OPTIMAL SELLING STRATEGIES: WHEN TO HAGGLE, WHEN TO HOLD FIRM†

By

John Riley and Richard Zeckhauser*

Discussion Paper Number 180
Revised September 1980

†The helpful comments of John McCall, Nicolaus Tideman and Barry Nalebuff and the financial support of the National Science Foundation are gratefully acknowledged. An earlier draft of this paper circulated as a Kennedy School of Government Discussion paper.

*Department of Economics, University of California, Los Angeles, and Kennedy School of Government, Harvard University.

ABSTRACT

A seller encountering risk-neutral buyers one at a time should, if commitments are feasible, quote a single take-it-or-leave it price to each. We demonstrate the validity of this proposition in comparison with any other possible selling strategy, whether there is learning or the distribution of buyer reservation prices is known, whether the buyer population is finite or infinite, whether there is one object for sale or many. If the population is finite or if there is learning, but not otherwise, it may be desirable to recall a refusing buyer for a second price quote after being refused by other buyers.

Though haggling may offer advantages in terms of price discrimination, these gains are more than offset by the losses it generates by encouraging buyers to refuse purchases at high prices. Any information gains haggling may offer can be reproduced at effective zero cost through a proposed mechanism that elicits the reservation prices of refusing buyers.

Sellers should only haggle when the buyers they face are risk averse, or when they can not make a commitment not to haggle--because of the lack of a reputation, for instance, or the occasional nature of a market encounter. Casual observation suggests that, for the most part, real-world sellers behave as if they understood these lessons.

Optimal Selling Strategies: When to Haggle, When to Hold Firm

John Riley and Richard Zeckhauser

How should sellers price their goods? In the bazaar or the agricultural market of a less developed nation, haggling is the norm. In most developed nations, on the other hand, posted non-negotiable prices are employed for most goods, though for a range of goods from autos to real estate there may be considerable flexibility in prices.*

Why do we find posted fixed prices? Should not vendors be willing to cut prices when a sale might otherwise be lost? A variety of explanations for the fixed-price phenomenon might be developed. In a perfectly competitive market, prices reflect marginal cost. A sale at a lower price would obviously not be worthwhile. In complex organizations, such as a modern department store, fixed prices solve problems of coordination. Without fixed prices, each salesman would have to receive extensive and detailed instructions relating to markups on different items, acceptable and unacceptable price cutting, how to judge customers, etc. Problems of collusion between salesmen and buyers might also arise. In any market, fixed prices dramatically reduce information costs -- you know immediately what price you will get in a store -- as well as costs of negotiation. Not surprisingly, many customers prefer to shop at stores with clearly posted prices.

<u>Goods yielding marginal profits</u>. Our focus in this paper is on goods that yield a direct profit to the vendor when they are sold. Such goods include

*Our work in this area has heightened our sensitivity to the question of firmness versus flexibility in prices. We have discovered that flexibility can be found in unexpected places, such as the prices of big-city hotels. As our analysis will show an ideal strategy for vendors who can get away with it is to proclaim inflexibility, but permit it when a sale may otherwise be lost. Antique stores may post prices and suggest on casual inquiry that they are fixed. These prices may be cut, or extras thrown in, for sophisticated buyers.

products sold on oligopolistic or monopsonistic markets. They also include goods sold on markets that are normally thought of as competitive or near-competitive, but where there is the slight imperfection that there is a cost of offering the good for sale. Such offer costs include the rent, heat and light of the store, and the salesclerk's salary (assuming that he is not 100% occupied); they must be covered through price. Any good whose sales price exceeds its marginal cost we call a marginal profit good.

The central question we shall explore in this paper is why a vendor should ever let a sale of a marginal profit good go by. That is, when he gets a refusal, why should he not cut the price of his good just a bit, though staying above marginal cost, to see if he can't secure the sale? One obvious answer is that once he develops the reputation for price cutting he will find it impossible to make sales at prices above marginal cost. A second answer might be that the cost of making still another offer (including the opportunity cost of his time) might be greater than the expected benefits to be reaped. In some circumstances, it might even pay the vendor to increase the cost he incurs to make an additional offer; it may provide a way to make his previous "final price" more convincing.

I. Formulation

This analysis focuses on the case where vendors can make firm commitments as to pricing strategy. A critical consideration in the model are the four costs the vendor faces: the cost of securing another customer, the cost of making an additional offer to the same customer, the cost (more an opportunity cost than a financial cost) of using up one of a finite number of customers, and the cost of securing another unit of the good to be sold. Weighing all these factors in all their permutations at once would be prohibitively complex. We shall normally assume that there are zero costs in making additional offers to the same customer, and that only one unit of the good to be sold is available (infinite cost for restocking inventory). The model is general, however, and can be readily extended to consider cases where there is a constant cost to secure another item for inventory.

The cost of bringing a customer into the store is assumed to be constant. Discounting is left aside for simplicity, though in a number of circumstances it could be included in the cost of securing a new customer. Sellers seek to maximize expected profits. They are risk neutral.

Purchases are assumed to be sufficiently infrequent that buyers make no attempt to establish reputations with sellers. Instead, their payoffs are very simple. They seek to maximize their expected consumer's surplus, the difference between their reservation price, i.e., the maximum take-it-or-leave-it price they would accept, and the price they pay. They are risk neutral. It costs buyers nothing to stay for another offer.* Finally, the reservation price of a buyer is assumed to reflect both his personal valuations of the object and his expectations about the price of substitute products offered by other sellers.

The seller's problem. A series of buyers will enter the seller's showroom at random. If there is a finite number of potential buyers, the seller knows how many there are. Each buyer will have a reservation value v. Making use of information gained through general experience, including what he learns from observing previous unsuccessful buyers, the seller forms a belief about a buyer's reservation value. These beliefs are described by a cumulative distribution function $F(v)$. Once the selling process starts, the seller has no source of information except the buyers who come through his showroom.

---

*While it may seem odd to ignore the time costs of haggling, the alternative assumption of equal positive time costs for each buyer simply makes haggling less profitable in comparison to selling at a fixed price. Our model does not, however, incorporate the possible incentives to haggle when time costs are known to vary.

Our analysis relies on the following critical assumption:

<u>A seller is able to make a commitment to any contingent strategy he wishes, and he can convey this binding commitment to each buyer</u>.

This assumption is reasonable for a great variety of selling situations, including any store or merchant who can establish a reputation.

In choosing his strategy, the seller compares the way buyers with different reservation prices will respond to each strategy. Given the pairing of a buyer's optimal strategy in response to an announced seller strategy, the probability of sale and the expected price conditional upon sale can both be computed. If a sale is not consummated, a new buyer enters the store and the process continues until a sale is made or until the seller withdraws the item. The situation can be understood as the two-player sequential game diagrammed in Figure 1.

INSERT FIGURE 1 ABOUT HERE

If the seller knows the distribution of buyer reservation prices the game starts at move 1. If there is learning, it starts at move 0 when the distribution of buyers' reservation prices is determined. The seller is not told which distribution is chosen, but he has prior beliefs about the likelihood of different distributions and updates those beliefs as he goes along.*

---

*We might think metaphorically of selecting an urn at random, with each urn containing a different distribution of balls indicating reservation prices. F(v) would be derived by computing a weighted average of the compositions of the different urns, the weights being the likelihoods assigned to the urns by the seller.--Buyers are only concerned with the seller's committed behavior. It does not matter whether they know his prior beliefs.

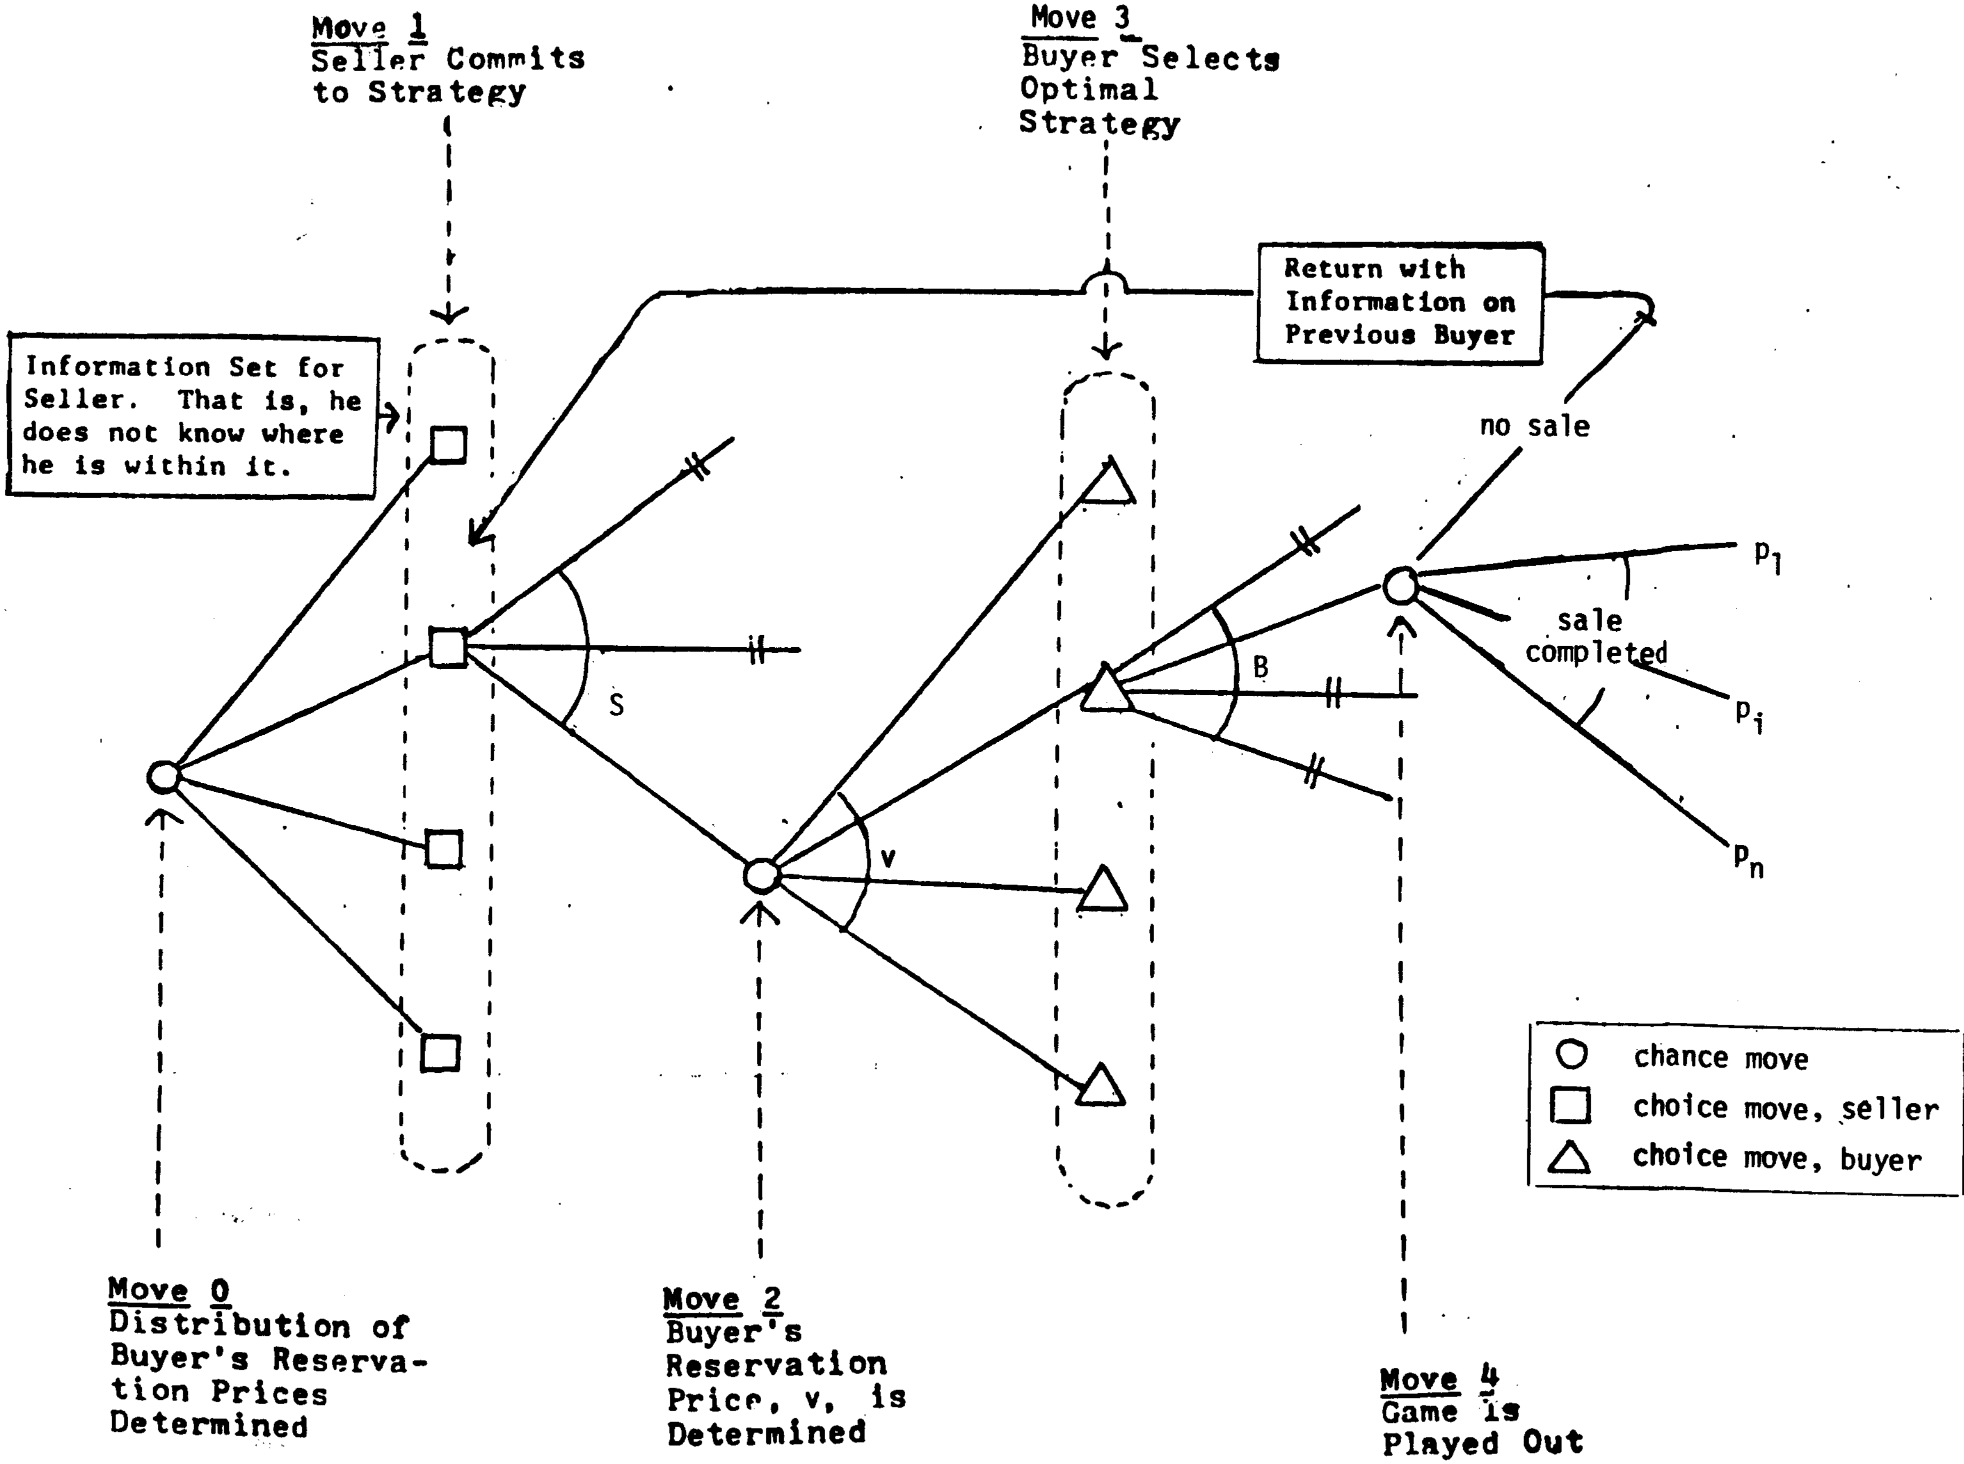

Figure 1: Seller and Buyer Interaction

At move 1 the seller commits himself to a contingent-pricing strategy S. The strategy details how he will behave in any situation that may arise in the future. For example, it might say that if he is turned down at a price of α, there is a 50% chance that he will quote a price of β and a 50% chance that he will request a new buyer. In effect, once a seller commits to S he turns over play to a capable computer, and then displays the computer program to each and every buyer. The program may be sophisticated, and allow him to change his practices as he acquires information. That is, it may commit to the way experience with buyer 1 will affect practices with buyer 2.

Move 2 is the chance move that determines the buyer's reservation price v. The buyer knows the result of this chance move; the seller does not. The seller's committed strategy, of course, may enable him to draw inferences about v from the buyer's behavior, and to act upon those inferences in quoting and accepting offers.*

At move 3, the buyer, knowing both v and S, selects his preferred strategy B. Now that the buyer's and seller's strategies are both determined, the game can be played out at move 4. That is, the bargaining and negotiation process can now be conducted. This aspect of the process might be compared to a poker game between two computers. Move 4 results either in no sale, or in a sale together with a selling price, p.

---

*The ordering of moves 1 and 2 is arbitrary, and they could be reversed. The seller does not learn the result of move 2 before he commits to a strategy, and the distribution of the buyer's reservation price is unaffected by the seller's chosen strategy.

If no sale is made, the process cycles back to move 1. The seller, having updated his information, announces his selling strategy to the next buyer. This may be a purely computational operation. The first time he arrived at move 1, he may already have committed himself to a super-strategy which told how he would behave towards buyer 2 depending on his buyer 1 experience, towards 3 depending on 2 and 1, etc.

The payoffs to the seller and buyer are indicated at the tips of the tree. The buyer will fold his payoff back to the triangle to select his preferred strategy. The seller folds back to his choice move, the square, to select the strategy to which he should commit. In computing his expected payoff for any strategy, the seller must also be concerned with the expected value of the game should there be a new play.

Possible bargaining formats. Our decision tree formulation allows for a rich array of possible bargaining formats; it captures all possibilities where, as would often seem reasonable, the seller as the more continuing and permanent participant can commit himself first.

Either the seller or the buyer or both can make price offers under our formulation. In traditional bargaining, the sellers and buyers alternate in making offers, with some probability of termination after each refusal. An alternate formulation would have the buyer make price bids, with the seller committing himself to particular probabilities of acceptance depending upon the last bid made, or possible upon the sequence of bids to date.

A third possibility would have the seller announce a price, with the buyer deciding whether to accept or reject. The seller would also announce a probability distribution on second-round prices (including possibly a raised price or withdrawing the item from sale) should the first price be refused, one on third-round prices contingent on second-round refusal, etc. The seller's fixed, posted-price strategy, a polar case of this third approach, turns out to be of particular interest. In this "take-it-or-leave-it" strategy, the seller announces a single price. If refused, no future price offer is made, the buyer leaves the store and the next buyer is called in.

Why consider haggling? Some readers, comfortable with the traditional fixed-price practices of Western merchants, might inquire why a vendor should ever consider haggling. There are two reasons why some alternative to a fixed price strategy might be preferable. First, by announcing an initial price and a probability less than one of continuing to a lower price if that price is refused, the seller can price discriminate. Given appropriate odds, a customer with a high valuation would find the initial asking price preferable to the risky second offer, whereas a potential buyer with a lower valuation would find it preferable on average to wait. Second, by adopting some form of discriminatory pricing policy over time, the seller might be able to gain valuable information about the distribution of reservation values. Note that the two classes of strategies we consider both allow, indeed require, the seller to make commitments as to future actions. In the extensions and generalizations section of our analysis, we inquire what happens when sellers can not make such commitments.

The primary result of our analysis is that in the circumstances described, a single "take-it-or-leave-it" price is optimal from the standpoint of the seller. This is reassuringly consistent with the commonly observed behavior of sellers. (Alternative explanations of fixed-price behavior seem insufficient to explain such a prevalent phenomenon.)

The seller's choice of an optimal strategy is described in Section II; we show that it reduces to a simple optimization problem. The solution is set forth in Section III, and its major implications are derived there. Section IV explores the question whether a seller should ever wish to recall a previously rejected buyer. In Section V, the question of risk aversion on the part of buyers is developed. Section VI addresses a variety of generalizations and extensions of our analysis. It explores in particular the class of situations where sellers can not make binding commitments.

Relation to the literature. In some respects, this analysis follows the tradition of the literature on optimal search. Indeed, in his seminal analysis of information and search, Stigler [1961] alludes to the "higgling process," which serves as a reference point for the present analysis. Most of the ensuing literature on search focuses on the decision of an agent whether to accept a present offer or to seek additional quotes elsewhere. See Chow and Robbins [1961] and the highly useful surveys by Rothschild [1973], and by Lippman and McCall [1976]. These analyses are well suited to describing the problem of a buyer, but not that of a seller. Some analyses, such as Arrow and Rothschild [1975] and Pratt, Wise and Zeckhauser [1979] allow sellers to optimize in setting fixed prices whose distribution in turn will influence buyers' search strategies. This analysis considers a richer array of seller strategies, from fixed-price to haggling, for cases where buyers do or do not make offers, as well as any possible pricing strategy to which the seller can commit himself.

II. The Seller's Choice of an Optimal Strategy

To keep matters simple at the outset, we make the following assumptions.

(A.1) A single object is offered for sale.

(A.2) It costs an amount c to bring a new buyer into the store.

(A.3) Recall of buyers is not permitted.

(A.4) Current information about the reservation value of the next buyer is summarized by a continuously differentiable distribution function $F(v)$, v scaled so that $F(0) = 0$, and $F(1) = 1$.

(A.5) The distribution function $F(v)$ is unaffected bythe seller's choice of strategy.

We have defined the seller's (possibly probabilistic) strategy as S. Once S is announced, the buyer selects his optimal response B. This response depends on his reservation value as well as the seller's strategy. Thus

$$B = b_S(v). \tag{1}$$

We begin by examining the optimal response of the buyer current in the store. To simplify the discussion somewhat we assume that money changes hands only if the object is sold.* Then the expected return to a response B' can be expressed as follows

$$\left\{\begin{matrix}\text{expected}\\\text{buyer}\\\text{gain}\end{matrix}\right\} = \left\{\begin{matrix}\text{probability}\\\text{sale is}\\\text{made}\end{matrix}\right\}\left[\left\{\begin{matrix}\text{reservation}\\\text{wage}\end{matrix}\right\} - \left\{\begin{matrix}\text{expected}\\\text{price}\end{matrix}\right\}\right]. \tag{2}$$

We now obtain simple expressions for both the probability of a sale and the expected buyer gain. Then, from (2), we are able to derive the expected payment of the buyer.

For any selling strategy S and buyer response B' there is some implied probability of sale $H_S(B')$ and expected price $\bar{p}_S(B')$ conditional

---

*This assumption is not critical. The main theorem holds even if we allow for possible payments by the buyer during the haggling process.

upon their being a sale. Furthermore, for any B' in the set of optimal responses there is some v' such that $B' = b_S(v')$. Then we may write the implied probability of sale and expected price as follows.

$$
\text{(3)} \qquad
\begin{aligned}
H(v') &\equiv H_S(b_S(v')) \, . \\
\bar{p}(v') &\equiv \bar{p}_S(b_S(v')) \, .
\end{aligned}
$$

The expected buyer gain, if his response is $B'(=b_S(v'))$ when his reservation value is v, can therefore be expressed as follows:

$$
\text{(4)} \qquad \phi(v',v) = H(v')(v-\bar{p}(v')) \, .
$$

Since we have defined $B = b_S(v)$ to be the buyer's optimal response it must be that $\phi(v',v)$ takes on its maximum at $v' = v$. Throughout we assume only that $H(v)$ and $\bar{p}(v)$ are piecewise differentiable. Then except at points of non-differentiability we have

$$
\phi_1(v,v) \equiv \frac{\partial\phi(x,v)}{\partial x}\bigg|_{x=v} = 0 \, .
$$

Also, from (4)

$$
\phi_2(v,v) \equiv \frac{\partial\phi(v,x)}{\partial x}\bigg|_{x=v} = H(v) \, .
$$

Then, except at points of nondifferentiability the total derivative of the buyer's maximized expected gain,

$$
\frac{d\phi(v,v)}{dv} = \phi_1(v,v) + \phi_2(v,v) = H(v) \, .
$$

While $\phi(v',v)$ is only piecewise differentiable, $\phi(v,v) = \max_{v'} \phi(v',v)$ is both piecewise differentiable and continuous.*

---

*For any v' and v > v'

$$
\begin{aligned}
\phi(v,v) - \phi(v',v') &\leq \phi(v,v) - \phi(v,v') \\
&= H(v)(v-v') \qquad \text{from (4)} \\
&\leq v - v'
\end{aligned}
$$

Then we can reintegrate to obtain the following expression for the buyer's maximized expected gain

$$\phi(v,v) = \int_0^v H(x)dx + \phi(0,0) \tag{5}$$

We are now ready to consider the haggling game from the seller's viewpoint. The expected payment by a buyer with reservation value v is just the probability of a sale times the expected price conditional upon their being a sale ($=H(v)\bar{p}(v)$). Then substituting from (5) into (2) the expected payment is

$$H(v)\bar{p}(v) = H(v)v - \int_0^v H(x)dx - \phi(0,0) \quad . \tag{6}$$

But, as far as the seller is concerned, v and hence $H(v)\bar{p}(v)$ is a random variable with density $f(v)$. Then the expected revenue of the seller is

$$\int_0^1 H(v)\bar{p}(v)f(v)dv = \int_0^1 H(v)vf(v)dv - \int_0^1 f(v) \int_0^v H(x)dxdv - \phi(0,0) \quad .$$

Integrating the second term by parts we have finally

$$\begin{matrix}\text{expected}\\ \text{seller}\\ \text{revenue}\end{matrix} = \int_0^1 H(v)j(v)f(v)dv - \phi(0,0) \tag{7}$$

where

$$j(v) = v - (1-F(v))/f(v) \quad . \tag{8}$$

Since buyers are free to exit from the store without purchase $\phi(v,v) \geq 0$ for all v. In particular for a buyer with reservation value equal to zero we require $\phi(0,0) \geq 0$. Then, since expected revenue is decreasing in $\phi(0,0)$ the seller will choose $\phi(0,0) = 0$. In economic terms he will never sell to a buyer who is not willing to pay anything for the object.

It remains to incorporate the expected gains to the seller in the absence of a sale to the current buyer. Once a buyer has been told he will not be sold the object, he has no incentive to conceal his true valuation. We therefore assume that if the buyer is not sold the object, he reveals its true value to the seller. (We shall show in section III that the seller can elicit such information from a self-interested buyer for essentially zero cost.) This assumption about full revelation merely simplifies our presentation. (A significant polar case of our analysis assumes that the seller learns nothing as he goes along, that in effect he has extensive information about the distribution of buyers' reservation values at the outset.) Using information accumulated to date, the seller computes the expected profit $\pi(v)$ from attempts to sell to future customers. Since $1 - H(v)$ is the probability of not selling to the current customer if his reservation value is v, overall expected future profit at this juncture is

$$(9) \qquad \int_0^1 (1-H(v))\pi(v)f(v)dv \ .$$

Combining (7) and (9) and rearranging, expected total profit is therefore

$$(10) \qquad \bar{\Pi} = [\int_0^1 \pi(v)f(v)dv - c] + \int_0^1 H(v)(j(v)-\pi(v))f(v)dv \ .$$

The bracketed term is the expected profit if the current customer is told that under no circumstances will he be sold the product. Therefore the final term is the increment in expected profit associated with the attempt to sell to this customer. Writing this as $\Delta\bar{\pi}$ it follows that a necessary condition for the maximization of expected seller profit is that $H(v)$

be chosen to maximize

$$(11) \qquad \Delta\bar{\Pi} = \int_0^1 H(v)(j(v)-\pi(v))dv \ .$$

Before proceeding to discuss the nature of the optimal solution, we should stop to say a word about the nature of our assumptions. In any specific application the future expected profits function, $\pi(v)$, will depend upon information gained from previously rejected buyers and the number of buyers remaining to be sampled. This function could also take account of the possibility that the early buyers coming into the store might be more eager and hence have higher reservation values.

If there were many objects that were to be offered for sale, with a constant cost, say x, as opposed to just one in our formulation, the analysis would simplify. We would replace $\pi(v)$ with x in all the equations. The lost future profits are simply the replacement cost of the asset.*

---

*Matters might be more complicated if refusing buyers did not reveal their reservation prices. In that case it might be worthwhile to employ a different class of optimal strategy in order to elicit more information than would otherwise become freely available. As suggested above, complete revelation is consistent with the model of self-interested behavior, though uncharacteristic of many real world situations.

III. The Nature of the Optimal Selling Strategy

We now ask what distribution function $H^*(v)$ maximizes (11), the contribution to expected profit associated with the attempt to sell to the current buyer, and then seek the selling strategy which implies such a distribution function.

We can summarize our answer to the first part of this question as follows.

Proposition 1*: The optimal selling strategy has an implied probability of sale function $H^*(v)$ of the form

$$H^*(v) = \begin{cases} 0, & v < v^* \\ 1, & v \geq v^* \end{cases}$$

where v^* is a root of $j(v) \equiv v - (1-F(v))/f(v) = \pi(v)$.

To simplify notation we begin the proof by defining

$$k(v) = (j(v) - \pi(v))f(v) \quad .$$

Expression (11) for expected current profit then becomes

$$\text{(12)} \qquad \Delta\Pi = \int_0^1 H(v)k(v)dv \quad .$$

Given assumption (A.4), $j(v)$ is negative at $v = 0$. Moreover, the seller always has the option of giving up his search for a buyer. Therefore the expected profit from future attempts to sell the product $\pi(v)$, must be non-negative. It follows that $k(v)$ is negative at $v = 0$ as depicted in Figure 2. However there are no further obvious restrictions on the form

*We are grateful to Barry Nalebuff whose comments on an earlier draft led to the following constructive proof.

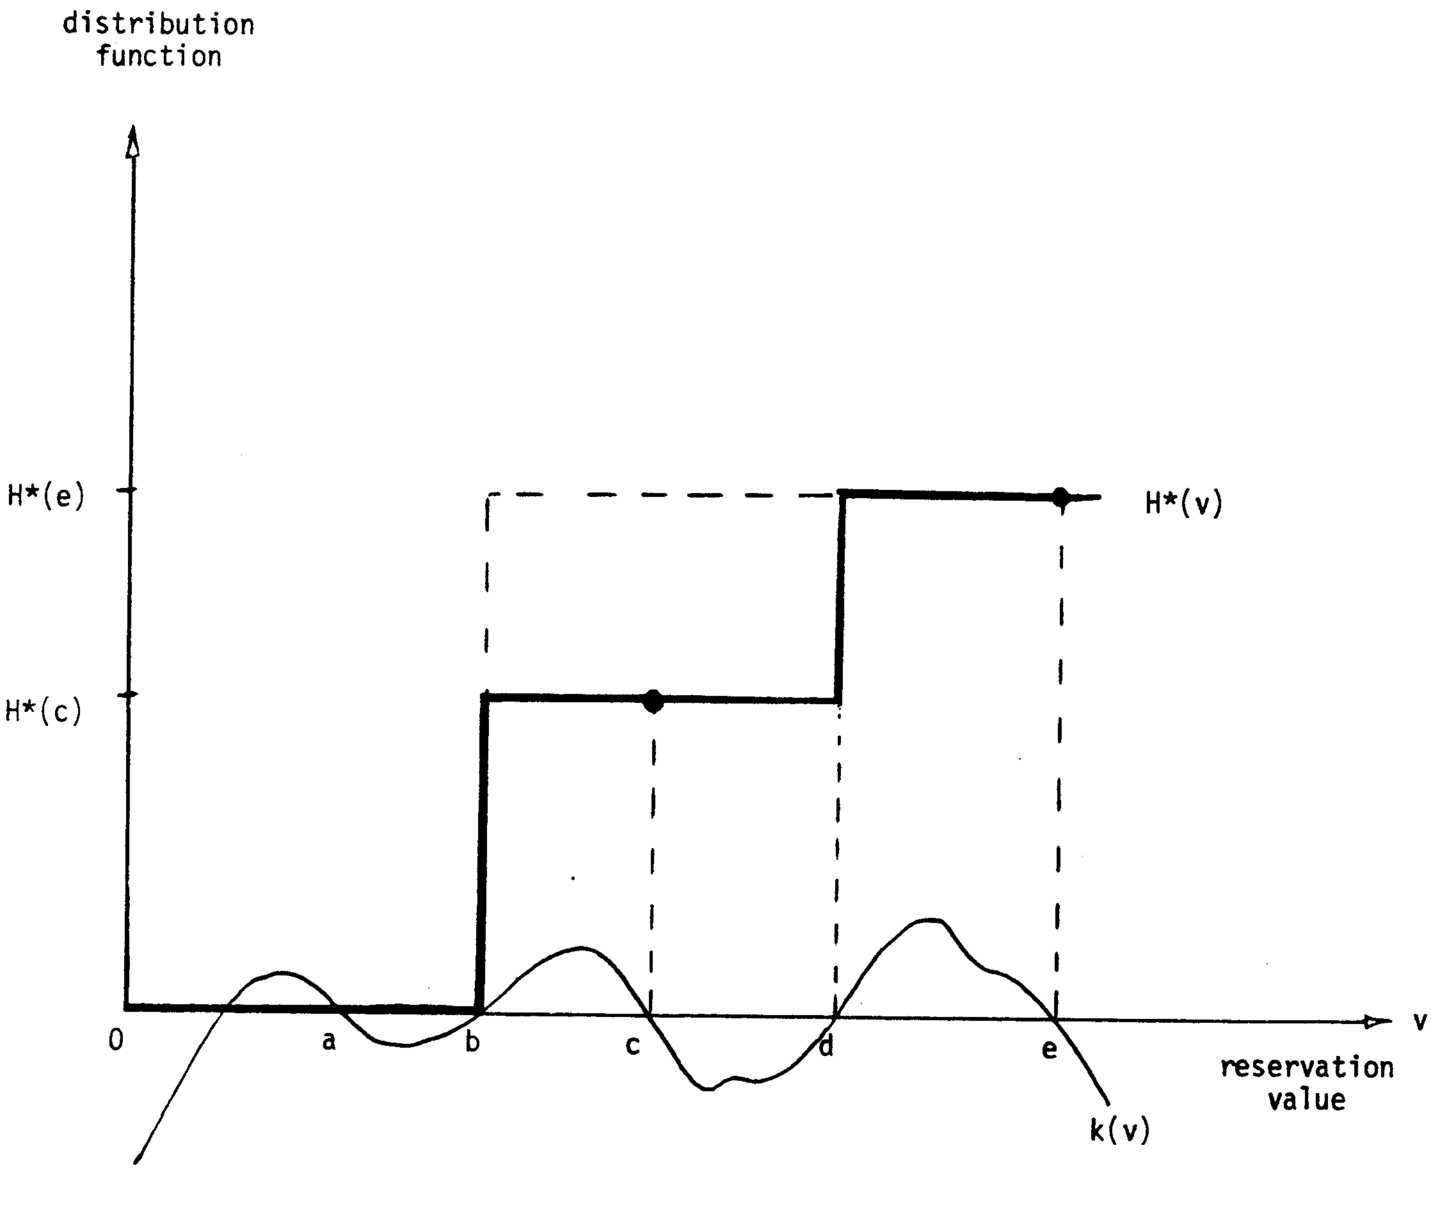

Figure 2: The Optimal Distribution Function

of k(v) and in particular it may change sign any number of times.

As a first step in solving for the optimal distribution function, H*(v), consider all the right-hand endpoints of subintervals over which k(v) is positive. Since H*(v) is necessarily an increasing function, if it is zero at every such point the integral $\Delta\bar{\Pi}$ must be non-positive. Then if search is worthwhile there must be some smallest right-hand end-point c such that H*(c) is positive. This is depicted in Figure 2.

Insert Fig. 2 about here

Given the definition of c we know that H*(a) = 0. Then we can rewrite $\Delta\bar{\Pi}$ as the following sum of integrals.

$$(13) \qquad \Delta\bar{\Pi} = \int_a^b H(v)k(v)dv + \int_b^c H(v)k(v)dv + \int_c^1 H(v)k(v)dv \ .$$

Since k(v) is negative on (a,b), $\Delta\bar{\Pi}$ is maximized by setting H*(v) = 0 on this subinterval. Similarly, since k(v) is positive on (b,c) it is optimal to make H(v) as large as possible on this subinterval. But for H(v) to be a distribution function it must be non-decreasing. Then H*(v) = H*(c) on (b,c).

Now let H*(e) be the optimal value of H at e, the right-hand end-point of (d,e), the next subinterval over which k(v) is positive. Arguing as above it follows that we should make H(v) as small as possible over (c,d) and as large as possible over (d,e).

Combining results, we have,

$$H^*(v) = \begin{cases} 0, & v < b \\ H^*(c), & b \leq v < d \\ H^*(e), & d \leq v < e \end{cases}$$

Then we can rewrite (13) as

$$\Delta\bar{\Pi} = H^*(c)\int_b^d k(v)dv + H^*(e)\int_d^e k(v)dv + \int_e^1 H(v)k(v)dv$$

If the first integral is negative $\Delta\bar{\Pi}$ is maximized by setting $H^*(c) = 0$. Since this contradicts the definition of c the integral must be non-negative. Then $\Delta\bar{\Pi}$ is maximized by setting $H^*(c)$ as large as possible, that is be setting $H^*(c) = H^*(e)$. The optimal distribution function thus has a single step over the subinterval $[0,e)$.

Finally we note that the same kinds of argument can be applied for each additional subinterval over which $k(v)$ is single signed. Therefore there is but a single step at $v = c$ and (13) can be rewritten as

$$\Delta\bar{\Pi} = H^*(c)\int_b^1 k(v)dv$$

It follows that if searching for a buyer is optimal ($\Delta\bar{\Pi} > 0$) $H^*(c)$ must be equal to 1, hence Proposition 1.

Proposition 1 tells us that an optimal strategy is one in which a sale is made if and only if the current buyer has a reservation value $v \geq v^*$. Of course this is none other than the "take-it-or-leave-it" strategy of announcing a fixed price of v^*. We have therefore shown that under the assumptions of our basic model it never pays to randomize or "haggle" over price.

Numbers of buyers. If the pool of potential buyers is finite, the expected profit from future attempts to sell the product, $\pi(v)$, will depend upon the number of unsampled buyers. Adding to this number cannot decrease and will generally strictly increase profit opportunities. That is, for all v, $\pi(v)$ is increasing with the size of the pool of unsampled buyers. We have therefore proved:

Proposition 2: The optimal selling strategy is to announce a single "take-it-or-leave-it" price v^* satisfying the conditions of Proposition 1. Other things equal, this price will be higher if there are more buyers remaining to be sampled.

Price behavior with F(v) known. If the reservation values are known by the seller to come from the distribution $F_0(v)$, so that there is no learning from customers, $j(v)$ is independent of the number of rejected buyers. Moreover, expected future profit after r buyers have been rejected, $\pi_r(v)$, is independent of the $(r+1)$th buyer's reservation value, that is,

$$\pi_r(v) = \bar{\pi}_r .$$

From Proposition 1 the r-th buyer will be offered a price v_r^* satisfying

(14) $$j(v_r^*) - \bar{\pi}_r = 0 .$$

Furthermore, expected future profit prior to bringing this buyer into the store is

(15) $$\bar{\pi}_{r-1} = (1 - F_0(v_r^*))v_r^* + F_0(v_r^*)\bar{\pi}_r - c .$$

Substituting (14) into (15) yields a first-order difference equation for v_r^*. Since $\pi_n = 0$ we can solve for v_n^* from (17) and hence, working backwards, solve for the complete sequence of asking prices $\{v_r^*\}$.

This result is illustrated in Figure 3,employing the assumption that v is distributed uniformly ($F_0(v) = v$). If the last buyer is before the seller the optimal price is 0.5. However, if there are many buyers remaining, the optimal price is close to $1 - \sqrt{c}$.

Insert Figure 3 About Here

Price behavior with F(v) unknown. Suppose that the seller begins with beliefs given by the distribution $F_0(v)$ but, after rejecting r buyers his beliefs are given by $F_r(v)$. While the actual distribution will depend upon the information obtained from the rejected buyers, suppose it is always the case that

(16) $$F_r(v) > F_{r-1}(v) \quad , \quad \text{for all } v \text{ and } r .$$

This assumption captures the notion that with every failure to sell the good probability mass is moved to lower values of v. Given (16), it follows that

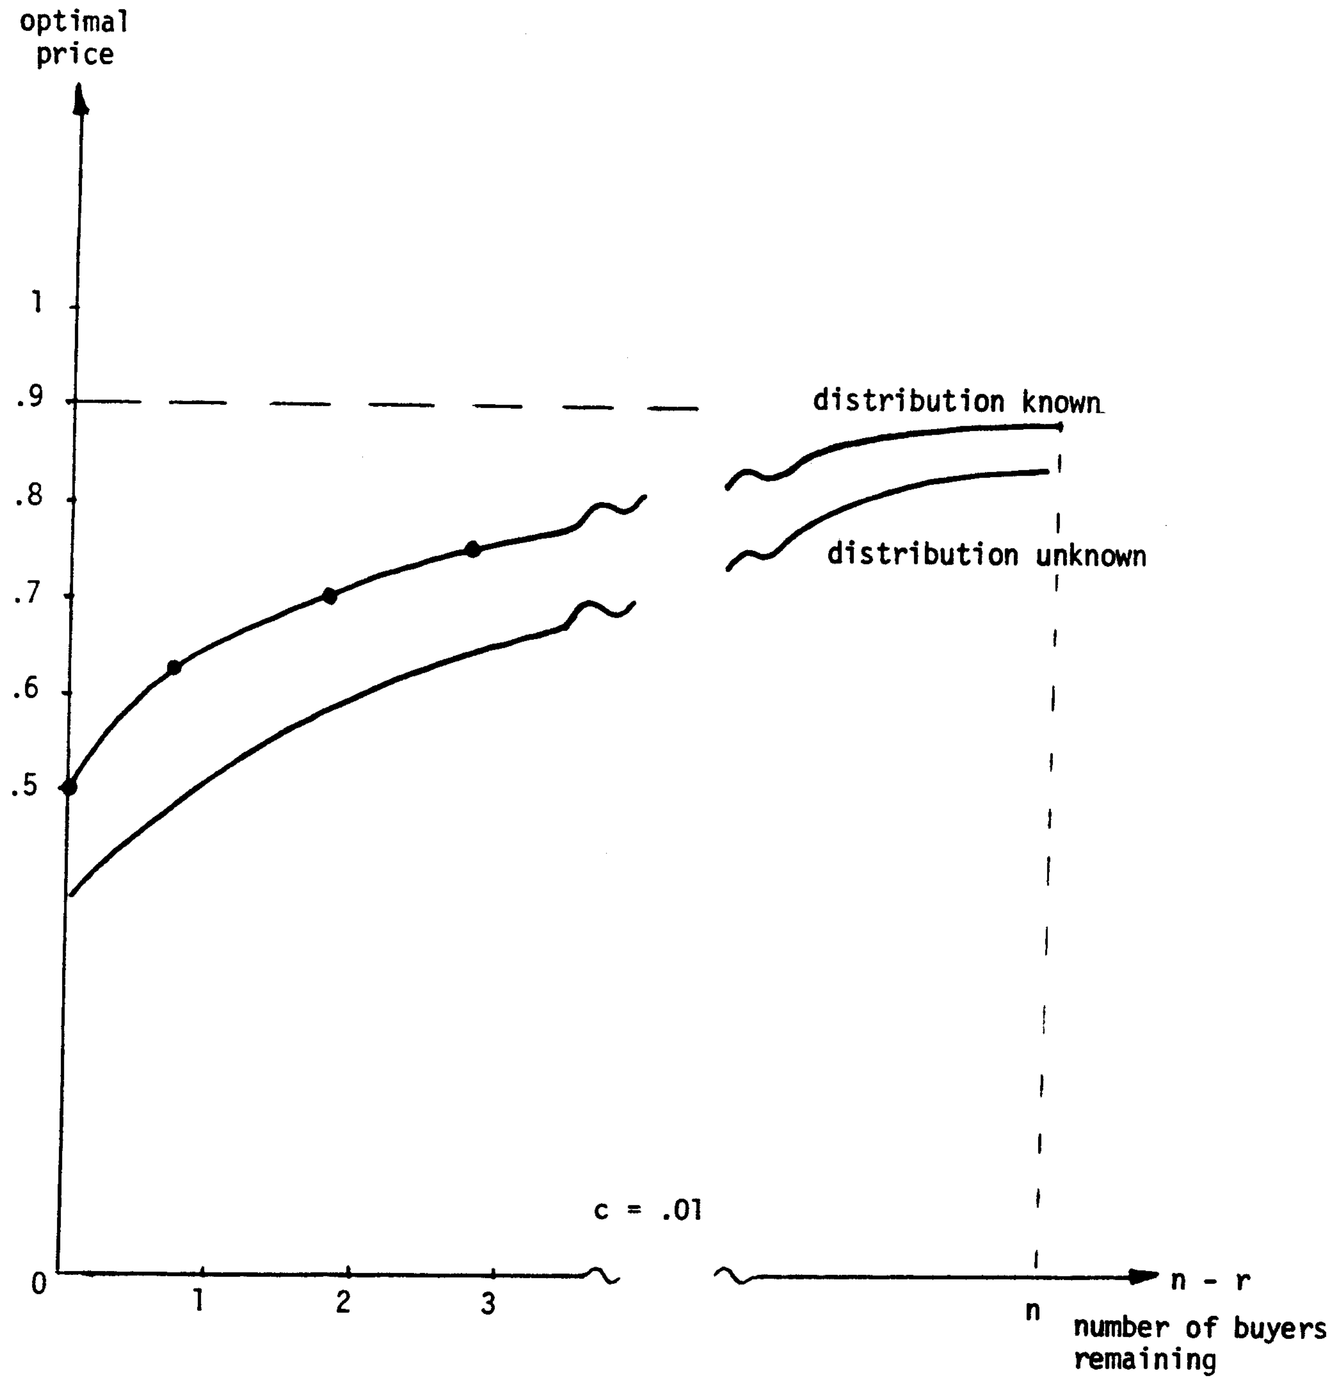

Figure 3: Optimal Pricing

for all v and r, expected future profit, $\pi_r(v)$, is less than $\pi_{r-1}(v)$. Moreover, for all v and r, $\pi_r(v)$ is less than $\bar{\pi}_r$, the profit if beliefs remain fixed at $F_0(v)$.

While generalizations are possible,we consider here only the special case:

(17) $$F_0(v) = v \quad , \quad F_r(v) = v^{\theta_r} .$$

To satisfy (16) we require

(18) $$1 = \theta_0 > \theta_r > \theta_{r+1} \quad , \quad r = 1,2,\ldots .$$

Given (18) it can be confirmed that

(19) $$j_{r+1}(v) > j_r(v) > j_0(v) = 2v - 1 \quad , \quad r = 1,2,\ldots .$$

Combining results it follows that

(20) $$j_{r+1}(v) - \pi_{r+1}(v) > j_r(v) - \pi_r(v) ,$$

and

(21) $$j_r(v) - \pi_r(v) > j_0(v) - \bar{\pi}_r .$$

Let $\{v_r^*\}$ and $\{v_r^{**}\}$ be the optimal price sequences without and with learning. From (20) we have

(22) $$v_{r+1}^{**} < v_r^{**} ,$$

and from (21)

(23) $$v_r^{**} < v_r^* .$$

Then the optimal price sequence with learning is as depicted in Figure 2, increasing in r and everywhere below the price sequence without learning. (Under some circumstances with learning, if matters turn out unfavorably the seller may stop quoting prices even with buyers remaining.)

While the assumptions utilized to compare pricing behavior with and without learning are relatively mild* it should be noted that they are not innocuous. Indeed with F(v) unknown the seller may wish to increase his price after rejecting some customers. Suppose, for example,** that the seller starts out by believing that with probability ρ the distribution of reservation values is given by

$$F_1(v) = \begin{cases} 0 & , \quad 0 \leq v \leq .1 \\ .1 & , \quad .1 \leq v < 1 \\ 1 & , \quad 1 \leq v \,, \end{cases}$$

and with probability $(1-\rho)$ the distribution of reservation values is

$$F_2(v) = \begin{cases} 0 & , \quad 0 \leq v < .5 \\ 1 & , \quad .5 \leq v \,. \end{cases}$$

If ρ is sufficiently small relative to the search cost, c, and c is not so large as to make search unprofitable--for example, if $\rho = 0.1$ and $c = 0.4$--the optimal first-round strategy is to announce a reserve price of just less than 0.5. If this is rejected by the first buyer the seller knows that the buyer is not drawn from $F_2(v)$. He therefore revises ρ upwards to unity and announces a second-round price of just less than unity. This leads to:

*These assumptions are especially mild if the only information the seller receives is that the buyer either accepted or rejected his offer.

**The example is adapted from one suggested by Rothschild [1973].

Proposition 3: The optimal selling price, v^*, may rise with the number of rejected buyers, unless the seller's beliefs about the distribution of reservation values are unaffected by sample information.

Eliciting reservation values. We have assumed that self-interested, rejected buyers will reveal their true reservation values. To justify this assumption, we must demonstrate that the seller can provide an incentive for truthful revelation from such buyers at arbitrarily low cost to himself. Consider a buyer who has just rejected the optimal take-it-or-leave-it price of v^*. The seller now asks the buyer to make an offer, m. The seller commits to accept the offer on a probabilistic basis according to the following rules:

$$\text{Prob[buyer's offer of } m \text{ is accepted]} = \alpha m.$$

For a bid of m, the expected gain to the buyer is

$$\alpha m(v-m) \quad . \tag{24}$$

This value is maximized by setting $m = v/2$. Therefore, by observing m, the seller can readily infer v. Substituting this optimal value in (24), the expected gain to a buyer with reservation value v is $\alpha v^2/4$.

We have found a way to elicit true information. Now we must show that the cost of this method can be made arbitrarily small, and that this scheme as a follow on to the original take-it-or-leave-it offer will not affect the buyer's original acceptance. Both objectives are achieved by making α arbitrarily small. As α shrinks towards 0, the expected buyer gain from the follow on elicitation goes to 0. The buyer should therefore accept the seller's initial offer unless v is very close to v^*. Moreover, as α approaches 0, the probability of selling at a price below v* approaches 0. Therefore, the expected loss in revenue to the seller also goes to 0. We have shown:

Proposition 4: At arbitrarily low cost, the seller can induce all buyers who are not willing to pay the optimal price, v*, to reveal their true reservation values.

The strict nature of the optimal seller strategy. Strictly speaking, the optimal strategy for the seller is to commit himself to a two-round game. On the first round he will quote a price. If this price is rejected, he will ask the buyer to quote a price, with a probability schedule for acceptance. If the seller rejects this price, the negotiations cease. In practice, the scheme turns out to be of the take-it-or-leave-it variety, with the additional twist that if the object is left, the self-interested buyer is induced to reveal his true reservation price.

We should be cautious not to exaggerate the virtues of this elicitation mechanism, which is predominantly a theoretical nicety. It is of low cost to the seller precisely because it is of low value to the buyer. Therefore, it seems plausible that information received by the seller would be subject to considerable noise. Many buyers may even refuse to play the game.

Optimal seller strategy when buyers only accept or reject. Recognizing the possible difficulty of implementing our elicitation scheme, we turn next to the question of optimal seller strategy when buyers only accept or reject price quotas, and therefore make no price quotes of their own. Seller strategies can be of two forms. With a probabilistic declining offer strategy, the seller announces an asking price together with a probability distribution on next price (possibly withdrawing the itme from sale) should that price be rejected, and a probability distribution after second-round rejection, etc. The alternative is a fixed-price, take-it-or-leave-it strategy. Interestingly, the take-it-or-leave-it strategy remains optimal in these circumstances.

With a probabilistic declining offer strategy, let $\hat{G}(p)$ be the probability that the seller will stop at an asking price of p or less. Also, let p(v) be the highest price that a buyer with valuation v will accept. We assume that $p(\cdot)$ is a non-decreasing function. Suppose that for the current buyer the seller has determined that he will withdraw his asking price at $\hat{p} = p(\hat{v})$. His expected current profit is therefore

$$\int_{\hat{v}}^{1} p(v)f(v)dv - c \quad .$$

If the product is not sold then the seller knows that the current buyer's valuation is less than $\hat{v}$. Let $\psi(\hat{v})$ be the expected profit from future attempts to sell the object given the information that the current buyer's valuation is less than $\hat{v}$. The seller's overall expected profit is therefore

$$\int_{\hat{v}}^{1} p(v)f(v)dv + \psi(\hat{v}) - c \quad .$$

Define

$$G(\hat{v}) \;=\; \hat{G}(p(\hat{v}))$$

$$= \text{Prob}\{\text{sale is made to an individual with valuation of at most } \hat{v}\}.$$

The expected profit from the declining offer strategy is therefore

$$(25)\qquad \bar{\pi} = \int_0^1 [\int_{\hat{v}}^1 p(v)f(v)dv + \psi(\hat{v})]g(\hat{v})d\hat{v} - c\ .$$

Obviously it is suboptimal to sell to someone with a zero reservation value. Then $G(0) = 0$. Furthermore there is no advantage in refusing to name any asking price. Then $G(1) = 1$. Integrating (25) by parts and making use of these endpoint conditions we therefore have

$$(26)\qquad \bar{\pi} = \int_0^1 p(v)G(v)f(v)dv + [\psi(\hat{v})G(\hat{v})]_0^1 - \int_0^1 \psi'(\hat{v})G(\hat{v})d\hat{v} - c$$

$$= \int_0^1 [p(v)f(v) - \psi'(v)]G(v)dv + \psi(1) - c\ .$$

Comparing equations (26) and (7) it follows directly that the seller's optimization problem has exactly the same structure as before. Thus once again the optimal selling strategy is to announce a "take-it-or-leave-it" price rather than attempt to discriminate by "haggling" with potential buyers. To summarize we have derived:

<u>Proposition 5</u>: If the only information available to the seller is whether or not a buyer will accept or reject his asking price, the "take-it-or-leave-it" pricing strategy dominates any probabilistically declining asking price strategy.

Since $\pi(v)$ was defined as the expected revenue from attempts to sell to future buyers, given that the current buyer has a reservation value of v, and $\psi(\hat{v})$ is the expected revenue from attempts to sell to future buyers, given that the current buyer has a reservation value of $\hat{v}$ or less, it must be the case that

(27) $$\psi(\hat{v}) \leq \int_0^{\hat{v}} \pi(v)f(v)dv \quad .$$

A sufficient condition for (30) is

(28) $$\psi'(v) \leq \pi(v)f(v) \; .$$

The latter inequality would imply that, under the conditions of Proposition 5, the seller's optimal take-it-or-leave-it price would be lower than when the rejected buyer reveals his true reservation value. However, (28) is certainly not necessary for (27) so we are unable to draw such an inference. Indeed we conjecture that if the seller is a Bayesian, updating his beliefs about the distribution of reservation values, the optimal price might turn out to be higher when buyers only accept or reject rather than when they reveal their reservation values after rejecting.

IV. The Possible Gains to Recall

Thus far we have assumed that there is no opportunity to recall individuals who leave the store. If the cost of locating such individuals, call it k, is sufficiently high, recall will never be undertaken. (It is quite possible that k could be higher than c, the cost of securing a new buyer.) Here we consider the circumstances under which the recall option might be used.

Results potentially differ depending on whether populations are finite or infinite, and whether there is learning about the distribution of buyers' reservation values. There are thus eight cases of interest.

Recall Potentially Beneficial

	$k \geq c$		$k < c$	
	Finite	Infinite	Finite	Infinite
No Learning	(A) yes	(B) no	(C) yes	(D) no
Learning	(E) yes	(F) no	(G) yes	(H) yes

Any positive result that applies for $k \geq c$ applies for $k < c$. Moreover, since no learning is a polar case of learning, any positive result that applies for no learning applies to learning as well. We shall first demonstrate the positive result for (A) using the easiest situation where $k = c = 0$. This implies positive results for (C), (E) and (G). Next we will derive a negative result for (F), which implies a negative result for (B). Then we will examine the negative result for (D) (which also implies a negative result for (B)). Finally, we present a successful example for (H).

No learning, finite population (A). A simple example illustrates the potential value of the recall option. There are two buyers. The common but independent distribution of their reservation prices is:

$$\text{Prob}\{v = 1/2\} = 1/4$$
$$\text{Prob}\{v = 3/4\} = 1/2$$
$$\text{Prob}\{v = 1\} = 1/4 .$$

The buyers enter sequentially. The cost of securing them, c, and the cost of recall, k, are both zero. With no recall the seller will always announce prices just below one of the possible values of v. The five alternative price sequences are as follows: {1,1}, {1,3/4}, {1,1/2}, {3/4,3/4}, {3/4,1/2}. It is easy to compute the expected profit for each sequence. We have,

$$\bar{\pi}\{1,1\} = 28/64$$
$$\bar{\pi}\{1,3/4\} = 43/64$$
$$\bar{\pi}\{1,1/2\} = 40/64$$
$$\bar{\pi}\{3/4,3/4\} = 27/64$$
$$\bar{\pi}\{3/4,1/2\} = 44/64 .$$

The preferred price sequence is {3/4,1/2} yielding an expected profit of 44/64.

Contrast this with the sequence employing recall of {7/8, 3/4, 1/2}. If the first buyer rejects the price of 7/8 and the second buyer rejects the price of 3/4, then the first buyer is offered the object at a price of 1/2. Suppose the first buyer has a reservation value of 1. If he rejects the price of 7/8 he knows that with probability 3/4 the second buyer will purchase at the price of 3/4. He can therefore obtain the object at a price of 1/2 with probability 1/4. This action thus yields an expected gain of

(1/4)(1 - 1/2) = 1/8. By setting the initial price no greater than 7/8 the seller thus provides an incentive for immediate purchase. The expected profit from this scheme is then

$$\bar{\pi}\{7/8, 3/4, 1/2\} = (1/4)(7/8) + (3/4)(3/4) + (3/4)(1/4)(1/2)$$

$$= 47/64 .$$

Therefore the recall option is valuable offering a gain of 47/64 - 44/64 = 3/64. A continuity argument establishes that the recall option remains valuable if search and recall costs are positive but small.

Optimality of the recall option in case (A). Is the pricing strategy described above optimal from the seller's viewpoint? The theory of optimal auctions proves helpful in answering this question. Maskin and Riley [1980] show that for discrete distribution functions the best the auctioneer can do is announce a finite set of prices and have each buyer submit a sealed bid consisting of one of these prices (or not submit a bid at all). The high bidder pays his bid and is awarded the object. In the case of a tie the winner is selected randomly.

For the simple example described here this set of prices is {23/28,11/16,1/2}. At these prices (or strictly speaking at prices which are just lower than these) a buyer with reservation value of 1 bids 23/28 while a buyer with reservation value of 3/4 bids 11/16. Expected profit is then

$$\bar{\pi}^* = \left\{ \begin{array}{l} \text{Prob}\{\max(v_1,v_2) = 1\}\ (23/28) \\ +\ \text{Prob}\{\max(v_1,v_2) = 3/4\}\ (11/16) \\ +\ \text{Prob}\{\max(v_1,v_2) = 1/2\}\ (1/2) \end{array} \right\} = 47/64 .$$

Since the sequential scheme described above matches the optimal auction scheme, it is itself optimal.*

A positive result in case (A) obviously implies a positive result as well in (C), where recall would if anything be a more attractive option. Moreover, satisfaction of (C) implies satisfaction of (G) as well. Strictly speaking, this is true trivially since (C) is a special case of (G). More important, a continuity argument assures that recall may be beneficial under (G) even if strict learning is required.

Learning, infinite population, (F). Since we get a negative result, consider the most favorable case where the cost of recalling a buyer who refused earlier equals the cost of securing a new buyer. At initial glance we might think that knowing an individual refused a higher price might tell us that he actually had greater (conditional) reservation price density at or above some present price quote. That advantage might outweigh the disadvantage that the possibility of recall could make individuals less willing to accept first offers.

A simple thought experiment reveals that recall can not be optimal in such circumstances. A number of buyers have been offered prices and refused. The choice is now between buyer X, an earlier refuser, and buyer Y, who has not yet been asked. To make the comparison we look at two situations. Would Y have accepted our earlier price when X refused, or would he have refused? If he would have accepted, then Y is to be preferred to X now, for he will surely accept. If he would have refused, then our knowledge of Y is the same

*We are grateful to T. Nicolaus Tideman for pointing out the equivalence between the sequential policy strategy and the optimal auction strategy. Unfortunately the equivalence appears to hold only for two and three point distributions.

as our knowledge of X. Given that we are indifferent in one case and prefer Y in the other, buyer Y should be our choice. Recall is never beneficial. This negative result for (F) implies a negative result for (B) as well, for (B) is a special case of (F).

<u>No learning, infinite population, $k < c$, (D)</u>. We have a negative result for (B), but what of (D), the situation where recall may be substantially less expensive than securing a new buyer? Consider first expected total profits $\bar{\pi}_N$ with no recall. Since the distribution of buyers is known, expected future profits in the absence of a sale to the current buyer, $\pi(v)$, must also be $\bar{\pi}_N$. Thus, making use of Proposition 1, the seller sets a price v^* in each period satisfying

$$j(v^*) - \bar{\pi}_N = 0 . \tag{29}$$

Given this strategy, the expected profit of the seller is the price v^* less the expected number of buyers times the cost c per buyer:

$$\bar{\pi}_N = v^* - \frac{c}{1-F^*} , \quad \text{where} \quad F^* = F(v^*) . \tag{30}$$

Together these equations determine v^*.

Consider now the effect on expected profit of allowing the first buyer to be recalled at some later time at no additional cost. For simplicity we assume that the seller can recall him after the departure of the second buyer. For any announced strategy by the seller the response $B = b(v)$ by the first buyer results in some probability

$$H_1(v) = \hat{H}_1(b(v))$$

of an immediate sale, and, in the absence of an immediate sale, a further probability

$$H_2(v) = \hat{H}_2(b(v))$$

of a later sale to the first buyer. Given the assumptions of no learning and an infinite population, the optimal reserve price for all other buyers remains v^*.

Then if $\bar{p}_i(v) = p_i(b(v))$ is the expected price paid by the first buyer conditional upon making a successful offer on the ith contact $(i = 1,2)$, expected seller profit at the outset with recall is

$$(31) \quad \bar{\pi}_R = E_v[H_1(v)\bar{p}_1(v) + F^*(1-H_1(v))H_2(v)\bar{p}_2(v)] + (1-F^*)E_v(1-H_1(v))\bar{\pi}_N + F^* E_v(1-H_1(v))(1-H_2(v))\bar{\pi}_N - c .$$

Furthermore, repeating the argument of Section II we can rewrite the first term in (31) as

$$(32) \quad \int_0^1 [H_1 + F^*(1-H_1)H_2]j(v)f(v)dv .$$

Making substitutions using (29), (30), and (31) we have finally

$$\bar{\pi}_R = \int_0^1 [H_1 + F^*(1-H_1)H_2][j(v) - j(v^*)]f(v)dv + \bar{\pi}_N - c .$$

Following the argument of Section III it can be confirmed that $\bar{\pi}_R$ is maximized by setting

$$H_1(v) = \begin{cases} 0 , & v < v^* \\ 1 , & v \geq v^* \end{cases} , \quad H_2(v) = 0 .$$

That is, in the infinite population case the seller's optimal strategy is to sell to the first buyer immediately or not at all. The optimal strategy has no recall.

Learning, infinite population, $k < c$, (H). Might the presence of learning turn the negative (D) result positive? Interestingly, the answer is yes, as a simple example makes clear. Assume that the seller and the buyers knew that all buyers had identical reservation prices, and that there were three equally likely situations, that that price was 1, that it was .6 and that it was .2. The cost of securing a buyer is positive, say .1. The cost of recall is 0.

The optimal strategy without recall is to quote a price of 1 minus a hair to the first buyer; .6 minus a hair to the second, should the first refuse; and a hair less than .2 to the third, should the second refuse as well. The expected payoff is

$$1/3(1 - .1) + 1/3(.6 - .1) + 1/3(.2 - .1) = .5,$$

leaving split hairs aside.

With recall possible, a number of superior strategies are available. One would be to make the same first two quotes, but then recall the first buyer for quote 3. Note in this case that if the first buyer's reservation price is 1 he will certainly buy at the first trial, for he knows that the second buyer's reservation price is also 1 and he will surely buy if given the chance. Better a certainty of a smidgen than a zero probability prospect of a large gain.

It may seem puzzling at first that learning turns a negative result positive. At first our intuition might suggest that it is because it is worthwhile to use a different strategy under learning so as to acquire more information. That, however, is not the case here. The reason that learning can make recall worthwhile is that the recalled buyer is purchasing a commodity that is less valuable to the seller than it was when first he refused. There is no such diminution in prospective value in the case without learning.

To sum up our results on recall, with finite populations recall must always be considered a possibility. With infinite populations, recall can only be advantageous when there is learning and the cost of recall is less than the cost of obtaining a new buyer. Might it be less expensive to recall than to secure a new buyer? It seems reasonable. The salesman can record the phone number of the refusing buyer, "just in case."

V. Randomized Strategies for Risk Averse Buyers

If buyers are risk neutral, the optimal strategy of the seller is to announce a single "take-it-or-leave-it" price. Here we demonstrate th.. this is not generally true when buyers are risk averse. Instead the seller can increase expected profits by utilizing a "probabilistically declining offer" strategy.

A simple example provides the proof. Suppose that there is just one buyer with reservation value v, a random draw from the uniform distribution on the unit interval 0 to 1. The best single-price strategy is to choose a price p to maximize

$$p \text{ Prob}\{v \text{ exceeds } p\} = p(1-p) .$$

Then $p^* = 1/2$ and the expected revenue is 1/4. From Section II we know that if the consumer is risk neutral this strategy is optimal. However suppose

$$u(v-p) = (v-p)^{\beta} , \quad 0 < \beta \leq 1 .$$

While introducing the risk aversion does not alter the optimal single price strategy it does change a buyer's response to other selling strategies. Suppose the seller announces that the object will be withdrawn from sale at a price of p or lower according to the probability function G(p), where

$$G(p) = p . \tag{33}$$

The buyer is asked to make a bid knowing that the lower the bid the lower the probability of being allowed to purchase the object. The buyer then chooses his bid p(v) to maximize his expected utility

(34) $$u(v-p)G(p) = (v-p)^{\beta}p \quad .$$

It can be readily confirmed that the optimal bid is

(35) $$p(v) = v/(1+\beta) \ .$$

For a buyer with reservation value v the probability of sale is G(p(v)) and the revenue of the seller conditional upon such a sale is p(v). Utilizing (33) and (35) expected seller revenue is therefore

$$\int_0^1 p^2(v)f(v) = \frac{1}{(1+\beta)^2}\int_0^1 v^2dv = \frac{1}{3(1+\beta)^2} \ .$$

Since expected revenue from the single price scheme is 1/4 it follows that the alternative scheme dominates if the buyer is sufficiently risk averse, that is,if β is sufficiently small.

The alternate scheme outlined here was called earlier a "probabilistically declining offer" strategy. Such a strategy has not been shown to be optimal, just superior to the fixed price, take-it-or-leave-it scheme in some circumstances. If the buyers' risk aversion levels were known, it would be possible to devise an optimal scheme.

VI. Extensions and Generalizations

Extensions and generalizations of this work could come in many areas. They could allow for:(1) multiple items for sale, (2) the seller's strategy to affect the distribution of buyers, (3) consideration of cases where sellers are unable to make commitments in advance, (4) alternative market structures where there may be competitive elements, (5) buyers moving first or simultaneous moves, and (6) examination of anecdotal and statistical evidence on the actual pricing behavior of firms. We shall just comment on the first three of these areas here.

Multiple items for sale. The models we have outlined above apply, albeit with a bit more complexity, when there is any number of items for sale. Consider the easiest case, no learning. Then expected profits with X items is simply X times the expected profit with 1; the identical fixed-price should be employed.*

An important special case would allow the store to sell as many items as it wished, but to pay a fixed cost per item. This formulation would actually simplify a number of calculations. Each buyer could in effect be treated separately. The store would simply set the fixed price that maximized:

Probability of Sale x (Price - Fixed Cost) .

The cost of bringing a customer into the store would become a sunk cost. It would not affect the optimal price; it might, however, induce the store to go out of business.

Situations in which there were increasing supply costs--of which a special case is any finite supply--would turn out to have properties like many depletable resource models. Attached to each item sold would be not only its

*If there is discounting with multiple items, then with an infinite population of buyers the optimal fixed price will rise as items are sold.

own immediate marginal cost, but a shadow price as well that indicates the (possibly discounted) increased cost in all future items sold due to its sale.

Learning creates no difficulties in the fixed-cost case. With rising supply prices, however, learning can make matters considerably more complex. In essence, experience with sales and attempted sales helps the seller assess more accurately the shadow price to be attached to low-cost supplies. Models of learning with increasing supply price are a challenging subject for future work.

Seller's strategy affects the distribution of buyers. Individuals with low reservation prices rarely wander into Gucci's; discount stores by contrast are disproportionately populated by bargain shoppers. If a store's committed pricing strategy can be communicated to the outside world, that will also have an effect on who comes to shop there.

To deal with situations such as these requires a more elaborate model that incorporates the cost of securing and dealing with buyers. If buyers were perfect at self selection, and if the cost of bringing a buyer into a store remained constant at c, then a store would simply commit itself to charge the highest reservation price of any buyer. Such a model is nonsensical. Buyers would enter the store only after very long intervals.

A more realistic model would make the cost of securing a buyer a function of the seller's strategy. The world implicit in the model might have buyers walk by the store, with only those who are interested entering. If all costs were merely waiting time, for example, rent and heat, then costs would be proportional to the number of individuals passing by, and all the previous models would apply.

Alterations would be required, however, if there were also costs of servicing. Intuitively it is clear that the possible avoidance of service costs would provide an additional argument for a fixed-price policy; such a policy would still be optimal. With service costs, however, the fixed price should be raised. Consider a situation where buyer reservation prices are distributed uniformly on the interval 0 to 1. Until the item is sold, each potential buyer imposes a cost of α. If the store sets a price of p, it will on average take $1/(1-p)$ trials to sell the object. The store sets p to maximize $p - \alpha/(1-p)$. The optimal $p = 1 - \alpha$. If, on the other hand α were all service cost, the store would charge a price of 1. Realistically, there will be some service cost and some cost of waiting for a buyer, and the optimal price will lie between the two extremes, i.e. above the optimal price when the seller's strategy does not affect the mix of buyers.

Commitment impossible, the inevitability of haggling. Our initial interest in this problem was sparked by consideration of the differences between markets where haggling was the order of the day, and others where fixed and posted prices seemed to rule. On a purely observational level, it seems that large stores with established reputations are most likely to employ fixed-price strategies. Though a variety of alternative explanations is possible, it struck us that this observation was consistent with the fact that established stores are much better able to make firm commitments to fixed prices.

Suppose a motorist walks into a country antique store on a back road and he refuses the owner's first price; it is difficult for the owner to maintain that by lowering his price he will ruin his reputation for the future. The seller's inability to commit himself tc a strategy in advance almost makes haggling the order of the day, unless, of course, the buyer finds such a process extremely unattractive. With just anecdotal evidence to support our theoretically guided insights, we conjecture that haggling will be discovered to be much more common when sellers encounter particular buyers infrequently and do not have reputations to maintain or establish.

This raises the possibility of another advantage of brand-naming or resale price maintenance. It may make it possible in certain circumstances for a seller to commit himself to a strategy which he otherwise would not be able to commit to. In particular circumstances where offer costs push average costs above marginal costs, buyers may welcome the availability to sellers of firm price commitments. Without such commitments, a weak variant of ruinous competition may make it difficult to find the product in the marketplace. No seller would take on the fixed costs that are required to make the product more conveniently available.

To model the haggling process is a most intriguing problem, a problem that will have to await a subsequent paper. A preliminary model based on our probabilistically declining offer model seems promising. The major difference from our previous model is that the seller's strategy must follow the optimality principle of dynamic programming. It must always be optimal from any point forward. Alas for the seller, this is not an optimality principle for two-person, non-zero-sum games. The inability to make binding commitments about future strategies, that is the inability to commit himself to behave non-optimally in the future, hurts the seller. One natural result of this forever-forward-optimality constraint is that the seller can not withdraw the item from sale until his reservation price is reached. (When there are costs of securing buyers, this reservation price may be a function of what he learns.)

VII. Summary and Conclusion

We have provided a strong theoretical justification for the pricing strategy found in a wide variety of stores. Prices are established and buyers can accept them or seek to buy elsewhere. So long as the seller can make a firm commitment in advance on his strategy, and so long as buyers are risk neutral, expected-value maximizers, this result is robust. The strategy is optimal in comparison to any other, including all forms of buyer involvement, price quoting behavior, etc. (The strict result requires that the seller induce the rejected buyer to reveal his reservation price by offering a microscopic probability of purchasing the item at a bargain price.)

Assuming that the buyer can only accept or reject price offers, which implies he cannot fully reveal his reservation price, the take-it-or-leave-it pricing strategy remains optimal.

A variety of special cases yield interesting optimal selling strategies. When there is learning or buying populations are finite, recall of rejected buyers may be desirable. When buyers are risk averse, take-it-or-leave-it strategies may not be optimal; haggling may be superior. Almost by indirection, we have answered the question in our title. You should haggle when buyers are risk averse.

What if buyers are risk neutral? You should haggle only when you can not make a convincing commitment not to. Stores without established reputations, or market encounters that are highly occasional must be expected to give rise to haggling behavior, although such behavior is not optimal for the seller. He just can not commit himself to something better. In general, sellers should seek ways to commit themselves to firm prices.

REFERENCES

Arrow, Kenneth and Michael Rothschild, unpublished manuscript. 1975.

Chow, Y.S. and Howard 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.

Lippman, Steven A. and John J. McCall, "The Economics of Job Search: A Survey" Econ. Inq., June 1976, 14(2), pp. 155-89.

Maskin, Eric and John G. Riley, "Auctioning an Indivisible Object" Discussion Paper, Kennedy School of Government, Harvard University, 1980.

Pratt, John, David Wise and Richard Zeckhauser, "Price Variations in Almost Competitive Markets," Quarterly Journal of Economics, May 1979, 189-211.

Rothschild, Michael, "Models of Market Organization with Imperfect Information: A Survey," Journal of Political Economy, Nov./Dec. 1973, 81(6), pp. 1283-1308.

Stigler, George J.,"The Economics of Information," Journal of Political Economy, June 1961, 69, pp. 213-25.