Optimal Selling Strategies: When to Haggle, When to Hold Firm
Author(s): John Riley and Richard Zeckhauser
Source: *The Quarterly Journal of Economics*, Vol. 98, No. 2 (May, 1983), pp. 267-289
Published by: Oxford University Press
Stable URL: https://www.jstor.org/stable/1885625

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

JOHN RILEY AND RICHARD ZECKHAUSER

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. This strategy is superior to any other for finite or infinite buyer populations, whether there is learning or the distribution of buyer prices is known at the outset, with one object for sale or many. Although 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.

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 nonnegotiable prices are employed in most cases, although for a range of goods from autos to real estate there may be considerable flexibility in prices.[1]

Why do we find posted fixed prices? In a perfectly competitive market, prices reflect marginal cost. A sale at a lower price would obviously not be worthwhile. Our focus in this paper is on goods that yield a direct profit to the vendor when they are sold. They include products sold on oligopolistic or monopsonistic markets. They also include goods sold on near-competitive markets that suffer the slight imperfection that there is a cost of offering the good for sale, such as rent or the salary of a sometimes-idle salesclerk, which must be covered through price. 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

* Eric Maskin, John McCall, Barry Nalebuff, William Samuelson, Nicolaus Tideman, and two referees provided helpful comments. The support of the National Science Foundation for Riley's research and the Harvard Faculty Project on Regulation for Zeckhauser's research is gratefully acknowledged.

1. Our work in this area has heightened our sensitivity to the question of firmness versus flexibility in prices. We have discovered flexibility in some 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.

The Quarterly Journal of Economics, May 1983 CCC 0033-5533/83/020267-23$03.30

negotiation. Not surprisingly, many customers prefer to shop at stores with clearly posted prices.

We offer a quite different explanation for fixed prices: In comparison to any haggling strategy, a commitment to a firm price is of benefit to the seller.

I. FORMULATION

A series of buyers will enter the seller's showroom until a sale is made or the object is withdrawn. It costs a buyer nothing to stay for another offer from the seller.[2] Buyers are risk-neutral and seek to maximize their expected consumer's surplus, the difference between their reservation value and the price they pay.

The seller is risk-neutral: He maximizes expected profit. Our analysis relies on the following:

ASSUMPTION. The seller is able to make a commitment to employ any strategy he wishes, and he conveys this commitment to each buyer.

In choosing his strategy, the seller compares the ways buyers with different reservation prices will respond to each strategy. Knowing a buyer's optimal strategy in response to a committed seller strategy, one can then compute the probability of sale and the expected price conditional upon sale. The situation can be understood as the two-player sequential game diagramed in Figure I.

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 buyer reservation prices is determined. In this case the seller is not told which distribution is chosen, but he has prior beliefs about the likelihood of different distributions and updates those beliefs on the basis of his experience with refusing buyers as he goes along.[3]

At move 1 the seller commits himself to a contingent-pricing strategy S. For example, S might say that if he is turned down at a

2. While it may seem odd to ignore the buyer's time costs of haggling, the alternative assumption of equal positive time costs for each buyer simply makes haggling still 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 among buyers.

3. We might think metaphorically of selecting an urn at random, with each urn containing a different distribution of balls indicating reservation prices. The posterior distribution on the urn's contents would be derived by computing a weighted average on 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.

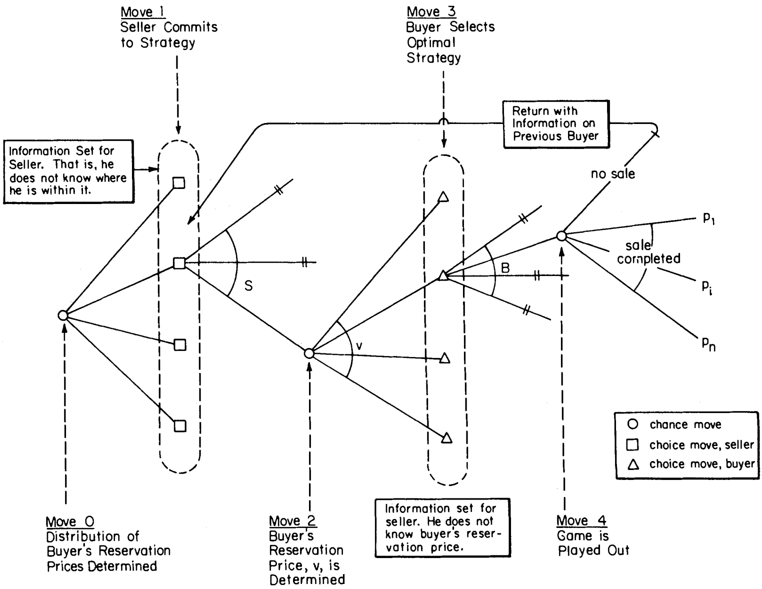

FIGURE I
Seller and Buyer Information

price of α, there is a 50 percent chance that he will quote a price of β and a 50 percent chance that he will request a new buyer. The strategy may allow for learning; for example, it may specify how 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.[4] 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 between two committed strategies can now be conducted. This aspect of the process might be compared to a poker game between two computers. If no sale is made, the process cycles back to move 1, with updated information for the seller.

Possible Bargaining Formats

The decision tree allows for all situations in which the seller, as the more continuing and permanent participant, can make the first commitment, as would often seem reasonable. Either the buyer or

4. 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.

seller can make price offers. For example, as some traditional models posit, they might alternate, with a probability of termination after each refusal. Or the buyer could make bids, with the seller committing himself to probabilities of acceptance depending upon the bid sequence to date. A third possibility has the seller announce a price together with a probability distribution on second-round prices (possibly higher) should the first price be refused, one on third-round prices contingent on second-round refusal, etc.

Some variants of these formulations would represent what is traditionally thought of as haggling. There are two reasons why haggling might prove beneficial. First, it allows the seller to price discriminate. That is, if the seller announces an initial price and a probability less than one of continuing to a lower price, a customer with a high valuation would find the initial asking price preferable to the risky second offer, whereas a potential buyer with a low valuation would find it preferable on average to wait. Second, by adopting some form of discriminatory pricing strategy over time, the seller might be able to gain valuable information about the distribution of reservation prices.

These reasons are insufficient to overcome the primary cost of haggling: It encourages buyers to refuse a high price in the hope of getting a lower one. Our principal result is that if commitment is possible, sellers should always use a fixed-price strategy, whereby a refusing buyer is shown the door and the next buyer is called in. It is reassuring that a "take-it-or-leave-it" strategy is optimal from the standpoint of the seller, since this "explains" the commonly observed behavior of sellers who can make firm commitments.

Relation to the Literature

This analysis builds on the literature on optimal search. In the seminal analysis of information and search, Stigler [1961] alludes to the "higgling process," a reference point for this paper. Much of the ensuing literature on search focuses on an agent's decision whether to accept a present offer or to seek additional quotes *elsewhere*. See Chow and Robbins [1961], Kohn and Shavell [1974], 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. At its heart lies the insight that a seller strategy can be thought of as a mechanism to get buyers to "reveal" their reservation prices. In this respect the paper bears parallels to recent research on optimum auctions by Harris and Raviv [1981a, 1981b], Myerson [1981], and Riley and Samuelson [1981].

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 by the 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 currently in the store. To simplify the discussion somewhat, we assume that money changes hands only if the object is sold.[5] Then the expected return to a response B' can be expressed as follows:

$$\{\text{Expected buyer gain}\} = \{\text{probability sale is made}\} \times [\{\text{reservation wage}\} - \{\text{expected price}\}]. \tag{2}$$

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.

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

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

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

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

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

Utilizing (4), we have

$$[\phi(v,v) - \phi(v',v)] + [\phi(v',v') - \phi(v,v')] = (H(v') - H(v))(v' - v).$$

Since we have defined $B = b_s(v)$ to be the buyer's optimal response, it must be that $\phi(x,v)$ takes on its maximum at $x = v$. Also $\phi(x,v')$ takes on its maximum at $x = v'$. Then each of the bracketed expressions must be nonnegative. Thus, the probability of sale function $H(v)$ must be nondecreasing so we may interpret it as a distribution function.

We now exploit the requirement that $\phi(x,v)$ take on its maximum at $x = v$ to obtain a simple expression for $\phi(v,v) = H(v)(v - \overline{p}(v))$, the maximized expected buyer gain.

Since $\phi(x,v)$ is continuous in v, it follows that the maximand $\phi(v,v)$ is a continuous function. Moreover, under the assumption that $H(v)$ and $\overline{p}(v)$ are both piecewise differentiable,[6] we may express $\phi(v,v)$ in integral form as

$$(5) \qquad \phi(v,v) = \int_0^v \frac{d\phi}{dx}(x,x)\,dx + \phi(0,0).$$

For $\phi(x,v)$ to have a maximum at $x = v$, we have at, points of differentiability,

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

Also, from (4),

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

6. Actually the assumption is not required. In the appendix to Riley and Zeckhauser [1981], it is shown that $\phi(v,v)$ is absolutely continuous. Since it is also nondecreasing, it follows immediately that it may be expressed as in equation (5).

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

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

Substituting into (5), we have finally

$$(6) \qquad \phi(v,v) = \int_0^v H(x)\,dx + \phi(0,0).$$

We are now ready to consider the haggling game from the seller's point. The expected payment by a buyer with reservation value v is just the probability of a sale times the expected price conditional upon there being a sale, that is, $H(v)\overline{p}(v)$. Then substituting from (6) into (2), we see that the expected payment is

$$(7) \qquad H(v)\overline{p}(v) = H(v)v - \int_0^v H(x)\,dx - \phi(0,0).$$

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

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

Integrating the second term by parts, we have finally

$$(8) \quad \text{Expected seller revenue} = \int_0^1 H(v)j(v)f(v)\,dv - \phi(0,0),$$

where

$$j(v) = v - (1 - F(v))/f(v).$$

Since buyers are free to leave 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 that $(0,0) \geq 0$. 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 to the seller its true value. (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

$$\int_0^1 (1 - H(v))\pi(v)f(v)\,dv. \tag{9}$$

Combining (8) and (9) and rearranging, we see that expected total profit is therefore

$$\overline{\Pi} = \left[\int_0^1 \pi(v)f(v)\,dv - c\right] + \int_0^1 H(v)(j(v) - \pi(v))f(v)\,dv. \tag{10}$$

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. A necessary condition for the maximization of expected seller profit is that $H(v)$ be chosen to maximize this increment. Denote it $\Delta\overline{\Pi}$, where

$$\Delta\overline{\Pi} = \int_0^1 H(v)(j(v) - \pi(v))f(v)\,dv. \tag{11}$$

Before discussing the nature of the optimal solution, we should 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 recognize the possibility that the early buyers coming into the store might be more eager and hence have higher reservation values.

If many objects were to be offered for sale, with a constant cost, say A, as opposed to just one in our formulation, the analysis would simplify. We would replace $\pi(v)$ with A in all the equations. The lost future profits are simply the replacement cost of the asset.[7]

7. 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 that 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[8]

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

Expression (11) for expected current profit then becomes

$$\Delta\overline{\Pi} = \int_0^1 H(v)k(v)\,dv. \tag{12}$$

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 nonnegative. It follows that $k(v)$ is negative at $v = 0$ as depicted in Figure II. However, there are no further obvious restrictions on the form 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 $\Delta\overline{\Pi}$ must be nonpositive. Then if search is worthwhile, there must be some smallest right-hand endpoint c such that $H^*(c)$ is positive. This is depicted in Figure II. Given the definition of c, we know that $H^*(a) = 0$. Then we can rewrite $\Delta\overline{\Pi}$ as the following sum of integrals:

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

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

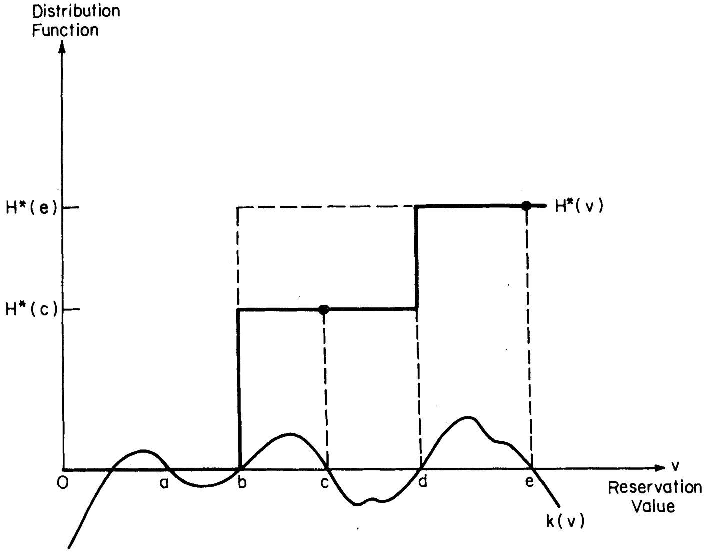

FIGURE II
The Optimal Distribution Function

Since $k(v)$ is negative on (a,b), $\Delta\overline{\Pi}$ 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 nondecreasing. Then $H^*(v) = H^*(c)$ on (b,c).

Now let $H^*(e)$ be the optimal value of H at e, the right-hand endpoint 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

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

Then we can rewrite (13) as

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

If the first integral is negative, $\Delta\overline{\Pi}$ is maximized by setting $H^*(c) = 0$. Since this contradicts the definition of c, the integral must be nonnegative. Then $\Delta\overline{\Pi}$ is maximized by setting $H^*(c)$ as large as possible, that is by 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

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

It follows that if searching for a buyer is optimal, $(\Delta\overline{\Pi} > 0)$, $H^*(c)$ must be equal to 1, hence Proposition 1. The proposition 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.[9]

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 seller knows that the reservation values 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,

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

9. In deriving Proposition 1, it is assumed that the cost of attracting a buyer is independent of the expected buyer gain. An alternative formulation would seek to maximize expected seller revenue subject to the constraint that expected buyer gain

$$\underset{v}{E}\, \phi(v,v) = \int_0^1 f(v) \int_0^v H(x)\, dx\, dv = \int_0^1 (1 - F(v))H(v)\, dv$$

is at least $\overline{\phi}$. Since the integrand is linear in H, the proof above is easily modified to establish that Proposition 1 continues to hold.

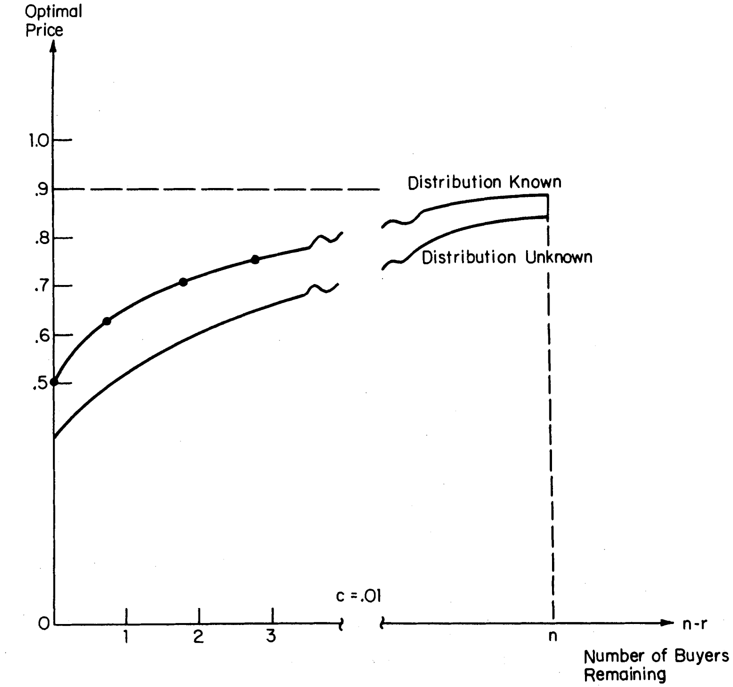

FIGURE III
Optimal Pricing

From Proposition 1 the rth buyer will be offered a price v_r^* satisfying

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

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

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

Substituting (14) into (15) yields a first-quarter 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 III, employing the assumption that v is distributed uniformly ($F_0(v) = v$). If the seller faces the last buyer, the optimal price is 0.5. If there are many buyers remaining, however, the optimal price is close to $1 - \sqrt{c}$.

Price behavior with $F(v)$ unknown. Suppose that the seller begins with beliefs given by the distribution $F_0(v)$, but after he rejects 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 that it is always the case that

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

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 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 $\overline{\pi}_r$, the profit if beliefs remain fixed at $F_0(v)$.

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

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

To satisfy (16), we require that

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

Given (18), it can be confirmed that

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

If results are combined, it follows that

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

and

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

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 III: 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 if some buyers remain.)

While the assumptions used to compare pricing behavior with and without learning may appear mild, they are in fact far from innocuous. Indeed with $F(v)$ unknown the seller may wish to increase his price after rejecting some customers. Suppose, for example,[10] that

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

the seller starts out by believing that with probability ρ the distribution of reservation values is given by

$$F_1(v) = \begin{cases} 0, & 0 \leq v \leq 0.1 \\ 0.1, & 0.1 \leq v < 1 \\ 1, & 1 \leq v, \end{cases}$$

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

$$F_2(v) = \begin{cases} 0, & 0 \leq v < 0.5 \\ 1, & 0.5 \leq v. \end{cases}$$

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 upward to unity and announces a second-round price of just less than unity. This leads to

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

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

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), a buyer with reservation value v may expect a gain of $\alpha v^2/4$.

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 its use will not affect the buyer's decision on the original take-it-or-leave-it offer. Both objectives are achieved by making α arbitrarily small. As α shrinks toward 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.

How the spirit of this formal result is applied in practice is a matter for future deliberations.

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 quotes, making 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 item 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 nondecreasing 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

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

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

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

Define

$$\begin{aligned} G(\hat{v}) &= \hat{G}(p(\hat{v})) \\ &= \text{Prob}\{\text{sale is made to an individual} \\ &\qquad \text{with valuation of at most } \hat{v}\}. \end{aligned}$$

The expected profit from the declining offer strategy is therefore

$$\bar{\pi} = \int_{0}^{1} \left[\int_{\hat{v}}^{1} p(v)f(v)\,dv + \psi(\hat{v})\right] g(\hat{v})\,d\hat{v} - c. \tag{25}$$

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 quote an asking price. Then $G(1) = 1$. Integrating (25) by parts and making use of these endpoint conditions, we therefore have

$$\begin{aligned} \bar{\pi} &= \int_{0}^{1} p(v)G(v)f(v)\,dv + [\psi(\hat{v})G(\hat{v})]_0^1 \\ &\quad - \int_{0}^{1} \psi'(\hat{v})G(\hat{v})\,d\hat{v} - c \\ &= \int_{0}^{1} [p(v)f(v) - \psi'(v)]G(v)\,dv + \psi(1) - c. \end{aligned} \tag{26}$$

From comparison of 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

PROPOSITION 5. If the only information available to the seller is whether 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

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

A sufficient condition for (27) is

$$\psi(v) \leq \pi(v) f(v). \tag{28}$$

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 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 r, is sufficiently high, recall will never be undertaken. When might recall be used? Results differ depending on whether populations are finite or infinite, and whether there is learning about the distribution of buyers' reservation values. There are four cases of interest, as shown in Table I.

TABLE I

RECALL POTENTIALLY BENEFICIAL

	Finite population	Infinite population
No learning	(A) yes	(B) no
Learning	(C) yes	(D) yes

We shall first demonstrate the positive result for (A). Since no learning is a polar case of learning, this implies a positive result for (C) as well. Then we shall examine the negative result for (B). Finally, we present a successful example for (D).

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 and the cost of recall are both zero. With no recall the seller will always announce a sequence of prices beginning with one of the possible values of v and proceeding to another value that is no higher. 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 that 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 $47/64 - 44/64 = 3/64$.[11]

A continuity argument establishes that the recall option remains valuable if search and recall costs are positive but small, including situations where the recall cost exceeds the cost of securing a new buyer. It is worth reiterating that a positive result for (A), a special instance of (C), implies a positive result for (C) as well.

No Learning, Infinite Population (B)

Even when the option to recall is costless, it will never be used in case (B). To simplify matters, we assume that a buyer is recalled at most once. Suppose that the seller's strategy implies that there is some anticipated nonzero probability π_1 that buyer 1 will later be recalled. Next consider the following modifications in the seller's strategy. Instead of allowing buyer 1 to leave the store according to the original rules, he is "recalled" immediately with probability π_1. Clearly this is equivalent from the viewpoint of buyer 1. Lowering the probability of later recall to zero will, in general, increase the probability of recall for some other buyers. We then suppose that these probabilities are simultaneously adjusted downward so that all buyers' expected gains are unaffected by the change. It follows that all buyers' strategies are unaffected so that expected seller revenue is unchanged.

But now buyer 1 is in the store only once, and from Proposition 2 it follows that the seller can do no better than offer this buyer a single take-it-or-leave-it price. Since exactly the same logic applies for each of the buyers, the recall option is indeed valueless. (By continuity, recall is likely to offer only negligible benefit for large finite populations.)

11. 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 for 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

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

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

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.

Learning, Infinite Population (D)

Might the presence of learning turn the negative (B) result positive? Interestingly, the answer is yes, as a simple example makes clear. Assume that the seller and the buyers know that all buyers have identical reservation prices, and that there are three equally likely situations, that the price is 1, that it is 0.6, and that it is 0.2. The cost of securing a buyer is $c < 0.3$. The cost of recall is $r < c/3$.

The optimal strategy without recall is to quote a price of 1 minus a hair to the first buyer; 0.6 minus a hair to the second, should the first refuse; and a hair less than 0.2 to the third, should the second refuse as well. Leaving split hairs aside, the expected revenue is

$$\tfrac{1}{3}(1) + \tfrac{1}{3}(0.6) + \tfrac{1}{3}(0.2) = 0.6,$$

while the expected search cost is

$$c + \tfrac{2}{3}(c) + \tfrac{1}{3}(c) = 2c.$$

With recall possible, a superior strategy is available. The seller can quote a price of 1 minus a hair to the first two buyers. Buyer 1 is told that if neither accepts the offer, he will be recalled and quoted a price of 0.6. Buyer 2 is told that if the object is not sold at a price of 1 or 0.6, he will be recalled and quoted a price of 0.2. Note that if the second buyer's reservation price is 1, he will certainly accept the quote of 1 since he knows that buyer 1's reservation price is also 1 and he will surely accept the quote of 0.6. But, recognizing this, buyer 1 will accept the quote of 1. Better a certainty of a smidgen than a zero-probability prospect of a large gain.

Expected revenue is therefore exactly as in the absence of recall, while the expected cost of search drops to

$$c + \tfrac{2}{3}(c + r) + \tfrac{1}{3}(r) = \tfrac{5}{3}(c + r).$$

It may seem puzzling at first that learning turns a negative result positive. An intuitive hypothesis to explain this result might be that sellers can 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.

If recall is as expensive as securing a new customer, it cannot be optimal, as a simple thought experiment reveals. 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 the two possible situations. If Y had accepted our earlier price when X refused, then Y is to be preferred to X now, for he will surely accept. If Y had refused the offer X received, then our knowledge of Y is the same as our knowledge or X. Since we are indifferent in one case and prefer Y in the other, buyer Y should be our choice.

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.

V. Extensions and Generalizations

Extensions and generalizations of this work could come in many areas. They could allow for the following: (1) multiple items for sale, (2) the seller's strategy affecting the distribution of buyers, (3) risk-averse buyers, (4) consideration of seller's abilities to make commitments, (5) alternative market structures where there may be competitive elements, (6) buyers moving first or simultaneous moves, and (7) examination of anecdotal and statistical evidence on the actual pricing behavior of firms. Here we shall just comment on the first four of these areas.

Multiple items for sale. Our models apply, albeit with a bit more complexity, when there are any number of items for sale and each potential buyer wishes to purchase only one unit. With no learning, expected profits with N objects is simply N times that expected with 1; the identical fixed-price strategy should be employed. Situations with increasing supply costs turn out to have many of the properties of depletable resource models. The cost of each item is its own marginal cost plus a shadow price indicating the (discounted) increased cost of future items sold.

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. A store's committed pricing strategy affects who comes to shop. 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 realistic model would recognize that the cost of securing a buyer is a function of the seller's strategy. As indicated in footnote 9, if expected buyer gain is held at some arbitrary level $\bar{\phi}$, maximized incremental profit, net of the cost of securing a buyer, $\Delta\overline{\Pi}(\bar{\phi}) - c$, is again achieved by utilizing a fixed price policy $\bar{p} = p(\bar{\phi})$. Therefore, with c also a function of $\bar{\phi}$, the seller chooses $\bar{\phi}$ and hence the take-it-or-leave-it price $p(\bar{\phi})$ to maximize incremental profit.

Risk-averse buyers. When buyers are risk-averse, the fixed-price strategy is no longer optimal for the seller. A more favorable scheme would be to ask the buyer to make a bid for the object, with the understanding that the seller will accept the bid on a probabilistic basis. The higher the price offered, the more likely the bid will be accepted. The strongly risk-averse buyer would bid just below his valuation v. A lower bid would offer him some chance of a larger gain, but would entail a greater risk of refusal. Thus, the seller can extract almost all the consumer surplus by exploiting the buyer's risk aversion.[12]

Possibilities for commitment and haggling. Observation suggests that large stores with established reputations are most likely to employ fixed-price strategies. Although a variety of alternative explanations are possible, it struck us that these are the stores best able to make firm commitments to fixed prices.

In this context we observe that one advantage of brand-naming or resale price maintenance may be to enable a seller to commit himself to a strategy he otherwise would not adopt. 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 might make it difficult to find the product in the marketplace.

Suppose that a motorist walks into a country antique store on a back road. If 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, or incur the wrath of some manufacturer. The seller's inability to commit himself to a fixed-price strategy virtually guarantees haggling unless, of course, the buyer finds such a process extremely unattractive, or social custom or tradition renders haggling behavior by the seller unacceptable. With just anecdotal evidence to support our theoretically guided insights, we conjecture that haggling will be much more common when sellers encounter particular buyers infrequently and do not have reputations to maintain or establish.

12. For further discussion of Selling schemes with risk-averse buyers, see Maskin and Riley [1982].

VI. Conclusion

We have provided a strong theoretical justification for the pricing strategy found in a wide variety of markets. Prices are established, and buyers can accept them or seek to buy elsewhere. This fixed price strategy—which requires that the seller be able to make commitments about his behavior—is optimal in comparison to any other, including all forms of buyer involvement, such as quoting offers. Where seller commitment is impossible, perhaps because reputations are difficult to establish or market encounters are highly occasional, sellers should be expected to haggle.

University of California, Los Angeles
Harvard University

References

Arrow, Kenneth, and Michael Rothschild, "The Existence of Equilibria Under Search," 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*, Vol. I (Berkeley: University of California Press, 1961), pp. 93–104.
Harris, M., and A. Raviv, "A Theory of Monopoly Pricing Schemes with Demand Uncertainty," *American Economic Review*, LXXI (June 1981), 347–65.
——, and ——, "Allocation Mechanisms and the Design of Auctions," *Econometrica*, (1981), 1477–99.
Kohn, M. G., and S. Shavell, "The Theory of Search," *Journal of Economic Theory*, IX (Oct. 1974), 93–123.
Lippman, Stevan A., and John J. McCall, "The Economics of Job Search: A Survey," *Economic Inquiry*, XIV (June 1976), 155–89.
Maskin, Eric, and John G. Riley, "Auctioning an Indivisible Object," Discussion Paper, Kennedy School of Government, Harvard University, 1980.
——, and ——, "Optimal Auctions with Risk Averse Buyers," M.I.T. Discussion Paper #311 (revised), September 1982.
Myerson, R. B., "Optimal Auction Design," *Mathematics of Operations Research*, VI (March 1981), 58–73.
Pratt, John, David Wise, and Richard Zeckhauser, "Price Variations in Almost Competitive Markets," this *Journal*, XCIII (May 1979), 189–211.
Riley, John G., and William F. Samuelson, "Optimal Auctions," *American Economic Review*, LXXI (June 1981), 381–92.
Riley, John G., and Richard Zeckhauser, "Optimal Selling Strategies: When to Haggle, When to Hold Firm," Discussion Paper 88D, Kennedy School of Government, Harvard University, April 1980, revised version, March 1981.
Rothschild, Michael, "Models of Market Organization with Imperfect Information: A Survey," *Journal of Political Economy*, LXXXI (Nov./Dec. 1973), 1283–1308.
Stigler, George J., "The Economics of Information," *Journal of Political Economy*, LXIX (June 1961), 213–25.