Price Differences in Almost Competitive Markets
Author(s): John W. Pratt, David A. Wise and Richard Zeckhauser
Source: *The Quarterly Journal of Economics*, Vol. 93, No. 2 (May, 1979), pp. 189-211
Published by: Oxford University Press
Stable URL: https://www.jstor.org/stable/1883191
Accessed: 27-05-2020 17:35 UTC

PRICE DIFFERENCES IN ALMOST COMPETITIVE MARKETS

JOHN W. PRATT
DAVID A. WISE
RICHARD ZECKHAUSER

I. INTRODUCTION

If positive costs of search insinuate an imperfection into otherwise perfectly competitive markets, a stable equilibrium may have sellers charging different prices. Whether this qualifying wrinkle on the competitive model turns out to be a major crease depends on the extent of the differences among quoted prices.

To gather evidence on diversity among prices, we chose thirty-nine standardized products according to a random procedure described in Section IV.[1] For each product an average of twelve price quotations was obtained. The differences among quoted prices proved much greater than expected.[2] Table I of Section IV gives some statistics on prices quoted for the products.

Many factors may have contributed to these notable differences in quoted price. Here we demonstrate the possibility of a simple explanation based on positive search costs, the fact that it costs a buyer something each time he secures a price quotation.

Sellers are assumed to be expected value maximizers. In setting their prices, they weigh higher profits per sale against the prospect of fewer sales. Buyers employ searching and buying strategies: In deciding whether to seek further price quotations, they balance the prospect of securing a lower price against greater incurred search costs. The confluence of these two sets of decisions is a market outcome.

The definition of an equilibrium is that no seller could increase his expected profits by changing to another permissible price. A consequence of this is that, aside from minor differences due to the indivisibility of sellers, all sellers make equal expected profits if all have equal costs. We often assume as well that all buyers are making optimal use of the information available to them, though this is not a requirement for an equilibrium.

The Quarterly Journal of Economics, May 1979 0033-5533/79/0093-0189$01.00

The central analytic questions we consider are as follows: Is there an equilibrium distribution of sellers' prices? At equilibrium, may sellers with identical characteristics charge different prices? We treat the first question primarily theoretically, the second primarily with analytic examples.

Assumptions

The analysis can swiftly become overwhelmingly complex. Many of its more interesting points, fortunately, can be illustrated using simple models. The following properties, except when noted to the contrary, are assumed throughout our analytical discussion.

Each buyer has constant search cost, possibly differing among buyers. Buyers can go back at no cost and purchase at a previously observed lower price. Buyers have prior beliefs about the distribution of prices. If they are uncertain about it, they update their beliefs as they search. Buyers are risk-neutral; they maximize expected consumer surplus net of search costs. Each buyer attaches sufficient value to one unit to induce him to enter the market. He buys at most one unit, since second and subsequent units are valued at zero.

Each seller has the same, constant marginal cost. The set of permissible prices is finite. Sellers, with long-term experience in the market, know the distribution of buyers' strategies. Those strategies may depend on the distribution of sellers' prices. Each seller picks a price to maximize expected profits, given buyers' strategies and other sellers' prices.

Use of Information

There is a variety of possible structures for the buyers' knowledge and use of information. We shall deal with two extreme cases. In the first, buyers know the exact distribution of sellers' prices. This type of assumption underlies a substantial proportion of the literature on information-related equilibria in both micro- and macroeconomics. In the second case, each of the buyers has prior beliefs about the distribution of sellers, beliefs that are independent of the actual distribution, but that they update as they receive price quotations in the market. We shall refer to the first case as knowledge on the part of the buyer, and to the second case as learning. We shall not explicitly consider the numerous intermediate possibilities where buyers' prior beliefs are influenced by the actual distribution, but not precisely determined by it.

Section II explores models with knowledge and shows that an equilibrium may involve differing prices. Section III turns to models with learning. After describing optimal strategies for buyers, we prove that for any distribution of buyers' strategies, optimal or not, there will be an equilibrium distribution of sellers' prices. We present an example of an equilibrium with differing prices and suggest a general algorithm for finding an equilibrium. Section IV reviews our empirical findings. Section V discusses implications and draws conclusions.

II. EQUILIBRIA IN MODELS WITHOUT LEARNING—THE CASE OF KNOWLEDGE

Once search costs contaminate a competitive world, the concepts defining the competitive market must be extended to include behavioral responses to information. We assume in this section that the buyers have exact knowledge of the probability distribution of sellers' prices and respond to it by searching and buying optimally, given their search costs. A distribution of buyers' strategies results, to which the sellers respond, producing an equilibrium distribution of sellers' prices. We require that this distribution be the same one the buyers used, so as to make their exact knowledge correct. In this sense, it is self-justifying. In brief, a distribution of sellers' prices is self-justifying if it is an equilibrium for the sellers when all buyers respond optimally.[3]

To be specific, let us look first at a buyer's behavior. Suppose that it costs him s to secure each price quotation, and the probability of getting a price quote of x on any trial is known by him to be $f(x)$. Following the dictates of sequential decision theory and its subdiscipline on optimal stopping rules, he maximizes his expected value by continuing to search until he receives a price quotation less than or equal to a certain price y, at which time he buys. The cutoff price y is the highest price at which the expected gain (price reduction) of one further trial is smaller than or equal to its cost,[4] the greatest value of y' with $f(y') > 0$ such that

$$s \geq \sum_{x=0}^{y'} (y' - x)f(x). \tag{1}$$

Note that in the case of knowledge the optimum strategy is myopic, looking only one trial ahead, and that returning to previous price quotations is never appropriate.

The value of y depends on the search cost s and the distribution of prices f. Different buyers may have different search costs. From the distribution of search costs, by tallying the y that goes along with each s over the range of s's, the distribution of cutoff prices can be determined. Let $q(u|f)$ be the probability that a random buyer has cutoff price y.

From this in turn, the expected return to a seller at each price can be determined as follows. Every seller whose price is at or below an individual buyer's cutoff price y has an equal chance to receive his business. If there are N sellers and $F(x)$ is the cumulative distribution of selling prices, then there are $NF(y)$ sellers quoting prices at or below y, so each has probability $1/NF(y)$ of securing the sale. To compute the expected number of sales, $S(x)$, for a seller at price x, we must cumulate these probabilities of sale over all buyers with cutoff prices $y \geq x$. If there are M buyers, this gives

$$S(x) = S(x|f) = (M/N) \sum_{y=x}^{\infty} q(y|f)/F(y).$$

If the marginal cost of provision is constant at c, then the expected return to a seller at price x is $S(x)(x - c)$. If we wish merely to compare either the probability of a sale or expected returns at two different prices, we need not worry about (M/N). Therefore, in subsequent discussion we shall employ the concept of a "standardized expected return," given by

$$\pi(x) = \pi(x|f) = (N/M)S(x)(x - c) \tag{2}$$

$$= (x - c) \sum_{y=x}^{\infty} q(y|f)/F(y);$$

it is the number of sellers times the expected net amount a seller quoting price x receives from a randomly selected buyer.

Case of an Infinite Number of Sellers

Suppose now that the number of sellers is infinite. Then equilibrium requires merely that the sellers all maximize $\pi(x)$, since each is small.[5] The distribution f is therefore in (possibly unstable) equilibrium, and thus self-justifying, if and only if $\pi(x)$ attains its maximum at every x where $f(x) > 0$. Equivalently:

1. All sellers receive the same standardized expected return; i.e., $\pi(x_i) = \pi(x_j)$ for all i and j such that $f(x_i) > 0$ and $f(x_j) > 0$.
2. No price not quoted by a seller would yield a greater expected return than a quoted price; i.e., $\pi(x_i) \geq \pi(x_k)$ for all i and k such that $f(x_i) > 0$ and $f(x_k) = 0$.

All sellers at the monopoly price is a stable equilibrium. (This is true even if demand levels may change with price. Inelastic demand was assumed for convenience in examples. Strictly speaking, it and our other assumptions, taken together, would imply that the monopoly price is the highest permissible price.) A central concern is whether some other equilibrium distribution of sellers' prices may exist, especially one not restricted to a single price.

Arrow and Rothschild [1975] have examined the possibility for multiple-price equilibria under somewhat different assumptions including identical demand curves for buyers, a minimum search cost of some positive amount, and continuous prices. With a large number of sellers they find there is a unique equilibrium with all sellers charging the monopoly price. With a small number of sellers, there may be no equilibrium. The all-at-monopoly price outcome may fail, for example, because sellers defect downward to capture a larger fraction of the market. If reality conformed to the Arrow-Rothschild formulation, there could not be differences among quoted prices at an equilibrium. Market equilibria at which more than one price is quoted are consistent with the optimal behavior models, however, in the realistic case where sellers can only charge prices belonging to a discrete set, like prices in dollars and cents. For illustration, we consider just two prices.

Examples of Equilibria in Two-Price Models

There are two prices, x_1 and x_2, with $x_1 < x_2$. Buyers know the fraction f_1 of sellers selling at price x_1. The optimal search strategy for a buyer will depend on f_1 and on his search cost s. Since there is no learning involved in searching, the optimal strategy for a buyer has one of two forms. Either he can search until he finds an x_1, or he can make one trial and accept his first price. By (1), the buyer will continue searching until he finds an x_1 if

$$s < (x_2 - x_1)f_1; \tag{3}$$

for higher costs, he will stop the first time. (This can also be seen directly by comparing the expected costs of the two strategies, which are $x_1 + s/f_1$ and $f_1x_1 + (1 - f_1)x_2 + s$, respectively, since the expected number of trials to find an x_1 is $1/f_1$.)

Let q_1 be the proportion of buyers who, given current search costs and price distributions, search until they find the lower price. Then the standardized expected returns π_1 and π_2 of sellers at x_1 and x_2 are, by (2),

$$\begin{aligned} \pi_1 &= (x_1 - c)(q_1/f_1 + 1 - q_1) \\ \pi_2 &= (x_2 - c)(1 - q_1). \end{aligned} \tag{4}$$

The key element of interest in this situation is the relationship be-

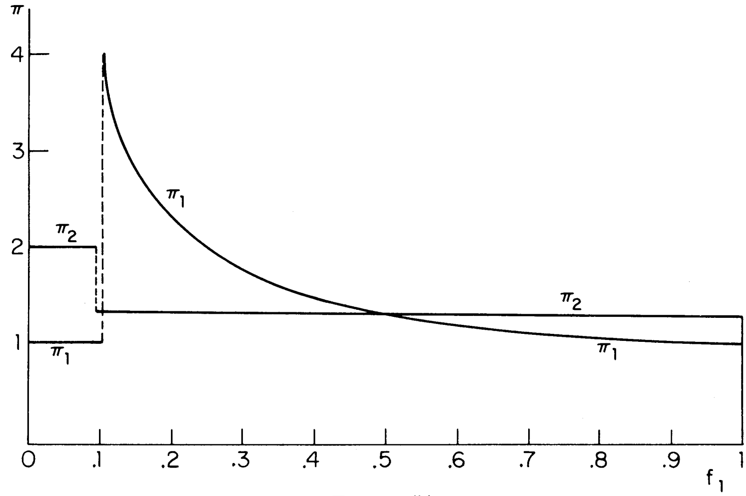

FIGURE IA
Profits of Two Groups of Sellers,
Two Groups of Buyers, Two Search Costs

tween the distribution of search costs and the equilibrium in the market. To fix ideas, we shall consider two simple examples.

Example 1. Let the two prices be 1 and 2. Assume that there are two groups of buyers. The individuals in the first group comprise one-third of the buyers and have a search cost $s_1 = 0.1$. By (3), they will search until they find the lower price if $f_1 > 0.1$. At $f_1 = 0.1$, they are indifferent between searching and not searching; by convention we assume that they do not search. For $f_1 < 0.1$, they accept their first price. Assume that search cost for the second group, the remaining two-thirds of the buyers, is greater than 1 so that they never search. Then $q_1 = 1/3$ for $f_1 > 0.1$, while $q_1 = 0$ for $f_1 < 0.1$. Finally, assume that all sellers have a marginal cost of zero. Then, the standardized expected returns are, by (4),

$$\pi_1 = 1 \quad \text{and} \quad \pi_2 = 2 \quad \text{if } f_1 \leq 0.1;$$
$$\pi_1 = 1/(3f_1) + 2/3 \quad \text{and} \quad \pi_2 = 4/3 \quad \text{if } f_1 > 0.1.$$

Returns of sellers in the two groups, as a function of f_1, are graphed in Figure IA. We see that there is a "noncorner" stable equilibrium at $f_1 = 0.5$. At this point, the return of all sellers is $4/3$, and no seller has an incentive to change his selling price. Stability is assured because the marginal seller at price 2 who switches to price 1 will lower the profits of group 1, inducing him to switch back to price 2, seeing that profits are now higher at that price. An analogous argument can be made for a seller at price 1 who switches to 2. Note that also at $f_1 =$

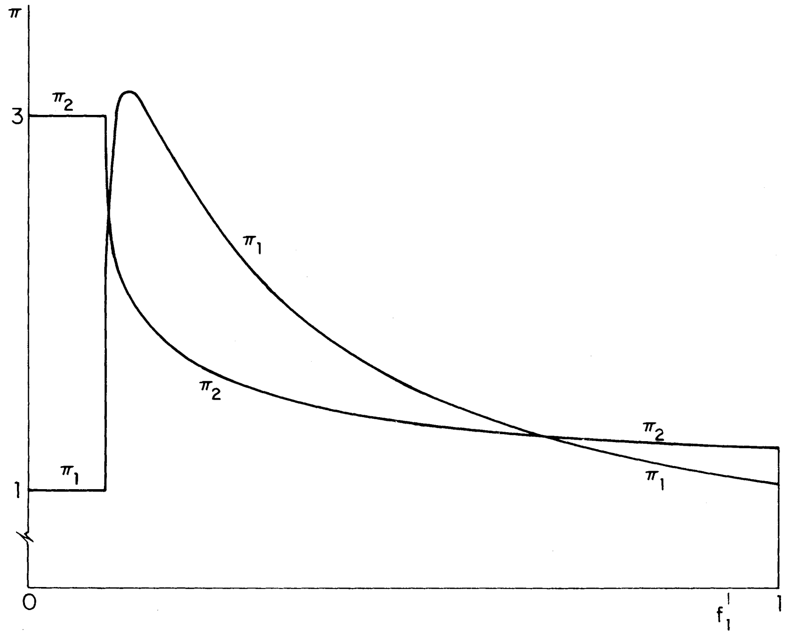

FIGURE IB
Profits of Two Groups of Sellers
Search Costs Distributed Gamma
$c = 1, \alpha = 0.25, \beta = 2, s_0 = 0.2$

0 there is a stable equilibrium with everyone selling at the higher price. At $f_1 = 0.1$, the relative profits of the two groups switch, but there is no equilibrium in the usual sense. However, if exactly $3/11$ of the buyers who are indifferent between searching and not searching should decide to search, then π_1 would equal π_2, and an unstable equilibrium would be achieved. There is no equilibrium at $f_1 = 1$ because a marginal seller at price 1 can increase his profits by switching to price 2.

Example 2. Assume that buyers' search costs follow a continuous gamma distribution above a minimum s_0. By (3) again, a buyer will search until he finds the lower price if $s < (x_2 - x_1)f_1$, that is if $f_1 > s/(x_2 - x_1)$. Thus, no one will search if $f_1 < s_0/(x_2 - x_1)$. For $f_1 > s_0/(x_2 - x_1)$, those with search costs less than $(x_2 - x_1)f_1$, will search, and accordingly the proportion of buyers who search is given by

$$(5)\quad q_1 = \int_{s_0}^{(x_2-x_1)f_1} \frac{1}{\beta^{\alpha}\Gamma(\alpha)} (s - s_0)^{\alpha-1} \exp\left(-(s - s_0)/\beta\right) ds,$$

where α and β are parameters of the gamma distribution. Assume that all sellers have constant marginal cost c. Then their standardized expected returns are

$$\pi_1 = x_1 - c \quad \text{and} \quad \pi_2 = x_2 - c \quad \text{if } f_1 \leq s_0/(x_2 - x_1);$$

and are given by (4) and (5) if $f_1 > s_0/(x_2 - x_1)$.

An example for $x_1 = 2$, $x_2 = 4$, $c = 1$, $\alpha = 0.25$, $\beta = 2$, $s_0 = 0.2$ is graphed in Figure IB. It exhibits one stable and one unstable internal equilibrium, at the right and left intersections, respectively. There is also a stable equilibrium at $f_1 = 0$, but no equilibrium at $f_1 = 1$.[6]

Case of a Finite Number of Sellers

Suppose that the number of sellers is finite. It is now possible that all at the monopoly price is not an equilibrium and that the only equilibrium has multiple prices, where profits may or may not be equal. In Example 1, for instance, if there are five sellers, the only equilibrium has two sellers at price 1 making $\pi_1 = 3/2$ and three sellers at price 2 making $\pi_2 = 4/3$. If there are six sellers, the only equilibrium has three at each price, all making $4/3$. The need to share markets among a finite number of sellers is what supports unequal profits. The possibility that a single seller can influence buyer behavior is what makes defection profitable if all are at the monopoly price.

If there are just two permissible prices, then an equilibrium always exists. Specifically, i out of N sellers at the lower price is an equilibrium for the largest i such that $\pi(x_1|i/N) \geq \pi(x_2|(i-1)/N)$. If no such i exists, then $\pi(x_1|1/N) < \pi(x_2|0)$, in which case all at x_2 is an equilibrium.

If there are three or more permissible prices, there may be no equilibrium, even when buyers have identical search costs.

Example 3. The permissible prices are 3, 4, and 5; there are just two sellers, with cost 0; all buyers have search cost 0.75 and hence search when the price difference is 2 but not otherwise. If both sellers are at 5, either would gain by cutting his price to 3 and taking all of the business. With sellers at 3 and 5, the seller at 5 would optimize by lowering his price to 4. If the sellers are at 3 and 4 or at 4 and 5, the lower-price seller can increase profits by raising his price, since he loses no business. Similarly, if both sellers are at 3 or both at 4, the only remaining possibilities, either seller can raise his price and increase his profits. From any configuration at least one seller has an incentive to move; there is no equilibrium.

III. Equilibrium in Models with Learning

We turn now to models in which buyers do not actually know the distribution of prices, but have prior beliefs about it. Finding a buyer's optimal strategy is a standard problem, but difficult, except in special situations when a myopic strategy is optimal. Part A discusses this briefly, and provides an example. Part B shows how to compute the sellers' profits for any distributions of buyers' strategies and sellers' prices. Part C proves that an equilibrium distribution of prices exists. Part D gives an example. Part E gives an algorithm for finding an equilibrium.

A. Optimal Buyer Behavior with Learning

A coherent, rational (Bayesian) buyer employs a prior distribution over the possible distributions of sellers' prices, which he updates as he receives price quotations. By suitable integration over his updated prior, he can obtain his forecast distribution for the next observation, or all future observations. Once the buyer decides to stop searching, he will, of course, buy at the lowest price to date if it is below his reservation price, and leave the market without buying otherwise. The problem is to determine when he should stop and when he should continue searching. This is an example of a problem of Bayesian sequential sampling.[7] Unfortunately, the solution is generally difficult to compute.

A simple strategy called "myopic" is to look only one stage ahead: after observing prices $X_1, \ldots, X_n$, defining $\hat{X}_n$ as the minimum of $X_1, \ldots, X_n$ and the reservation price, compute the expected one-stage price reduction $E\{\hat{X}_n - \hat{X}_{n+1} | X_1, \ldots, X_n\}$ using the forecast distribution of X_{n+1}, and the relation $\hat{X}_{n+1} = \min(\hat{X}_n, X_{n+1})$; if the expected one-stage price reduction exceeds the search cost s, continue searching; otherwise stop. A sufficient condition for myopia to be optimum is that once myopia says to stop, if a further search is made, it is sure to say stop again, whatever price is found (and hence the same holds forever). Specifically,

THEOREM. Suppose that, for all n and all $X_1, \ldots, X_n$, if

$$E\{\hat{X}_{n-1} - \hat{X}_n | X_1, \ldots, X_{n-1}\} \leq s,$$

then

$$E\{\hat{X}_n - \hat{X}_{n+1} | X_1, \ldots, X_n\} \leq s.$$

Then myopia is optimum.

Proof. After myopia says to stop, every time you continue, you lose expected value. Hence no stopping rule can have positive expected gain. Technically, $\hat{X}_n - \hat{X}_{n+1} - ns$, the gain forms a lower semimartingale (submartingale) once myopia says to stop.[8]

COROLLARY. The foregoing sufficient condition holds, and hence myopia is optimum, if further observations never increase the expected one-stage price reduction, i.e., for all n and all $X_1, \ldots, X_n$,

$$E\{\hat{X}_n - \hat{X}_{n+1} | X_1, \ldots, X_n\} \leq E\{\hat{X}_{n-1} - \hat{X}_n | X_1, \ldots, X_{n-1}\}.$$

With suitable definitions of the gain function, this theorem and corollary obviously apply to rather general sequential stopping problems, though their conditions may rarely be satisfied. They are satisfied, however, when the buyer knows (or thinks he knows) the price distribution exactly.

If myopia is not optimum, one approach is to define recursively the expected k-stage gain as

$$\begin{aligned} G_k(X_1, \ldots, X_n) \\ = E\{\hat{X}_n - \hat{X}_{n+1} + G_{k-1}(X_1, \ldots, X_{n+1}) | X_1, \ldots, X_n\} - s, \end{aligned}$$

when this is positive, 0 otherwise. If at most k further observations are allowed, this is the optimum expected gain, and it is optimum to continue search when it is positive and stop otherwise. If search is unrestricted, it is clearly worth continuing if it is worth continuing for some k. In this problem, the converse also holds, and the optimum expected gain is lim $G_k(X_1, \ldots, X_n)$ as $k \to \infty$. This approach, though intuitively natural, may be computationally difficult because it requires consideration of all k and because of the unbounded dimension of $X_1, \ldots, X_n$ (the "state" in the terminology of dynamic programming). Unfortunately, no easy approach is available. We can, however, replace $X_1, \ldots, X_n$ by $\hat{X}_n$ and a vector of sufficient statistics if the buyer's model possesses one. The following example illustrates both this and optimal myopia.

Example 4. Suppose that the buyer assumes prices have an exponential density with known "decay rate" but unknown left endpoint, say $\exp(\theta - x)$, $x \geq \theta$ (0 elsewhere). If his prior density is $g(\theta)$, then his posterior density is proportional to $g(\theta) \exp(n\theta)$, $\theta \leq \hat{X}_n$ (0 elsewhere), and[9]

$$E\{\hat{X}_n - \hat{X}_{n+1} | X_1, \ldots, X_n\} = \int_0^{\hat{X}_n} (\hat{X}_n - \theta + \exp(\theta - \hat{X}_n) - 1) g(\theta) \exp(n\theta)\, d\theta \Big/ \int_0^{\hat{X}_n} g(\theta) \exp(n\theta)\, d\theta.$$

If $g(\theta)$ is uniform on the interval $A \leq \theta \leq B$ (0 elsewhere), then

$$E\{\hat{X}_n - \hat{X}_{n+1} | X_1, \ldots, X_n\}$$

$$= \hat{X}_n - B + \frac{n}{n+1} \exp(B - \hat{X}_n) \cdot - \frac{n-1}{n} + \frac{n[\exp(B - \hat{X}_n) - \exp(A - \hat{X}_n)]/(n+1) + A - B}{\exp(n\hat{X} - nA) - 1}$$

for $\hat{X}_n \geq B$; for $\hat{X}_n \leq B$ the same expression applies with B replaced by $\hat{X}_n$. (The buyer is a priori certain that $\hat{X}_n > A$ always.) The myopic strategy would continue as long as this exceeded s, but stop the first time it was less.[10] For $A = 0$, $B = 5.0$, and $s = 0.4$, the cutoff points for successive n are shown below.[11]

Cutoffs for exponential example

n	1	2	3	4	5	10	50	∞
Cutoff	3.41	5.49	5.69	5.79	5.84	5.95	6.03	6.05

B. Sellers' Profits

Suppose that the actual price distribution is discrete and the proportion of sellers at price x is $f(x)$. For convenience, we assume that there are m possible prices $x_1, \ldots, x_m$, with $x_1 < x_2 < \cdots < x_m$, and work with the vector $f = [f_1, \ldots, f_m]$, where $f_i = f(x_i)$ is the proportion at price x_i.

A buyer will buy at price x_j if and only if his reservation price is at least x_j and he observes some sequence of prices $X_1, X_2, \ldots, X_n$ in his stopping region with $\hat{X}_n = x_j$. The probability that he will observe a particular such sequence is $\Pi_{i=1}^{n} f(X_i)$. The probability that he will buy at x_j is therefore

$$P_j = \sum \prod_{i=1}^{n} f(X_i),$$

where the sum (which would be an integral if prices were continuously distributed) is over all sequences $X_1, \ldots, X_n$, for all n, which have minimum x_j and fall in the buyer's stopping region. Furthermore,

$$\prod_{i=1}^{n} f(X_i) = \prod_{i=j}^{m} f_i^{r_i},$$

where r_i is the number of occurrences of x_i in the sequence $X_1, \ldots, X_n$. Therefore, P_j is a polynomial or infinite series (depending on whether the buyer's searches are limited or unlimited) in $f_j, f_{j+1}, \ldots, f_m$ with nonnegative integer coefficients, and the power of f_j is at least 1 in every nonzero term. Note that the distribution of selling prices below x_j is irrelevant to this probability. Intuitively, this is because once a buyer observes a price below x_j he will not buy at x_j, so his probability of buying at x_j is unaffected by what price below x_j he observes and what further prices he observes thereafter.

The probability that a seller at x_j will sell to the buyer just discussed is the foregoing probability P_j divided by the number of sellers at x_j which is f_j times the number of sellers. Hence, the standardized profit π_j for sellers at x_j is $(x_j - c)$ times the average over all buyers of P_j/f_j. Thus, π_j is a polynomial or infinite series in $f_j, f_{j+1}, \ldots, f_m$ with all coefficients nonnegative. The coefficient of

$$\prod_{i=j}^{m} f_i^{r_i}$$

is $x_j - c$ times the probability that a randomly chosen buyer has reservation price at least x_j and stopping region including the sequence $X_1, X_2, \ldots, X_n$ summed over all sequences $X_1, X_2, \ldots, X_n$ consisting of $r_j + 1$ occurrences of x_j and r_i occurrences of x_i, $i = j + 1, \ldots, m$, where $n = 1 + \sum_{i=j}^{m} r_i$. For example, the standardized profit for sellers at the highest price x_m is

$$\pi_m = \pi_m(f_m) = (x_m - c) \sum_{n=1}^{\infty} f_m^{n-1} a_n,$$

where a_n is the proportion of buyers who have reservation price at least x_m and who would stop after the nth search if they found the price x_m every time.

C. Proof of the Existence of an Equilibrium When Buyers Have Arbitrary Search Strategies and the Number of Sellers is Infinite

Let buyers have any searching and purchasing strategies whatever. Some may employ irrational strategies, assume information that they do not possess, etc. Let the number of sellers be infinite. We shall show that there will always be an equilibrium distribution of sellers' prices if the possible price quotations are restricted to a finite set. This finiteness involves no real loss in generality, since we can pack the possible prices as tightly as we wish over a range as large as we wish. Thus, we can generate a theoretical equilibrium that might be thought reasonably to approximate conditions in the real world.[12]

Specifically, let the possible prices be $x_1, \ldots, x_m$, let f_i be the proportion of sellers at price x_i, and let π_i be the standardized profits of a seller at price x_i. Then π_i is a function of $f_1, \ldots, f_m$, say $\pi_i = \pi_i(f)$, where $f = [f_1, \ldots, f_m]$. We showed above that π_i is a polynomial or infinite series in only those f_i pertaining to prices of x_i and higher. The proof below uses only the continuity of the π_i, however. Consequently, the result clearly holds under much more general conditions than we have set up here. In particular, it applies when buyers may leave the market without buying (already allowed), when they may want more than one unit, and when they can only purchase at the last price found.

THEOREM. Provided that each π_i is a continuous function of f, there exist $f_1, \ldots, f_m$, with

$$f_i \geq 0, \ \sum_{i=1}^{m} f_i = 1,$$

such that $\pi_i(f) = \max_j \pi_j(f)$ for every i with $f_i > 0$. This is by definition an equilibrium distribution.

Proof. According to a lemma of Knaster, Kuratowski, and Mazurkiewicz [Scarf, 1973], if $A_1, \ldots, A_m$ are closed subsets of the simplex,

$$\left\{f \colon f_i \geq 0, \ \sum_{i=1}^{m} f_i = 1\right\},$$

if each A_i contains the face $\{f \colon f_i = 0\}$, and if the union of the A_i is the whole simplex, then the A_i have a point in common. To apply this lemma, let

$$A_i = \{f \colon \pi_i(f) = \max_j \pi_j(f) \text{ or } f_i = 0\}.$$

The A_i are easily seen to be closed by the continuity of the π_j; each contains the face $\{f \colon f_i = 0\}$ by definition, and their union is the whole simplex, since for every f, some j maximizes $\pi_j(f)$, and f belongs at least to that A_j. Therefore, by the lemma there exists an f belonging to every A_i. This f satisfies the conclusion of the theorem.

D. An Example

Knowing that an equilibrium will exist, we might like to see one with more than one price. Consider a situation in which there are three possible prices, $x_1 < x_2 < x_3$, and three types of buyers. Type 1 buyers will search until they find a price of x_1; type 2 buyers will search until they find a price of x_2 or lower; type 3 buyers will accept the first price they are quoted. Represent the proportion of buyers of each type as q_1, q_2, and q_3. Let the profits per sale (i.e., price minus marginal cost) for sellers at the three prices be r_1, r_2, and r_3, respectively. The standardized profits for the sellers will then be

$$\pi_1 = r_1\left(\frac{q_1}{f_1} + \frac{q_2}{f_1 + f_2} + q_3\right),$$

(6)
$$\pi_2 = r_2\left(\frac{q_2}{f_1 + f_2} + q_3\right),$$

and

$$\pi_3 = r_3 q_3.$$

For a given q_1, q_2, q_3, we wish to find f_1, f_2, f_3 such that $\pi_i = \max_j \pi_j$ whenever $f_i > 0$. For an internal equilibrium, one where all $f_i > 0$, we must have $\pi_1 = \pi_2 = \pi_3$, and therefore

$$f_1 = \frac{r_1 r_2 q_1}{r_3(r_2 - r_1)q_3},$$

(7)
$$f_2 = \frac{r_2 q_2}{(r_3 - r_2)q_3} - f_1,$$

and

$$f_3 = 1 - (f_1 + f_2).$$

To insure that all $f_i > 0$, we must also have

(8)
$$\frac{r_1 r_2 q_1}{r_3(r_2 - r_1)} < \frac{r_2 q_2}{r_3 - r_2} < q_3.$$

If (8) is satisfied, then the f_i given by (7) provide an internal equilibrium. It can be seen from (6) that the equilibrium is stable and unique. Similarly, relationships can be derived to describe equilibria at which one or more of the $f_i = 0$.

By way of numerical illustration, let $q_1 = 0.1$, $q_2 = 0.2$, and $q_3 = 0.7$. Assume that marginal cost is zero and that prices are 1, 2, and 3, so that $r_1 = 1$, $r_2 = 2$, and $r_3 = 3$. There will then be a unique equilibrium that is stable, with $f_1 = 2/21$, $f_2 = 10/21$, and $f_3 = 9/21$. The standardized profits per seller at the equilibrium will be 2.1. It is worth noting as a sidelight that the search strategies for the three groups of buyers would be optimal, assuming full knowledge of the distribution of sellers, if they had search costs less than 2/21, between 2/21 and 14/21, and greater than 14/21, respectively. (This assumes that their expected consumer surplus makes it worthwhile for them to enter the market initially.)

Figure II reveals the structure of the problem. It shows which group of sellers receives the highest standardized profits, as a function

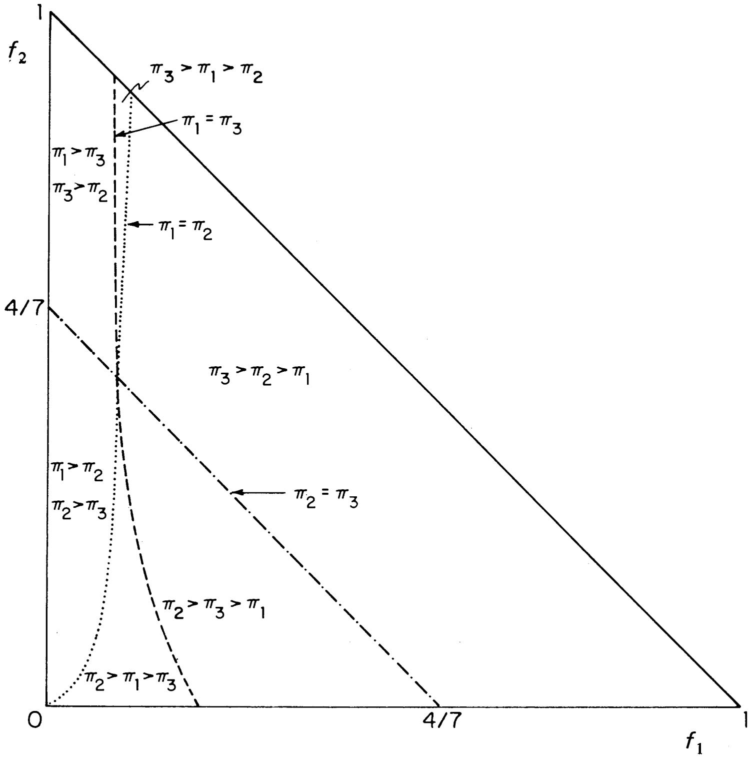

FIGURE II
An Equilibrium with Three
Prices and Three Types of Buyers

of the proportions of the three groups of sellers. The broken lines in the figure show where the standardized profits to two groups of sellers will be equal. It can be seen that only at point E are the profits to the three groups of sellers equal. That there are no corner equilibria can be determined in the figure by seeing that some price not being quoted always has a higher level of standardized profits than some quoted price at any side or corner point. Thus, the only equilibrium has multiple prices.

E. An Algorithm for Finding an Equilibrium

Having proved the existence of an equilibrium, we should like to have an algorithm to find one. One successful approach is to start with the two highest potential prices for sellers. For any total proportion, $f_m + f_{m-1}$, at these two prices, compute all possible equilibrium combinations of f_m and f_{m-1}, and the standardized profits associated with each. Such calculations are possible, since standardized profits to sellers at these prices depend only on the total proportion of sellers at lower prices, not on their distribution. Knowing the possible equilibrium profits as a function of the total proportion $f_m + f_{m-1}$, we can now consider the next lower possible price and compute all possible equilibrium combinations of $f_m + f_{m-1}$ and f_{m-2} and the standardized profits associated with each of them. This procedure for cascading downward can be continued to compute an equilibrium for any arbitrarily large number of prices.[13]

IV. Empirically Observed Distributions of Prices Quoted by Different Sellers

The attempt by one of the authors to find the "best" price for sand revealed surprisingly large differences among the prices quoted by sellers in the Boston area. To determine whether or not this phenomenon was peculiar to a single product, or some kinds of products but not others, for example, low priced items but not expensive ones, we undertook a more systematic investigation of products traded in "competitive" markets. A sample of fifty products was compiled by selecting pages at random from the Yellow Pages of the Boston telephone directory; the first product or service line starting on that page was included in the sample. Some, such as funeral director or airport construction, were eliminated for the collection of price information could be seen to be, or proved to be, difficult. Within each of the remaining thirty-nine categories, a relatively standardized product was selected, for example, a particular brand and model number.[14] Then every listed seller, their number ranged from four to twenty-two, was called by telephone and asked for the price at which he would be "willing to sell" the product.

Summary statistics for the prices of each product class are presented in Table I. For most of the products, including expensive products, there were substantial percentage differences among quoted prices. For eighteen of the thirty-nine products, the highest price was over twice the lowest. To provide another description of the price distributions, we fitted the prices for each product to a gamma distribution by the maximum likelihood method. Seven of the shape parameter estimates are less than one, indicating a distribution more skewed than the exponential, with a mode at the left of the range. Fourteen estimates were between one and nine, indicating distributions that are still substantially skewed to the right, but have interior modes. Eighteen estimates were greater than nine, indicating nearly normal skewness.[15]

We found a large positive relationship between standard deviation and mean price. A reasonable representation of the observed relationship is provided by the least squares fit of the logarithm of the estimated standard deviations to the estimated means. It is given by

$$\ln SD = -1.517 + \underset{(0.059)}{0.892} \ln \mu, \quad R^2 = 0.870.$$

By exponentiation it is seen that a doubling of the estimated mean price is associated with approximately an 86 percent increase ($2^{0.892} \approx 1.86$) in the estimated standard deviation of prices. This 86 percent figure might seem puzzlingly high, for unless search costs increased dramatically with the price of the product, the expected gains from searching would lead to ratios between standard deviation and price that declined rapidly with mean price. The explanation may lie with the infrequent purchase of expensive products, which reduces the incentive of a buyer to search. Less searching, in turn, allows greater variability among prices.[16]

V. Qualifications, Implications, and Conclusions

The surprisingly large difference among prices that we observed can be explained solely by positive search costs or significantly divergent beliefs among buyers or both. Other factors, however, assuredly played a role as well. Although conscientious attempts were made to select standardized products, there undoubtedly were quality differentials our phone inquiries did not detect, such as seller location, credit politics, or accompanying services. (It was surprising that none of the providers queried made any claims in this regard.) Sellers may have quoted different prices because they confronted different situations with regard to demand conditions, costs, modes of operation, prospects for tie-in sales, etc., or because they competed in different geographic submarkets. Many alternative hypotheses about buyer behavior could help explain the observed price distributions. Because the model does not consider such aspects as advertising or other means for sellers to send price information, it may be incomplete or unrepresentative for some products.

The observation that quoted prices vary significantly from seller to seller in markets we thought to be competitive may have implications both as to the concepts we employ to describe market behavior and our attitudes toward phenomena affecting this performance, such as advertising and government regulation. These are matters for future investigation. Here, we shall merely touch on the subject of ef-

TABLE I

DIFFERENCES AMONG PRICES FOR 39 PRODUCTS[a]

Product	Mean price ($'s)	Standard deviation ($'s)	Minimum	Maximum ÷ minimum (×100)	Number of firms	Value of[b] shape parameter
Bicycle	144.77	6.34	135.00	111	7	24.3
Boat	602.87	104.89	425.00	197	15	9.6
Aquarium	17.38	3.66	12.95	216	12	3.8
Service stations	13.69	1.46	11.00	148	15	16.9
Lumber	10.91	1.04	9.50	137	14	15.0
Cameras	329.12	29.78	260.00	142	15	28.8
Pet washing	15.63	2.00	12.00	150	8	16.7
Liquor stores	7.47	0.65	6.50	123	11	16.8
Office furniture	109.45	9.03	89.00	135	15	22.0
Paint	8.19	0.58	7.15	139	22	26.7
Carnations	0.33	0.14	0.20	250	7	0.5
Mufflers	28.58	4.97	22.95	157	19	3.5
Chlorine	1.79	0.79	0.60	500	7	2.6
Horoscope	30.00	16.83	10.00	500	4	0.6
Vocal instruction	14.79	5.17	10.00	265	12	0.5
Teeth cleaned	16.85	3.76	12.00	208	13	2.8
Printers envelopes	39.54	9.12	24.00	250	14	7.2
Air conditioners	16.74	4.23	10.00	250	14	5.5
Canvas goods	40.46	28.87	14.89	638	14	0.5
Diamond appraisals	7.71	2.01	5.00	200	12	6.2
Peanuts	0.53	0.17	0.43	209	8	0.6
Auto tune-up	39.57	7.27	30.00	183	15	7.0

Styling brush	4.33	2.01	1.50	667	12	3.9
Developer	1.44	0.14	1.20	133	15	18.0
Clean 2-piece suit	2.08	0.35	1.18	212	20	10.4
Flying instructions	29.67	4.15	20.00	170	9	15.7
Watch cleaned	15.78	5.98	8.00	350	12	2.3
Turntable	68.55	12.37	58.00	130	9	0.9
Stationery	92.82	5.62	83.50	125	15	20.7
Towing	14.04	3.57	8.00	250	10	5.2
Repair clarinet	44.28	13.78	30.00	233	8	0.9
Microwave oven	451.50	26.15	425.00	116	5	15.9
Rebuilt alternator	38.02	2.96	34.95	129	14	2.4
Fuel oil	0.38	0.11	0.36	111	15	18.7
Skates	42.83	3.53	36.00	139	15	18.1
Electrodes	30.71	9.55	19.60	286	12	1.2
Board poodle	4.00	0.68	3.00	183	13	8.4
Concrete	38.50	2.01	25.00	126	11	26.9
Calculator	123.44	12.27	102.00	142	12	16.5

a. Product and further description: 1. Bicycles: Raleigh Grand Prix, 10-speed. 2. Boats: Johnson 9.9 hp outboard motor, 15″ shaft, manual start. 3. Aquariums: Fishtanks, 20 gal. 4. Service stations: Lube job, grease, oil & filter change, '67 Camaro 327 V-8. 5. Lumber: Plywood, Fir AC $\frac{1}{2}$″ × 48″ × 96″. 6. Cameras: Konica, Autoreflex T(3), f1.7 lens, chrome finish, split img. foc. 7. Pet washing: French cut and shampoo, standard poodle. 8. Liquor stores: Beer, Budweiser, case (12 quart bottles). 9.Office furniture: 4-drawer letter size, HON #214, black. 10. Paint: Acrylic flat latex paint, Benjamin Moore, Regal Wall Satin Super White. 11. Carnations: White carnations, 200 loose, undelivered. 12. Mufflers: 1967 Camaro, single exhaust. 13. Chlorine: Liquid, 15% strength, a few gallons, per gallon price. 14. Horoscope: Basic charting, includes approximately a one-hour discussion. 15. Vocal music instruction: Length of classes vary, price/hour. 16. Teeth cleaned: Dental office. 17. Printers envelopes: 2000 #10 size w/4 lines copy, flat printed, 24 lb., white-wove, sulfate p. 18. Air conditioners: Repair. 19. Canvas goods: Cover for pickup, 70″ × 8′2″, 3 materials, glove snaps. 20. Diamond appraisal: 3 one-karat diamond rings, price/item. 21. Peanuts: Roasted in the hull, 100#, price/pound. 22. Auto tune-up: '67 Camaro, V-8, without air conditioning, parts, labor, plugs, timing, condenser, points. 23. Dennman styling brush: Trow. 24. Dektol developer: 1 lb. 2 oz. 25. Clean 2-piece suit: Dry cleaned. 26. Flying instructions: Aircraft with instructor, rates/hour. 27. Watch cleaned: Clinton 25 jewel, self-winding, with day-date. 28. Turntable: BSR 2310X. 29. Stationery: Southworth 33C, permanent record, 10 reams. 30. Towing: Initial charge, daytime, round trip of 5 miles. 31. Repair clarinet: Repad, Selmer clarinet. 32. Microwave oven: RK4D, Amana. 33. Rebuilt alternator: 1967 Camaro, 327 cu. in., V-8, with air conditioning. 34. Fuel oil: 500 gallons, price/gallon. 35. Skates: Ice, "Black Panther." 36. Electrodes: E6013, $\frac{1}{8}$″, 50 lbs. 37. Board poodle: Standard, approximately 50 lbs., price/day. 38. Concrete: 3000 lbs. (type or grade), price/cubic yard. 39. Calculator: Texas Instruments SR-50.

b. Maximum likelihood estimate of the shape parameter of the gamma distribution, α, to which the observed prices were fit. The gamma distribution was chosen because it is reasonably flexible, having as special cases the exponential ($\alpha = 1$) and the normal ($\alpha \to \infty$) distributions. Estimates of α less than one indicate exponential-like distributions. Values between one and nine indicate distributions that are substantially skewed to the right. Distributions with values greater than nine are essentially normal looking. See Pratt, Wise, and Zeckhauser [1975].

ficiency, present one curiosity relating to economic theory, and finish with some brief conclusions.

A. *Efficiency and Search Costs*

Search costs have complex efficiency implications. Zero search costs obviously lead to traditional competitive outcomes and maximum efficiency. No general conclusion should be extrapolated from this result. A simple example shows that efficiency may actually fall as per-trial search costs decrease.

Example 5. Permissible prices are 1 and 2. The ten sellers have zero marginal costs; the ten buyers have unit search costs of 0.6. There is a stable self-justifying equilibrium with all sellers at price 2. (There is another stable equilibrium at price 1, but sellers would have no incentive individually or collectively to migrate from 2 to 1.) Each buyer will get precisely one price quotation. The search cost of one of the buyers now falls to 0.09. It will be worthwhile for one of the sellers to lower his price to 1, reaping 1.9 units of the standardized profit, as opposed to 1.8 for the sellers at price 2. This is an equilibrium, in the sense that no seller can benefit by changing his price. (Since sellers are assumed to be indivisible, profits are not fully equalized.) It will take the low search cost buyer ten times on average to find the price 1 seller. Expected search costs, the sole indicator of efficiency if quantities purchased do not vary, will have increased from 0.6 at the former equilibrium to 0.9 at the new one for this buyer, with no change for the others.

B. *The Howard Johnson's and Cheap Motel Phenomena*

Search costs invalidate many general propositions about competitive models. For example, a low-price seller may prefer to have a high-price competitor lower his price. Though this will split off some of the low-price market, it will also expand that market by inducing buyers to search for low-priced sellers. (This conjunction of effects allows for multiple stable equilibria as illustrated by Example 1 or unstable equilibria as seen in Example 2.)

Example 6. Say that there are four suppliers selling at 2 and six selling at 1, 100 buyers with search costs of 0.1 and 100 buyers with search costs of 0.75. The world is in a self-justifying equilibrium in the sense that no seller can switch by himself and do better. A single seller at price 2 dropping his price to 1 would not change search behavior. But if two sellers at 2 dropped, there would then be an 80 percent chance of finding a price of 1 on a further trial. People with search costs of 0.75 would find it worthwhile to search, the returns to the 1's would go up. The additional sellers at 2, following their profit-seeking noses, would cascade into the ranks of those selling at 1. A stable equilibrium would be reached with everyone at 1. During the transition, some 1's were making extra profits; now both they and the original 2's will have lost 25 percent of the original profits. (We leave aside such considerations as sellers exiting or demand changing as prices fall.)

Real-world entrepreneurs are well aware of this phenomenon. The orange roofs of Howard Johnson's have long peppered the country. A traditional explanation of the importance of a nationwide chain such as this one would be that it can provide some assurance of quality at a reasonable price to travelers not familiar with a particular restaurant. The search model suggests an additional, possibly complementary, explanation. If there were just one or two such restaurants in the country, motorists would not pass up inferior restaurants in the hope of finding a Howard Johnson. But with the orange-roofed restaurants scattered liberally about, it would be worthwhile to drive on by, knowing that for not too great a search cost a better deal could be found.

Many potential significant changes in the economic landscape may fail to occur because of the substantial numbers needed to get a new price or new quality going, i.e., to get to a new equilibrium. Consider the plight of the motel operator who observes that his neighboring motels are two-thirds empty charging their \$12 prices. Since marginal cost between occupied and unoccupied rooms is low, why not charge \$8 and have a full motel? The answer is straightforward. If this will be the only \$8 motel along a great road, motorists may stop resignedly at more expensive ones and demand will not increase. The \$8 motels can only make it when there are enough of them to repay search for one.

C. Conclusions

The data presented here show that relatively standardized products vary substantially in price from seller to seller. The persistence of such situations means that there are deviations from the competitive model. Nonnegligible search costs for consumers would be one such deviation. Their possible role as a contributor to price differences is explored in this analysis.

For the case where consumers know the distribution of sellers' prices, there may nevertheless be an equilibrium with some seller-to-seller price differences. Indeed, such an equilibrium may be unique. Thus, price differences may be necessary for equilibrium.

A more realistic case has buyers less than fully informed about the distribution of sellers' prices. Buyers formulate sequential search strategies, gathering information as they go. For this case of learning, we show, there will always be an equilibrium distribution of prices, and again it may display price differences even when it is unique.

HARVARD UNIVERSITY

NOTES

1. Bryn Zeckhauser indirectly provided the original impetus for this research. Her school playground required sand. Mac Gaither dialed diligently in collecting information on our sample of standardized products. Norman Bookstein and Thomas Shemo provided computer assistance. Truman Bewley, Elon Kohlberg, Donald Rosenfield, and Roy Shapiro gave us mathematical advice. Tradin' Sam Peltzman gave us the benefit of his extensive purchasing experience. This research was supported in part by NSF grants to Harvard: SOC76-15546 (Pratt), SOC76-81989 (Wise), and SOC77-16602 (Zeckhauser) and to MIT SOC75-14258 (Zeckhauser). A preliminary version of this paper was presented at the December 1975 AEA Convention.

2. If individuals search for lower prices, of course, the dispersion of purchase prices will be less than the dispersion of quoted prices. There is also the possibility that quoted prices are negotiable and would tighten down to the competitive price if bargaining were conducted.

3. The self-justifying concept can be extended to include any situation in which buyers respond to the distribution of sellers' prices in some well-defined manner, whether or not the responses are optimal, and likewise for well-defined sellers' responses, whether or not in equilibrium.

4. Equality holds in (1) at a unique y'. Though ordinarily not one of the permissible prices, it is the price at which the buyer would be indifferent between buying and continuing optimally. It equals the optimal expected cost of search plus price paid. See Chow and Robbins [1961] and DeGroot [1970, especially Section 13.5]. Our cutoff price is the highest price with positive probability not exceeding the indifference price.

5. This requires some argument, since a "small" change in f may produce an appreciable change in q and hence in π, for instance when an appreciable fraction of buyers are indifferent between continuing and stopping at a price y, and thus would be induced to continue instead of stopping by a "small" increase in the probability of a price below y. (If they continue when indifferent, a similar discontinuity occurs in the opposite direction.) However, if a price reduction never decreases the probability of a buyer's purchasing at the reduced price, and if $\pi(x)$ at the highest (or any other) quoted price is smaller than at some lower price, quoted or not, then some seller can gain by reducing price. Furthermore, if a price increase to the highest quoted price never decreases the probability of a buyer's purchasing at that price, and if $\pi(x)$ at the highest quoted price is greater than at some other quoted price, then some seller can gain by increasing price. Both conditions on buyers' responses are natural and hold here. In addition, $\pi(x) = 0$ above the highest quoted price. Hence all quoted prices maximize $\pi(x)$ in equilibrium. It is possible for all quoted prices to maximize $\pi(x)$ without equilibrium in our earlier sense; we shall refer to this as an unstable equilibrium.

6. For further examples, see Pratt, Wise, and Zeckhauser [1975].

7. See Blackwell and Girshick [1954, Chs. 9 and 10], DeGroot [1970, Ch. 13], and Chow, Robbins, and Siegmund [1971]. Kohn and Shavell [1974] obtain some properties of the solution by a somewhat different approach, but do not provide a method of finding the exact solution.

8. See Doob [1954, pp. 294, 300].

9. The discussion here applies only after a price below the reservation price has been found. We shall not discuss stopping without buying.

10. In conjunction with this paper, Rosenfield and Shapiro [1971] have developed general conditions for the optimality of myopia and show optimality for this case and for an exponential prior as well.

11. If the optimal cutoff point for some n is below B, as it is for $n = 1$ here, then it is the same for all larger values of B, because B is irrelevant once $\hat{X}_n$ is below B, and the gain from continuation is greater for $\hat{X}_n$ above B than below B. As $n \to \infty$, the difference between the optimal cutoff point and B approaches a positive constant, independent of the value of B.

12. We should contrast this situation with that considered in Section II. There the distribution of sellers was assumed to be known and hence to affect buyers' strategies. Here buyers are assumed to formulate strategies independent of the distribution of sellers. The actual distribution of sellers will affect what buyers learn and the prices at which they will ultimately buy, but not their strategies.

13. Standardized profits may be a multivalued function of the proportion of sellers above any price, though we conjecture, in accord with our observations, that the graph relating standardized profits to this proportion will always be connected.

14. If the category happened to be one including no readily identifiable standardized product, it was eliminated from further investigation. Voice instruction, for example, was selected by the page drawing, but lessons were not priced.

15. Considerable further detail on our empirical results and estimation methods is contained in Pratt, Wise, and Zeckhauser [1975].

16. In fact, Sam Peltzman, who commented on this paper at the December 1975 AEA meetings, selected all of our products that were "unambiguous" consumer goods—32 products in all—and distinguished "frequently" purchased goods from the others. Using a dummy variable D to represent frequently purchased goods, he obtained the following result:

$$\ln SD = -1.161 + \underset{(0.084)}{0.836} \ln \mu - \underset{(0.369)}{1.015}\, D, \qquad R^2 = 0.870.$$

Thus, the standard deviation of a frequently purchased good is estimated to be about 36 percent (exp $(-1.015) \approx 0.36$) of that of an infrequently purchased good with the same mean.

Less frequent purchase also implies that buyers will know less about the distribution of prices. For any individual buyer this would make it more worthwhile to gather information through search, a factor tending to reduce price variability. Cutting in the opposite direction, with less frequent purchase, buyers will have less incentive to search hard for a low price. So too, with less known there will be more divergent beliefs among buyers, which allows for greater variability in quoted prices.

References

Arrow, K., and M. Rothschild, unpublished manuscript. 1975.

Blackwell, D., and M. A. Girshick, *Theory of Games and Statistical Decisions* (New York: Wiley, 1954), Chs. 9–10.

Chow, Y. S., and H. Robbins, "A Martingale System Theorem and Applications," *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability* (Berkeley: University of California Press, 1961), Vol. I, 93–104.

Chow, Y. S., H. Robbins, and D. Siegmund, *Great Expectations: The Theory of Optimal Stopping* (Boston: Houghton Mifflin, 1971).

DeGroot, M. H., *Optimal Statistical Decisions* (New York: McGraw-Hill, 1970), Chs. 12–13.

Doob, J. S., *Stochastic Processes* (New York: Wiley, 1953), Ch. 7.

Kohn, M. G., and S. Shavell, "The Theory of Search," *Journal of Economic Theory*, IX, No. 2 (Oct. 1974), 93–123.

Pratt, J., D. Wise, and R. Zeckhauser, "Price Variations in Almost Competitive Markets," Discussion Paper No. 37D, Cambridge: Kennedy School of Government, December, 1975.

Rosenfield, D. B., and R. D. Shapiro, "Optimal Price Search with Bayesian Extensions," Harvard Business School, Division of Research: HBS-77-23, 1977.

Scarf, H., *The Computation of Economic Equilibria*, Cowles Foundation Monograph 24 (New Haven: Yale University Press, 1973).