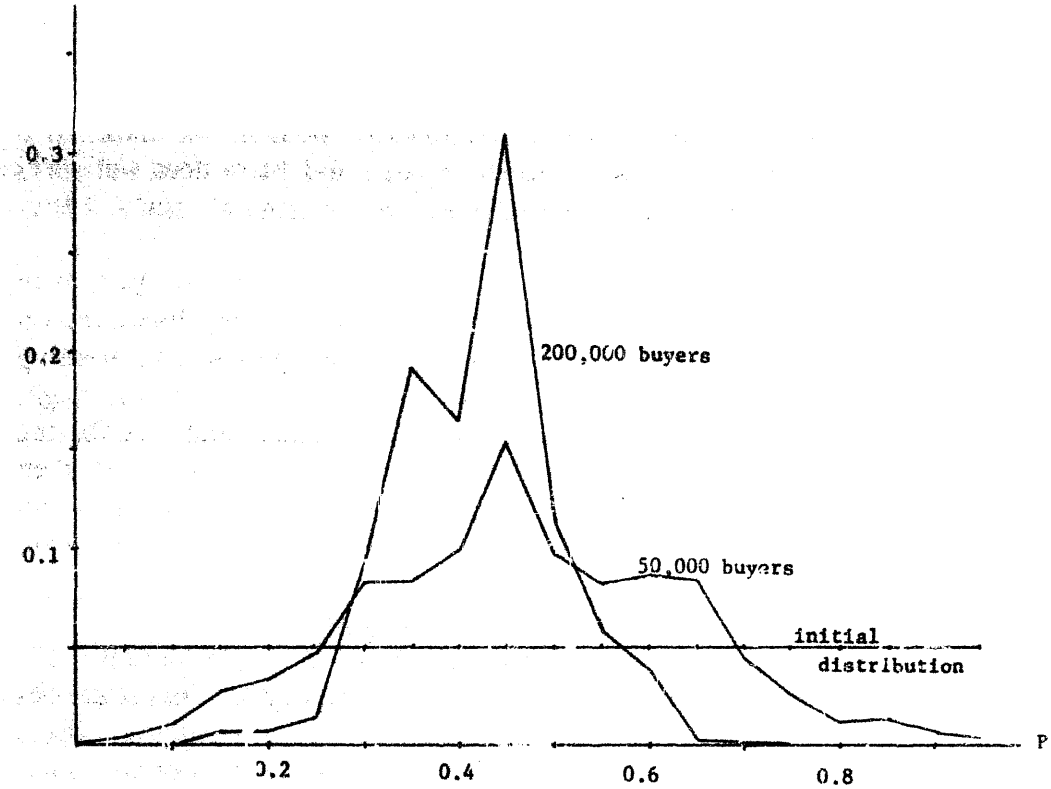

Fig. 2. Simulated density function over market price.

functions were linear (with $\alpha=0.2$), while consumer search costs were 0.1 and consumer learning functions were also linear (with $\alpha=0.2$). $V(v)$ was rectangular for $0\leqq v\leqq 1$. The initial price distribution was rectangular and had converged to the distribution shown after 200,000 consumers had entered the market. (Incidentally, even with this dispersion in pricing behavior, household searching was indistinguishable from that shown for $C=0.1$ in fig. 1.)

7. Extensions of the model

Throughout this discussion, search costs have been represented by a constant, and this is an oversimplification in at least two respects. First, it would certainly be more plausible to treat C as an increasing function of n. Indeed, to many consumers, first encounters with automobile, camera, or house sellers may be so interesting and informative as to justify the search costs by themselves. Further investigations, however, can be sheer tedium, and to many of us, successive searches may become so exasperating that total search costs exceed their market cost components by substantial margins. Second, should a search process be terminated at a dealer whose offer is not the best, some cost must be incurred in retrieving the optimal bid. Were this fact taken into account, consumers in our model presumably would show less inclination to continue searching whenever a 'best' bid has been encountered. For the sake of simplicity, neither of these two modifications has been introduced into the model, and there does not seem to be any reason to expect substantial changes in our general results if they were.

A more awkward problem arises as a consequence of the introduction of product differentiation into the theory. Following the existing literature on search behavior, we have placed most of our emphasis upon pricing, treating dealer characteristics as given. In a more general model, however, one might follow the lead of Lancaster (1966), Rosen (1974) and others and describe the optimal choice of product characteristics. The learning process would then apply to product type as well as to price. Presumably, the model would operate essentially as before, except that convergence to equilibrium from some disequilibrium state would take even longer than in the simple model because of the need for a multidimensional learning process.

This additional variable was not introduced into the model because it is evident that the optimal distribution of prices and product characteristics depends heavily upon the circumstances under which differentiation takes place: the effect of differentiation on preferences, the distribution of consumer tastes, and even on the number of firms in the market. In only a few of the models of optimal product differentiation which are possible, moreover, will it prove to be the case that (a) all firms face identical market conditions, and (b) all consumers see the same distribution of product values. Since these conditions are essential to the derivation of simple optimal search rules which could be compared to the results of the learning model, the question of optimal product differentiation was put aside. Of course, simulation of the learning process could handle these more general models quite readily, but there would have been no means for direct comparisons with the more conventional search rules.

One of the advantages of learning as contrasted to optimization theory is that models of behavior under dynamic conditions may be obtained without a theory of expectations formation. Indeed, it is the absence of an expectations theory that generates most of the objections to Stigler's original model or to optimal stopping rules for sequential search. Both of these require well-specified prior distributions, even though we are provided with no insights into the genesis of such priors. Although it is not necessary, it may be interesting to modify the learning model so that some forms of expectations could be introduced. It is possible, for example, that an optimistic household, upon encountering an unexpectedly low Δ_1, may be inclined to do more searching than another household with more realistic expectations. A family used to paying low rents may move to a more expensive area and search extensively for accommodations, whereas another, accustomed to even higher rents may be satisfied relatively easily. Some of the flavor of such 'satisficing' behavior may be reintroduced into the learning model by means of the distinction between ϕ_0 (the stopping probability with which the search process begins) and ϕ_1, the likelihood that search will stop at the first dealer. Thus we might write $\phi_1 = \phi(\phi_0, \Delta_1 - \Delta^e)$, where Δ^e reflects the consumers' 'expectations', $\partial\phi_1/\partial\Delta_1 > 0$, and $\phi_1 = \phi_0$ if $\Delta_1 = \Delta^e$. Introduction of such a relation into a simulated learning model such as that used to generate the distributions of fig. 1 would not be particularly difficult.

Finally, any theory of learning must address the question of what it is that is to be learned. This paper has focussed upon the simplest variables: whether or not to search, and what price to charge. It is at least as plausible to suggest that what is learned is *rules* of behavior rather than the behavior itself. Consumer search could be governed by learned rules of thumb which may or may not use arguments which are related to optimal search models, and firms might learn to apply specific mark-up formulae, or to tie price-changing decisions to recent market experiences. The variety of possible models is very large, and obviously many of them would enable firms and consumers to do better for themselves than they do using the naive behavior described in this paper. Nevertheless, the essential principle which we have demonstrated for our simple examples would apply in any case, and that is that in a stochastic environment, what is 'learned' by any individual will inevitably incorporate some measure of error. Some consumers will search more than others simply because they happen to have had favorable searching experiences. Some firms may charge more than others because they have been lucky enough to have been visited recently by low-searching consumers. These 'errors' then go on to enrich the stochastic environment in which the learning of others takes place. This point, of course, applies to optimal Bayesian search procedures [as in Rothschild (1974b)] as well as to the psychological process described here. However, if it is a 'rule' that is to be learned rather than the behavior itself, these errors may include the learning of rules which have no 'rational' basis: the consumer may learn to search whenever $\Delta_t > \Delta_{t-1}$ and to stop otherwise, or he may learn to go on searching whenever the last salesman encountered wore glasses. Many firms will learn to keep prices high as Christmastime, but some may learn to maintain high prices whenever the moon is full. In short, one might modify our model by making rules the objects of the learning process, but it would be inconsistent with the spirit of this paper to assume that consumers learn to apply only 'good' rules. With our simple models, we were able to demonstrate the existence of equilibrium distributions of behavior probabilities, but the description of equilibrium distributions of wrong rules of behavior would be well beyond the aspiration level of this paper.

Appendix: Proof of Proposition 1

Define $H(\phi_T, \phi) \equiv (1/\phi_T) h(\phi_T, \phi)$. By Property I, $H(\phi_T, \phi)$ is bounded by $\lim_{\phi_T \to 0} H(\phi_T, \phi) = 0$ and hence from Properties E, F, and G, $H(\phi_T, \phi)$ has a maximal value, M, for some pair (ϕ_T, ϕ) in the interior of $[0,1]^2$. Naturally, $H(\cdot)$ shares Properties E, F, and G with $h(\cdot)$.

Define $\bar{\phi} \equiv \int_0^1 \phi g(\phi)\,\mathrm{d}\phi$ and a function $f(\phi_T) \equiv \bar{\phi} \int_0^1 \phi H(\phi_T, \phi) g(\phi)\,\mathrm{d}\phi$. Since $g(\phi) \geqq 0$, we have $f(\phi_T) \leqq \bar{\phi}^2 M$, and since $H(\phi_T, \phi) \geqq h(\phi_T, \phi)$ for all $\phi_T \in (0, 1]$, we have $f(\phi_T) \geqq \bar{\phi} \int_0^1 \phi h(\phi_T, \phi) g(\phi)\,\mathrm{d}\phi$.

Define $k \equiv \int_0^1 f(\phi_T)\,\mathrm{d}\phi_T$, where we now have

$$\bar{\phi}^2 \leqq k \leqq M \bar{\phi}^2. \tag{A.1}$$

Let $G(\phi_T) = f(\phi_T)/k$ and write the identity

$$G(\phi_T) = (1 - \phi_T) G(\phi_T) + \frac{1}{k} \phi_T f(\phi_T).$$

Multiplying the ϕ_T into $f(\phi_T)$, we obtain

$$G(\phi_T) = (1 - \phi_T) G(\phi_T) + \frac{1}{k} \bar{\phi} \int_0^1 \phi h(\phi_T, \phi) g(\phi)\,\mathrm{d}\phi. \tag{A.2}$$

Integrating over ϕ_T and rearranging terms:

$$\bar{\phi}_T \equiv \int_0^1 \phi_T g(\phi_T)\,\mathrm{d}\phi_T = \frac{\bar{\phi}^2}{k}. \tag{A.3}$$

We note using (A.1) that $\bar{\phi}_T \geqq 1/M$. If $\bar{\phi} \geqq 1/M$, then $k \geqq 1/M^2$ always. $\bar{\phi} \leqq 1$ since $\int_0^1 g(\phi)\,\mathrm{d}\phi = 1$, and hence $0 \leqq G(\phi_T) \leqq M^3$ for all ϕ_T in $[0, 1]$. Define I as a set of continuous functions $g(\cdot)$ with $g(0) = g(1) = 0$, $g(\phi) \geqq 0$, $\int_0^1 \phi g(\phi)\,\mathrm{d}\phi \geqq 1/M$, and $\int_0^1 g(\phi)\,\mathrm{d}\phi = 1$. I is a closed, bounded, convex subset of L^1, and applying the bound on $G(\phi)$, we can use the Ascoli Theorem to show that the function $G(\phi_T)$ defines a compact mapping from I to I [see Munkres (1975)]. Hence by Schauder's fixed point theorem [Smart (1973)] there exists a fixed point $g^*(\phi)$. From (A.3) the value of k at $g^*(\phi)$ is $\bar{\phi}$, and substituting this into (A.2) we see that $g^*(\phi)$ is a fixed point for eq. (16) as required.

Now suppose there are two fixed points $g_1(\phi)$ and $g_2(\phi)$. Define

$$\delta(\phi) \equiv \min[g_1(\phi), g_2(\phi)],$$

$$B_i(\phi) \equiv g_i(\phi) - \delta(\phi), \qquad i = 1, 2. \tag{A.4}$$

At a fixed point, eq. (16) becomes $g_i(\phi_T)=\int_0^1 \phi H(\phi_T,\phi)g_i(\phi)\,\mathrm{d}\phi$, and using (A.4), we obtain

$$B_i(\phi_T)=\int_0^1 \phi H(\phi_T,\phi)B_i(\phi)\,\mathrm{d}\phi+\int_0^1 \phi H(\phi_T,\phi)\delta(\phi)\,\mathrm{d}\phi-\delta(\phi_T).$$

For every ϕ_T in [0, 1], there is an i which gives $B_i(\phi_T)=0$. By Property G, $\int_0^1 \phi H(\phi_T,\phi)B_i(\phi)\,\mathrm{d}\phi\geqq 0$ with the strict inequality holding for some of those values of ϕ_T for which $B_i(\phi_T)=0$ [i.e., those values of ϕ_T which are no more than a units above a ϕ with $B_i(\phi)>0$]. Therefore, $\int_0^1 \phi H(\phi_T,\phi)\delta(\phi)\,\mathrm{d}\phi-\delta(\phi_T)\leqq 0$ for all ϕ_T with the strict inequality holding for some ϕ_T. Multiplying by ϕ_T and integrating, we obtain $\int_0^1 \phi\delta(\phi)\,\mathrm{d}\phi-\int_0^1 \phi_T\delta(\phi_T)\,\mathrm{d}\phi_T<0$, and since this is impossible, we must have $B_i(\phi)=0$ for all ϕ, which means the fixed point is unique.

References

Archibald, G.C., 1961–1962, Chamberlin vs. Chicago, Rev. Ec. Studies, 2–28.

Butters, G., 1977, Equilibrium distributions of sales and advertising prices, Rev. Ec. Stud. 465–492.

Coombs, C.H., R.M. Dawes and A. Tversky, 1970, Mathematical psychology (Prentice-Hall, Englewood Cliffs, NJ).

Cross, J.G., 1973, A stochastic learning model of economic behavior, QJE 87, May, 239–266.

Diamond, P.A., 1971, A model of price adjustment, J. Ec. Theory 3, June, 156–158.

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

Himmelweit, S., 1976, A behavioral model of learning in production, Rev. Ec. Studies 43, June, 329–346.

Kohn, G. and S. Shavell, 1974, The theory of search, J. Ec. Theory 9, Oct., 93–123.

Lancaster, K., 1966, A new approach to consumer theory, JPE 74, 132–156.

McCall, J.J., 1970, Economics of information and job search, QJE, Feb., 113–126.

Munkres, James R., 1975, Topology (Prentice-Hall, Englewood Cliffs, NJ).

Norman, M.F., 1972, Markov processes and learning models (Academic Press, New York).

Phelps, E.S., 1970, Money wage dynamics and labor market equilibrium, in: E.S. Phelps et al., eds., Microeconomic foundations of employment and inflation theory, 124–166.

Reinganum, U.F., 1979, A simple model of equilibrium price dispersion, JPE 87, Aug., 851–858.

Rothschild, M., 1973, Models of market organization with imperfect information: A survey, J. Poli. Ec. 81, Nov./Dec., 1283–1308.

Rothschild, M., 1974a, Searching for the lowest price when the distribution of price is unknown, JPE 82, July/Aug., 689–712.

Rothschild, M., 1974b, A two-armed bandit theory of market pricing, J. Ec. Theory 9, Oct., 185–202.

Rosen, S., 1974, Hedonic prices and implicit markets, JPE 82, 34–55.

Salop and Stiglitz, 1977, Bargains and ripoffs, A model of monopolistically competitive price dispersion, Rev. Ec. Stud., 493–510.

Smart, D.R., 1973, Fixed point theorems (Cambridge University Press, Cambridge).

Stigler, G.J., 1961, The economics of information, JPE 69, June, 213–225.

Whipple, D., 1973, A generalized theory of job search, JPE 81, Sept./Oct., 1170–1188.