NBER WORKING PAPER SERIES

RETAIL PRICING, THE TIME DISTRIBUTION OF TRANSACTIONS, AND CLEARANCE SALES

Edward P. Lazear

Working Paper No. 1446

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 1984

Many useful comments were provided by Dennis Carlton, John P. Gould, Robert Hall, Brian Hanessian, Michael Mussa, Peter Pashigian, Sam Peltzman, Melvin Reder, and Sherwin Rosen, and George Stigler. I am especially indebted to Kevin M. Murphy, who caught a number of errors in the first draft. This research was supported by the National Science Foundation. The research reported here is part of the NBER's research program in Labor Studies. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

NBER Working Paper #1446
September 1984

Retail Pricing, the Time Distribution of Transactions, and Clearance Sales

ABSTRACT

Sellers of new products are faced with having to guess demand conditions to set price appropriately. But sellers are able to adjust price over time and to learn from past mistakes. Additionally, it is not necessary that all goods be sold with certainty. It is sometimes better to set a high price and to risk no sale. This process is modeled to explain retail pricing behavior and the time distribution of transactions. Prices start high and fall as a function of time on the shelf. The initial price and rate of decline can be predicted and depends on thinness of the market, the proportion of customers who are "window shoppers," and other observable characteristics. In a simple case, when prices are set optimally, the probability of selling the product is constant over time. Among the more interesting predictions is that women's clothes may sell for a higher average price than men's clothes, given similar cost, even in a competitive market. Another is that the initial price level and the rate of price decline are positively related to the probability of selling the good. Other observable relationships are discussed.

Edward P. Lazear
Graduate School of Business
University of Chicago
1101 East 58th Street
Chicago, IL 60637

A large department store wants to sell a "one-of-a-kind" designer gown. Although the manager has some idea about the price that the gown can command, there is generally some guesswork associated with the process. How should he choose his initial price? If the dress does not sell at that price after the first few weeks on the rack, he has the option of trying a new price. When does he change price and how does the new price relate to the old? How does the decision depend on the characteristics of the gown and on conditions in the clothing market? Is the gown more likely to sell during the first few weeks on the rack, is the pattern of expected transactions smooth over time, or do most sales occur later, when the seller is most frantic about getting rid of the gown?

Firms face a very similar problem when they market a new product that is not unique. Imagine a computer firm that introduces a new model. How should the firm select a time path of prices for the computer, recognizing that the company is not completely certain about the market for the new item? Is it best to start with a high price and lower it over time, or should it do the reverse? Are there any circumstances for which a constant price over time is the appropriate strategy? How is its price contingent on the number of sales made in the first days that the product is on the market? Are the number of transactions likely to be larger at the beginning and then taper off, or might they be smooth over time? How do these patterns vary with the characteristics of the goods and the nature of the buyers? When does the firm announce a "clearance sale," which is an attempt to move merchandise at a price (significantly) below its original price?

This paper provides a simple framework that permits the analysis of these issues. It is an attempt to explain pricing and transaction patterns over time. A number of market phenomena are explained. Among the more interesting ones are:

1. Differences in pricing behavior by characteristics of the goods.

For example, it is commonly alleged that women's clothes are more expensive than men's clothes, given cost conditions. Also, "designer" items often carry extremely high initial prices, which fall rapidly if the good does not sell. If men's suits do not exhibit such volatile price behavior, what might explain this pattern? What does "fashion" have to do with this and is there an objective definition of "fashion" that yields predictions?

2. Prices may be more or less variable depending upon the thinness of the market.

For some items, say, a $2 million mansion, transactions are relatively rare events. How does the pricing of these infrequently traded items differ from that of goods that turn over often?

3. Do strategies vary with the uniqueness of the good?

Designer dresses and Picasso paintings are unique. One and only one item of its exact type is for sale. But there are many copies of a new computer model and the sale of one machine does not preclude the sale of another identical one to another buyer. How does pricing and selling strategy differ in these two cases?

4. Price reduction policies as a function of time on the shelf.

Some famous department stores have an announced policy of halving the price of an item for each week that it remains on the floor. Such "bargain basement" behavior can be predicted and the price cutting rule can be specified as well. When is a rigid rule of this sort an optimal pricing policy?

The goal is to relate these pricing and selling strategies to underlying, observable characteristics of the market in order to explain the differences. Factors relating to the heterogeneity of the goods, the heterogeneity of buyer preferences, and search costs are discussed.

The idea behind the model is that the ability to sell goods over time allows richer strategies for two reasons. First, if the good does not sell during the first period, the seller still has a chance of selling it during the next period. Second, the outcome of the first period provides the second-period seller with additional information. The amount and nature of that information depends on the characteristics of the market and the number and attributes of the buyers. This can be modeled in a very easy way and all of the questions posed above can be addressed.[1]

A distinction that plays an important role is the one between the expected selling price and the expected revenue associated with a good. The former is the price, conditional on a sale. The latter takes into account that a sale does not always occur. It generally pays to set prices in a way that sometimes leaves the good unsold at the end of the period.

The most important point to bear in mind is that this is a model of "retailing." Retailing, as defined in this paper, describes a selling pattern with an announced price that is maintained for some period of time. The seller agrees implicitly to sell to the first person (or in the case of non-unique goods, to any person) who comes along and is willing to pay that price. The good sells at the stated price. No haggling occurs and auctions, which pit one buyer against another, are not held. Since the analysis and some results bear a close relation to the auction literature, some comparisons are made below. Although the retailing paradigm is taken as given and exogenous, some attempt to explain why retailing is used over other selling schemes is presented in that section as well.

The design of the paper is to start with the simplest model and then to introduce complications as necessary to explain the data. The effects of market competition and strategic behavior of buyers are all considered in turn.

I. Intertemporal vs. Single-Period Pricing

The ability to change price after the first attempt to sell the product fails produces a richer set of strategies and changes the problem facing the firm. To see this, let us begin with the most basic characterization of the firm's pricing problem in a single-period context.

A. A One-Period Model

Suppose that the firm will encounter one and only one buyer who is willing to pay V for the good, but no more. The firm does not know V with certainty, but has some prior notion of the density of V denoted $f(V)$ with distribution function $F(V)$.[2] (The prior may be based on an examination of the selling prices of similar goods, but for now, its source is unimportant.) The risk-neutral firm's problem is to maximize expected profits or

$$\max_{R} R[1 - F(R)] \tag{1}$$

where R is the price and $1 - F(R)$ is the probability that V exceeds R so that a sale is made. For the purposes of expositional simplicity, suppose that the prior on V is uniform between zero and one. Then $F(R) = R$ so that the optimum is at $R = 1/2$, yielding expected profits of $1/4$.

An alternative formulation is that the firm has a large number of similar, but not identical items that it wishes to sell. It knows that the distribution of demand prices is given by $f(V)$, but it does not know which items correspond to high values and which to low. An example is a line of dresses, which come in different colors or have different trim. Ex ante, the seller does not know whether it is the yellow or the red one that has $V = 1$. The one-period pricing rule is again, set $R = 1/2$.

B. A Two-Period Model

Now suppose instead that if the good does not sell during the first period, the seller faces another buyer during the second period who is identical to the one he saw during the first period. The firm now has two chances to sell the good. Furthermore, the failure of the good to sell in period 1 at price R_1 tells the seller something about the V. In this simple case, it implies that $V < R_1$, because if $V > R_1$, the good would have sold.[3]

Using Bayes' Theorem, this implies that the posterior distribution which the firm carries into period 2 is uniform between 0 and R_1, so that $F_2(V)$, the posterior distribution, $= V/R_1$. The choice of R_1 affects the problem in two ways. First, it affects the probability of a sale in period 1. Second, it determines what the firm can infer from no sale. For example, if $R_1 = 1$, then the fact that the good did not sell is uninformative because the firm was certain that $V < 1$ at the outset. Similarly, $R_1 = 0$ is certain to result in a sale during the first period so there is no learning that occurs with this choice either.

The firm's problem then is to choose R_1 and R_2 where R_2 is the price that the firm tries in period 2, given that the good did not sell in period 1. In this example, the good is unique so that a sale in period 1 eliminates any concern about period 2. That problem can be written as:

$$\max_{R_1, R_2} \; R_1[1 - F(R_1)] + R_2[1 - F_2(R_2)]F(R_1) \; . \tag{2}$$

The first term is the price charged in period 1 times the probability that the good sells in period 1. The second term is the price charged in period 2 times the probability of a sale in period 2 at that price, given the information from period 1, times the probability that the good does not sell in period 1.

It is instructive to think of this as a dynamic programming problem and to consider the firm's optimal strategy in period 2, given that the good did not sell in period 1 at the price R_1. The firm's problem in period 2 is

(3) $$\max_{R_2} R_2[1 - F_2(R_2)] \ .$$

Given the prior on V, this becomes

$$\max_{R_2} R_2(1 - R_2/R_1)$$

which has an optimum at $R_2 = (R_1)/2$.

This makes obvious intuitive sense. For any given R_1, if the good did not sell during the first period, then the seller can rule out the possibility that $V > R_1$. The distribution that the seller uses in period 2 is uniform between zero and R_1 so the second period's problem is equivalent to the one facing a firm with only one period to sell and with a prior between 0 and R_1. The solution to that problem is to select $R_2 = (R_1)/2$.

Thus, for any given R_1, if the good remains unsold after one period, the rule is to cut the price in half next period. (This halving is specific to the assumed distribution, of course.) Substitution of $R_2 = (R_1)/2$ into (2) and maximizing with respect to R_1 yields[4]

$$R_1 = 2/3$$
$$R_2 = R_1/2 = 1/3 \ .$$

This illustrates a number of important points. First, prices fall over time. A retailer puts a gown on the market at a high price ($R_1 = 2/3$), hoping that it will sell at that price. If it does not sell, he can revise his price downward during the next period. The reverse pattern would never be optimal because once the gown sold at 1/3, the seller has eliminated the chance that he will get 2/3 for it. (If $V > 2/3$, then it exceeds 1/3 as well so period 2 becomes irrelevant since no gown that will sell in period 2 is ever available after period 1.)[5] Stated alternatively, the gowns in the "prettiest" colors with $V > 2/3$ sell in period 1. The seller must revise downward his opinion of the value of remaining gowns. The distribution at the beginning of period 2 has lower value gowns than those at the start of period 1 because the best ones have already been picked off.

This abstracts from any investment in brand name recognition associated with charging a lower initial price. For example, new firms frequently charge lower prices than their rivals to induce customers to try the new product. The difference between the observed price and the optimal one as calculated in this problem can be thought of as advertising and is ignored throughout. It also abstracts from contagion or network effects. The value, V, is assumed to be independent of the number of others who have similar items.[6]

Second, the comparison with the one-period solution is interesting. There, the solution was to set price equal to 1/2. Now, because a disappointed seller has another chance, a first-period price that exceeds 1/2 is justified.[7] If he wanted to, he could always select a price of 1/2 during the second period because he still has one chance left. Of course, given what he has learned from period 1, a price of 1/2 is no longer optimal in round 2. So the prices charged straddle the one period optimum.

Furthermore, expected profits are higher as a result of having a second chance. In the one-period problem, expected profits were 1/4. In the two-period problem, expected profits are 1/3 (substitute $R_1 = 2/3$, $R_2 = 2/3$ into (3)). This is because the expected probability of a sale is higher in the two-period problem. The expected probability that a sale occurs in one of the two periods is

$$(1 - F(R_1)) + (1 - F_2(R_2))F(R_1)$$

or

$$1 - 2/3 + [1 - (1/3)/(2/3)](2/3) = 2/3 \ .$$

In the one-period problem, the expected probability of a sale was $(1 - F(R))$ $= 1/2$. (The expected selling price is the same in both cases).[8]

C. Heterogeneous Consumers and Thin Markets

The previous problem was made simple because the inference problem was so trivial. If the good did not sell during the first period, the firm knew with certainty that it overpriced the good. In reality, other factors make the inference problem more difficult. Specifically, two factors are important. The first is the number of customers who come into the store during the first period. Intuitively, if only a few customers arrive during the first period, the firm should be less certain about its inference than if a large number examine the good and reject it at price R_1. Second, heterogeneity among consumers may be important. If some consumers are willing to pay V, while others will pay an amount below the firm's reservation price, then the problem is more complicated. The good might not have sold not because the price was too high, but because that period's customers were all of the wrong type.

This can be parameterized as follows. Suppose that in period 1, N "customers" examine the good. Of those, a proportion P are just "shoppers" whose value of the good is less than the seller's reservation price, and $1 - P$ are "buyers" who are willing to pay V. As before, V is unknown to the seller and his goal is to select R_1 and R_2 to maximize profits, given his prior beliefs on V. In what follows, "customers" refers to the total number of individuals who inspect the good, "buyers" refers to the subset with value equal to V, and "shoppers" refers to the subset with value equal to zero. (Important is that a given individual does not know whether he is a buyer or shopper until he has inspected the good. No one who knew that he was a shopper would ever bother to look.[9])

The problem is similar to the one in (3) except for two points: First, $F_2(V)$ is different here. Second, expected sales depend on P. More formally, the seller wants to maximize

(4) $\underset{R_1, R_2}{\text{Max}}\ R_1(\text{Prob. sale in 1}) + R_2(\text{Posterior prob. sale in 2})(\text{Prob. no sale in 1})$

Now, the probability of a sale in period 1 is

$$(1 - F(R_1))(1 - P^N)$$

because the probability that <u>every</u> customer is a shopper is P^N so that $1 - P^N$ is the probability of encountering at least one buyer. It only requires one buyer to make the sale as long as $R_1 < V$. Similarly, the posterior probability of a sale in period 2 is

$$(1 - F_2(R_2))(1 - P^N)$$

and the probability of no sale in period 1 is

$$1 - [(1 - F(R_1))(1 - P^N)] \ .$$

It is now necessary to derive $F_2(V)$. Bayes' Theorem states that the posterior probability is proportional to the probability of the sample, given the parameter, times the prior probability of the parameter. The sample in this case is the observation that no one bought during period 1. For $V < R_1$, the probability of no purchase is 1. For $V > R_1$, there is only one reason why the good did not sell during period 1 and that is that all customers were shoppers. This happens with probability P^N.

It is easy to show that the normalization required to make the integral of the density function equal to 1 is

$$1/[R_1(1 - P^N) + P^N]$$

so that the density is given by

(5) $$f_2(V) = \frac{1}{R_1(1 - P^N) + P^N} \quad \text{for} \quad V \leq R_1$$

$$= \frac{P^N}{R_1(1 - P^N) + P^N} \quad \text{for} \quad V > R_1 \ .$$

Integrating this yields the distribution function

(6) $$F_2(V) = \frac{V}{R_1(1 - P^N) + P^N} \quad \text{for} \quad V \leq R_1$$

$$= \frac{R_1 + P^N(V - R_1)}{R_1(1 - P^N) + P^N} \quad \text{for} \quad V > R_1 \ .$$

To obtain the probability of a sale in period 2, $(1 - F_2(R_2))$ must be multiplied by $(1 - P^N)$ since that is the probability that at least one buyer is encountered in the group of customers.

Substitution of these expressions into (4) yields the following maximization problem:

(7) $$\max_{R_1, R_2} \ R_1[1 - P^N - (1-P^N)R_1] + R_2[P^N + (1-P^N)R_1 - R_2](1 - P^N) \ .$$

As before, it is instructive to solve this as a dynamic programming problem, deriving the optimal R_2 for any R_1, given that period 2 is reached. This problem is written

(8) $$\max_{R_2} \ R_2(1 - F_2(R_2))(1 - P^N)$$

or

$$\max_{R_2} \ R_2\left[1 - \frac{R_2}{R_1(1 - P^N) + P^N}\right](1 - P^N) \ .$$

Differentiating with respect to R_2 and setting the derivative equal to zero yield the optimum R_2 which is given by:

$$R_2 = \frac{1}{2}\left[R_1(1 - P^N) + P^N\right] . \tag{9}$$

When $P^N = 0$, the problem in (7) and (8) reduces to the simpler problem, which is a special case, considered earlier. Indeed, substitution of $P^N = 0$ into (9) yields the earlier solution that $R_2 = (1/2)R_1$.

Given the solution for R_2 in terms of R_1, (8) can be rewritten as

$$\max_{R_1} R_1[1 - P^N - (1 - P^N)R_1] + \frac{1}{4}(1 - P^N)[P^N + (1 - P^N)R_1]^2 \tag{10}$$

which yields the solution

$$R_1 = \frac{2 + P^N(1 - P^N)}{4 - (1 - P^N)^2} . \tag{11}$$

When $P^N = 0$ so that the problem reduces to the simple one, the solution is again $R_1 = 2/3$ and $R_2 = (1/2)(2/3) = 1/3$. As P^N goes to 1, however, the solution goes to $R_1 = 1/2$ and from (9), $R_2 = 1/2$. That is, as P^N goes to 1, prices remain constant over time.

The intuition behind this result is straightforward. When $P^N = 0$, all customers are buyers (there are no window shoppers) so the inference problem becomes perfect. If the good is left on the shelf after the first period, it can only be because the good was priced higher than V. Therefore, all $V > R_1$ can be ruled out. But as P^N approaches 1, almost all of the customers are merely shoppers. Thus, little can be inferred from the fact that no one bought the good after the first period. Even if R_1 were less than V, there is a very good chance that the good would still remain on the shelf after one period in this climate of browsers. Under these circumstances, having two consecutive periods is no different from having two independent one-period problems, since nothing is learned from the first period. This implies that the solution to the single period problem, namely price = 1/2, applies in both periods.[10]

From (11), it can be shown that

$$\frac{\partial R_1}{\partial P^N} = \frac{-4P^N + (1 - P^N)(2P^{2N} - 1)}{[4 - (1 - P^N)^2)]^2} < 0 \tag{12}$$

and that

$$\frac{\partial R_2}{\partial P^N} = \frac{1}{2}\left[1 - R_1 + (1 - P^N)\frac{\partial R_1}{\partial P^N}\right] > 0 . \tag{13}$$

The implication is that when P^N is small prices start higher and fall more rapidly with time unsold. For P^N close to 1, prices tend to be constant over time. Now, P^N has at least two real world interpretations.

First, N is the number of customers per period of time. As N increases, P^N gets small so that as N increases, prices start high and fall more rapidly. When there are a lot of customers per period, there is more information contained in the fact that the good did not sell so that the strategy moves toward that used when perfect inference is available. On the other hand, if N is small, less is learned from the fact that the good remains unsold after one round. This implies that the prices of goods in thin markets should start lower and fall less rapidly (relative to the prior distribution) than prices of goods where markets are dense.

Consider, for example, the problem of selling a house. Suppose there are two different types of houses. One is a $2 million mansion. Such houses turn over very infrequently and there are very few buyers. Another is a $50,000 high-rise condominium in a building where one of the 300 apartments is sold weekly. The implication of this section is that prices of mansions should be less sensitive to time on the market than prices of condominiums. The reason is that the owner of the mansion cannot infer that his house was overpriced from the fact that it has been on the market for two months without selling.

There are very few potential buyers of mansions. But the owner of the apartment can quickly and precisely infer that if the apartment did not sell it is not because he encountered few buyers, but instead because it was overpriced. Information comes with each genuine buyer and there are fewer of these per period of time in the case of mansions. This suggests that prices of lower quality goods adjust more rapidly to time on the market during which the good remains unsold.

The sale of a house does not quite fit this model, since the process is one of haggling over price, rather than strict retailing. Still, the intuition of the example is appealing. Goods for which the markets are thin have more rigid prices; "clearance sales" are less common. How is thin defined? Since N is the number of customers that any one seller faces, thinness must be defined in some relative sense. Probably the most easily measured aspect of thinness relates to the transactions per unit of time. Consider the house example. If there are 100 houses of the low-priced variety and 5 of the high-priced variety, then in equilibrium, 100 families live in the former and 5 in the latter. Thinness would be the same unless each of the 100 turns over more frequently per unit of time. Suppose that those who live in the low-priced type move twice as often as the high-price residents. Then the number of customers that visit the low-priced houses per unit of time exceeds that at the high-priced houses so an unambiguous measure of thinness can be obtained.

The second interpretation of P^N relates to search cost and information. For a given N, P is the proportion of customers who have a purchase price below the seller's reservation price (in the example above, it was zero). If customers have much information about the good before they inspect it, then few shoppers will show up and all of the customers will be buyers. Consider wholesale versus retail buyers. It is possible (although not obviously true) that purchasers in the wholesale market have better prior information than those in the retail market. If true, this implies that wholesale prices fall more rapidly with inventory time than retail prices because the seller at the wholesale level can infer more than the seller at the retail level about his pricing policy. What should be true under any circumstances, is that for a given number of customers, an increase in the proportion of those who do not buy after having examined the good reduces the speed with which prices fall. Both of these variables, the number of customers and the proportion who do not buy after looking, are observable, at least conceptually.

Related to this is the idea that search costs are important in determining the speed with which prices fall as a function of time on the shelf. Consider a good for which search is costly, for example, a piece of land in the middle of Alaska. For a given number of customers, a very small proportion will be window shoppers. Because inspection is so expensive, most who inspect the good are likely to be buyers rather than shoppers. As such, the seller of that parcel of land can infer a great deal from the decision by any customer not to purchase the land. Thus, the listing price of the land should drop rapidly each time a customer opts against purchase.

Contrast this with a house in the middle of Chicago. The proportion of shoppers-to-buyers is much higher here because search costs are low. Even individuals who are likely to value the good at zero rather than V may consider taking a look to be certain. Thus, less can be inferred from a given customer's decision not to buy the house. This implies that price is less sensitive to N in Chicago than it is in Alaska. Of course, for a given period of time, N, the number of customers, is likely to be higher for the house in Chicago than for the land in Alaska. This means that prices may fall more rapidly with time even though not with N for the house in Chicago. Both time on the shelf and N are observable.

Additionally, goods for which repeat purchases are made are likely to have informed customers and sellers. The prior on V is tight and P is likely to be small. More is said on this below.

D. Competition

Competition has not been mentioned in the preceding analysis. Implicit in the fact that goods must be examined to assess value, is some ex post monopoly power that is created by the imperfect information. Still, most of the goods that are the subject of this paper are sold in what would normally be thought of as a competitive environment. How does competition affect the story?

There are two ways that competition affects the analysis. The first is that the existence of competition is likely to alter the firm's prior. Even with uncertainty in the world, a seller in a competitive market might reasonably assume that the distribution of V lies to the left of that for a seller in a monopolistic market for the standard reasons. If a perfect substitute exists, and if all customers know its price, then no V in excess of that price is feasible.

There is a more interesting way to think of competition in this context. What is meant by competition is that even though firms may not be able to compete directly on the exact good, (e.g., only one store may carry one particular designer's latest dress), the customer's search is usually for a good among a class of goods rather than for any one particular item. (E.g., a woman is looking for a new party dress, not for Yves St. Laurent's latest.) In this context, a key parameter, P^N, has been assumed to be exogenous. But competition tends to make P^N endogenous. Ex ante, stores are identical even though there may be some ex post differences. For example, a street may contain ten art galleries, all of which look identical from the outside. But since paintings are unique, each gallery has a different selection of paintings. Positive profit in the gallery business causes more firms to enter so that the N that each firm sees falls until expected profit equals zero.[11]

More formally, substitution of (9) and (11) into (7) yields the firm's revenue function in terms of P^N. Differentiation with respect to P^N yields the obvious conclusion that expected revenue falls as P^N rises.[12] (When $P^N = 1$, revenue is zero because all customers are shoppers. When $P^N = 0$, revenue is maximized because all customers are buyers and, so long as N equals or exceeds 1, a sale is certain for $R_1 < V$.) One way to think about competition is that firms enter the industry, reducing the per period flow of N that any one firm faces until profits are driven to zero. Since expected revenue falls when N falls (it moves inversely with P^N), competition reduces N until expected revenue equals costs.[13]

The competitive equilibrium yields some testable predictions on pricing. To the extent that competition reduces profits to zero by decreasing N and increasing P^N, competition also affects the choice of initial price and its fall over time. From (12), it is clear that increases in P^N as a result of competition reduce R_1. From (13), the same increase in P^N implies an increase in R_2. This implies that for the same prior on V, firms in competitive markets choose lower initial prices and reduce them less rapidly than firms in monopolistic markets. Put differently, this suggests that prices are more sensitive to inventory time in monopolistic markets than in competitive ones.

This raises a more general point. Prices of, say, nonperishable items in supermarkets do not seem to exhibit much time variation at all. How does the model explain this phenomenon? Competition among supermarkets, which results in higher P^N, is capable of providing one explanation for the phenomenon. But a more basic force is at work.

Recall that the reason that prices fall over time is that learning has taken place during the relevant period. The amount of learning that can occur depends on the dispersion in the prior on V. But if the same good has been sold for a long period of time, the relevant prior is likely to be extremely tight. Thus, little learning occurs and prices remain rigid as a result. Both factors, the length of the time horizon and the amount of dispersion in the prior, are analyzed more rigorously below.

E. Observable Time Patterns of Price and Quantity

The theory yields predictions of pricing behavior as a function of three factors: the number of customers, N; the proportion of customers who are shoppers rather than buyers, P; and the firm's beliefs about the market, parameterized through the prior on V. With the exception perhaps of the last of the three, these variables are observable, at least in theory. However, it is likely to prove quite difficult to obtain information on P and N.

Quite aside from data considerations, it is useful to be able to relate price time paths and quantity time paths to some observable characteristics, as well as to each other. The relation of R_1 and R_2 to P and N has already been discussed. Recall that as P^N goes from zero to one, R_1 moves from 2/3 to 1/2 and R_2 moves from 1/3 to 1/2: R_1 falls and R_2 rises so the ratio of R_1 to R_2 falls as P^N increases, i.e., as inference becomes more difficult.

The pattern of expected transactions over time is somewhat less intuitive. The probability of a sale in period 1 is

$$\text{Prob. sale in 1} = 1 - P^N - (1 - P^N)R_1 \tag{14}$$

$$= (1 - P^N)(1 - R_1) \; .$$

The (unconditional) probability of a sale in period 2 is

$$\text{Prob(sale in } 2 \mid R_2\text{, no sale in 1)} \cdot \text{Prob(no sale in 1)} .$$

This is the second term of (7) without the price R_2 as a scalar. Substituting (9) into this part of (7) yields

$$\text{Prob. sale in } 2 = \frac{1}{2} [P^N + (1 - P^N)R_1](1 - P^N) \tag{15}$$

Division of (14) by (15) gives the ratio of sales in period 1 to those in period 2. That ratio is

$$\frac{2(1 - R_1)}{P^N + (1 - P^N)R_1} .$$

After substituting (11) into this expression, the ratio reduces to 1. That means that the unconditional probability of a sale in period 1 is equal to that for period 2. Expected sales are smooth over time.

Additionally, since expected sales are equal in each period, the probability that the good sells is given by twice the probability that it sells in period 1 or by

$$\text{prob. of sale} = 2(1 - P^N)(1 - R_1) .$$

This varies with P^N as

$$\frac{d(\text{Prob. of sale})}{d\ P^N} = -2(1 - R_1) - 2(1 - P^N) \frac{\partial R_1}{\partial P^N} .$$

After substitution of (12), it can be shown that

$$\frac{\partial(\text{Prob. of sale})}{\partial\ P^N} < 0 .$$

Table 1 simulates some values.

Table 1

Expected Price and Quantity Relationships

P^N	R_1	R_2	Expected Revenue	Probability of Sale	R_1/R_2
0.0000	0.6667	0.3333	0.3333	0.6667	2.0000
0.0500	0.6610	0.3390	0.3220	0.6441	1.9500
0.1000	0.6552	0.3448	0.3103	0.6207	1.9000
0.1500	0.6491	0.3509	0.2982	0.5965	1.8500
0.2000	0.6429	0.3571	0.2857	0.5714	1.8000
0.2500	0.6364	0.3636	0.2727	0.5455	1.7500
0.3000	0.6296	0.3704	0.2593	0.5185	1.7000
0.3500	0.6226	0.3774	0.2453	0.4906	1.6500
0.4000	0.6154	0.3846	0.2308	0.4615	1.6000
0.4500	0.6078	0.3922	0.2157	0.4314	1.5500
0.5000	0.6000	0.4000	0.2000	0.4000	1.5000
0.5500	0.5918	0.4082	0.1837	0.3673	1.4500
0.6000	0.5833	0.4167	0.1667	0.3333	1.4000
0.6500	0.5745	0.4255	0.1489	0.2979	1.3500
0.7000	0.5652	0.4348	0.1304	0.2609	1.3000
0.7500	0.5556	0.4444	0.1111	0.2222	1.2500
0.8000	0.5455	0.4545	0.0909	0.1818	1.2000
0.8500	0.5349	0.4651	0.0698	0.1395	1.1500
0.9000	0.5238	0.4762	0.0476	0.0952	1.1000
0.9500	0.5122	0.4878	0.0244	0.0488	1.0500
1.0000	0.5000	0.5000	0.0000	0.0000	1.0000

The relationships illustrated in table 1 provide empirically testable predictions. As P^N goes from zero to one, the price ratio, R_1/R_2, falls. Similarly, as P^N goes from zero to one the probability of an eventual sale falls. This implies that in markets where prices fall rapidly as a function of time on the shelf, the probability that the good will go unsold is relatively low. The prediction in the housing sample is that mansions, for which P^N is high, should have slowly declining prices and should be more likely to be taken off the market after an unsuccessful attempt to sell than inexpensive condos, for which P^N is low.

This implication is not an obvious one. Since R_1/R_2 is high when R_1 is high, the logic implies that for a given prior, goods for which price starts high are actually more likely to sell. The reason is that the high initial price reflects low P^N. It is not a matter of calling out prices randomly. The high initial price is a response to conditions that also imply that a sale is likely.

Note that the expected price at which the good sells is always 1/2, irrespective of P^N. This follows because (prob. sale in 1)/(prob. sale in 2) = 1 and because $(R_1 + R_2)/2 = .5$. This also illustrates the important distinction between the expected price at which a good sells and expected revenue. Although expected price, given a sale is independent of P^N, expected revenue falls with P^N. The probability that the good remains unsold (and is returned to the supplier as scrap) increases with P^N. The point that not all goods are sold and that there is a systematic relationship between pricing and the probability of a sale is fundamental to this analysis. It plays an essential role in reconciling some phenomena described below.

F. Recapitulation

The ability to readjust price as a function of past sales provides the firm with a richer strategy set. This is especially important when the firm is more uncertain about the value that consumers attach to the good in question. Not only does intertemporal pricing permit more than one chance to attract buyers, but it also allows the firm to learn about the nature of demand in the market.

An important implication is that prices start high and fall with time on the shelf. The level of initial price and speed with which price falls are positively related to the number of customers that it encounters per period and to the proportion of real "buyers" in the group. Thin markets have lower initial and more rigid prices.

Competition among firms for customers reduces the number of potential buyers that any one seller encounters. This drives profits to zero, but in the process, alters the optimal pricing rule. Department stores that sell somewhat distinct items will select lower initial prices and lower those prices more slowly when there is competition between stores for buyers. This is true even when the good in question is available at only one store.

The time pattern of transactions tends to be smooth over time. The probability that the good eventually sells is positively related to initial price and the rate of price decline along the optimal price trajectory.

II. Heterogeneous Goods, Fashion, Obsolescence and Discount Rates

This section builds on the earlier ones to explain how prices vary with factors like product heterogeneity, obsolescence rates and time discounting. In most of this section, it will be assumed that all customers are buyers, that is, that $P = 0$ so that the less complex formulation of the model can be used.

A. Heterogeneity Among Goods

Is there any sense to the claim that women's clothes cost more than men's, even for given cost conditions? This is a direct implication of different product heterogeneity across the type of good.

Formally, what this section examines is how dispersion in the prior on V affects pricing policy and the probability that a sale is made. Assume that $P = 0$ so that the firm's problem becomes the one in (2) (which is the special case of (4) with $P = 0$).

Consider a mean preserving spread. For expositional convenience, let us be specific. Suppose that the prior on V for, say, men's clothes, is uniform between .5 and 1.5, but for women, it is uniform between 0 and 2. (We ignore the endogeneity of the prior throughout.) The idea is that to the extent that women's clothes take on more variations, it is more difficult to predict the value of any particular item. Both distributions have the same mean value and it would seem that average prices, average revenues, and expected revenues might be the same. This is not the case.

Given the distributions, the prior distribution function for men's clothes is

(16a) $$F(V) = V - 1/2 \quad \text{for} \quad 1/2 \leq V \leq 3/2$$

and this results in a posterior for any given R_1 of

(16b) $$F_2(V) = (V - 1/2)/(R_1 - 1/2) \quad \text{for} \quad 1/2 \leq V \leq R_1 \ .$$

Similarly, the prior for women's clothes is

(17a) $$F(V) = V/2 \quad \text{for} \quad 0 \leq V \leq 2$$

and this results in a posterior for any given R_1 of

(17b) $$F_2(V) = V/R_1 \quad \text{for} \quad 0 \leq V \leq R_1 \ .$$

Substitution of (17a,b) into (2) yields the solution that the initial price for women's clothing, R_1, equals 4/3 and the period 2 price, R_2, is 2/3. This makes sense since the prior on V is simply a rescaling of the original prior, where solutions were 2/3, 1/3.

Substitution of (16a,b) into (2) yields the solution that the initial price of men's clothing, R_1, equals 1 and the period 2 price, R_2, is 1/2. Note that the period 2 price is the lower bound of the posterior (and prior) distributions so that the optimum in this case is to make the sale a certainty in period 2. (Of course, this result is dependent on the shape of the density function.)

Given the prices and the priors, it is obvious that the probability that the woman's garment sells at 4/3 is 1/3 and the (unconditional) probability that it sells at 2/3 is also 1/3. This results in an expected price of 1, and the good is sold 2/3 of the time so expected revenue is 2/3.

For men's clothes, the probability that the garment sells at $R_1 = 1$ is 1/2 and the probability that it sells at 1/2 is also 1/2. The expected price is 3/4 and expected revenue is 3/4.

Although this is only one example, it illustrates a number of important points: First, the more disperse prior results in a higher expected price for a given mean. Thus, women's clothes cost more than men's clothes. Second, because women's clothes remain unsold more often than men's clothes, expected revenues can be lower, even though the price, given a sale, is higher. In competition, firms enter the men's clothing industry until expected revenue is equal across sectors. In the more general model, where $P \neq 0$, this reinforces the result that men's clothes sell for lower prices. The key to this result is that at the optimum prices, more women's clothes remain unsold. The seller either retains the good (as in the case of an unsold house), or wholesales it off.

A similar story might apply to goods that are very new or rapidly changing over time. To the extent that the prior is more diffuse for these goods, their prices should start higher, but fall faster than those on more traditional items. This predicts more variance over time in the prices of new computers than in the prices of standard typewriters. Another reason for high price variance in the computer market may be the importance of obsolescence. The next section examines that issue.

As an empirical matter, economists who construct price indexes tend to focus on the price, given a sale, and ignore the probability of a transaction.

What this points out is that pricing and sale probabilities are linked. For many purposes, when the probability of a sale is less than one, expected revenue-per-good might be a better metric than expected price, given a sale. The former is more closely related to what the firm generally cares about, even though the latter is what consumers care about.

B. Fashion, Obsolescence and Discounting the Future

Some goods go out of style very quickly whereas others seem to retain their popularity for long periods of time. Again, the example of men's and women's clothes may be relevant. It may be true that men's suits change lapel widths less frequently than women's clothes change style. That phenomenon is assumed exogenous for the purposes of this paper, but it is interesting to know how fashion, or obsolescence as it might be termed in other markets, affects the choice of initial price and the rigidity of prices over time.

This is easily treated in the current framework. Let us think of obsolescence or fashion as taking the following form: During the first priod, the good is worth V, but in the second period it is worth V/K, where $K \geq 1$. The seller still does not know V, but he does know that whatever it is, it will retain only $1/K$ of its worth in period 2.

All that changes is the value that is inserted into the period 2 density function. That is, the individual buys the good in period 2 when

$$V/K > R_2$$

or when

$$V > KR_2 .$$

During period 1, nothing is changed so the firm's maximization problem in (2) now becomes

$$\text{(18)} \qquad \max_{R_1, R_2} \; R_1[1 - F(R_1)] + R_2[1 - F_2(KR_2)]F(R_1)$$

Assuming that the prior is uniform between zero and one, the optimum prices are from the first-order conditions

(19a) $$R_2 = \frac{R_1}{2K} .$$

and

(19b) $$R_1 = \frac{2K}{4K - 1} .$$

For $K = 1$, the solutions are identical (as they must be) to those obtained without obsolescence, namely, $R_1 = 2/3$, $R_2 = 1/3$.

What is clear from (19a) is that prices fall faster with time on the shelf when K is large. The reason, of course, is that the seller knows that if the good was worth V in period 1 it is only worth V/K in period 2, so period 2's price adjusts accordingly.

Equally intuitive is that the price in period 1 is lower when K is large. The more obsolete the good becomes, the more anxious is the seller to get rid of it in period 1. As a result, he trades off this sense of urgency against the price that would provide him with the best posterior to carry into period 2.

Stated alternatively, a "classic," defined as a good that does not go out of style, carries a higher initial price, independent of any resale considerations. Its price is less sensitive to inventory than a good that goes out of style rapidly. This is true even for a given set of cost conditions.

Time discounting, although seemingly similar, is somewhat different. The reason is that even though the seller might think of a sale in period 2 at V as worth only V/K in present value, the posterior density function is still on V, not V/K because buyers are willing to pay V in period 2. Thus, the objective function is not (18), but is instead

(20) $$\max_{R_1, R_2} R_1(1 - F(R_1)) + \frac{R_2}{(1 + r)} (1 - F_2(R_1))F(R_1) .$$

the present value of the period 2 price is $R_2/(1 + r)$, but the customer continues to buy the good as long as $V > R_2$. The solution when V is uniform between zero and one is given by

(21a) $$R_2 = R_1/2$$

(21b) $$R_1 = \frac{2(1 + r)}{4(1 + r) - 1} .$$

Note that R_2 is $R_1/2$, which differs from (19a). The price does not fall more rapidly when the discount rate is positive. This is as it should be. Given that the firm gets to period 2, the best that it can do is take the information from period 1 (that $V < R_1$) and optimize. Discounting is irrelevant to that decision. This was not true when the good became obsolete in the second period.

But implications about urgency are similar. As r gets large, the firm is anxious to make the sale in period 1, not because the good will become obsolete, but for reasons of time preference. At the extreme, as r goes to infinity, $R_1 = 1/2$. The value of the second period is zero so the firm behaves as it would in the one-period problem, setting $R_1 = 1/2$. But if it does get to period 2, the best policy now is to cut price to 1/4 because it knows (with certainty) that $V < 1/2$.

Time discounting reduces the initial price, but does not change the rate at which prices fall as a function of time on the shelf. Obsolescence reduces the initial price too, but also increases the rate at which prices fall as a function of time on the shelf.

C. Longer Horizons

Two periods have been assumed throughout the analysis. Time discounting was one way to modify that assumption, but it is useful to consider more directly how a change in time horizon affects pricing strategy. In many respects, this is another way to treat obsolescence, but there is more to it than that.

Consider a firm that has T rather than 2 periods during which to sell its product. The problem in (2) generalizes to

$$\max_{R_1, R_2, \ldots, R_T} R_1[1 - F(R_1)] + R_2[1 - F_2(R_2)]F(R_1) + R_3[1 - F_3(R_3)]F(R_2) + \ldots + R_T[1 - F_T(R_T)]F(R_{T-1}) \tag{22}$$

where $F_t(V)$ is the posterior after $t - 1$ periods. As before, $F(V)$ refers to the prior distribution before period 1. Each term on the right-hand side has as one of its components $F(R_{t-1})$ because this is the probability that the good was not sold before period t. The problem yields a system of recursive first-order conditions given by

$$\begin{aligned} \frac{\partial}{\partial R_1} &= 1 - 2R_1 + R_2 = 0 \\ \frac{\partial}{\partial R_2} &= R_1 - 2R_2 + R_3 = 0 \\ &\vdots \\ \frac{\partial}{\partial R_{T-1}} &= R_{T-2} - 2R_{T-1} + R_T = 0 \\ \frac{\partial}{\partial R_T} &= R_{T-1} - 2R_T = 0 \end{aligned} \tag{23}$$

These yield the general solution that

$$R_T = 1/(T + 1) \tag{24a}$$

$$R_t = (T - t + 1)R_T = \frac{T - t + 1}{T + 1}. \tag{24b}$$

These solutions are quite intuitive. First, as T gets large so that the horizon lengthens, eq. (24a) implies that the price in the last period goes to zero. Second, eq. (24b) implies that price drops by a smaller amount each period with increases in T. Thus, price changes less rapidly per period. But the initial price is higher as T increases so that a larger total range of prices is covered. As T goes to infinity, $R_1 = T/(T + 1)$ goes to 1. The price starts at the top and moves down trivially each period until V is hit precisely. As T goes to 1, we are back to the one-period problem and R_1 is 1/2.

Stated simply, as the firm's selling horizon lengthens, initial price is higher and prices fall off less rapidly each period. However, the price in the final period is lower as the time horizon increases. This also implies that the probability that the good sells before the end is reached increases in T because $1 - F(R_T)$ increases in T.

The difference between adding periods and merely lengthening the time associated with each period is that learning takes place and a new price can be chosen each period. This comes back to an essential feature of "retailing" as defined in this paper. The price is fixed for a given length of time (which is likely to depend on the number of customers encountered per unit of time). Price changes only occur at the end of that interval. No attempt is made to call out the highest possible price, and lower it until the customer agrees to purchase. There are good reasons for not doing this, and those reasons are discussed below.

D. Non-unique Goods

There is another respect in which the time horizon can be lengthened. The situation that many firms face in marketing new products is somewhat different from the one analyzed so far. Above, it was assumed that once the good is sold, there are no others to sell. This is appropriate for a painting or designer dress, but what of a new computer model put out by an established company? The fact that the good sells in the first period does not preclude additional sales in the second period. How should the prices be set under these circumstances?

Again, for simplicity, return to the two-period horizon problem and continue to assume that $P = 0$: all customers are buyers. Now, three prices are relevant: The seller must select a price in period 1, R_1; he must choose a price in period 2 given that no purchases were made in period 1, R_2; and he must choose a price in period 2 given that at least one purchase was made in period 1, $\tilde{R}_2$. (Under the assumptions about consumer homogeneity, knowing the exact number of items sold provides no additional information.[14]

Normalize such that one item is available for sale in period 1 and N_2 are available in period 2. If no sale occurs in period 1, then $N_2 + 1$ are available. N_2 may be greater or less than 1. The preceding analysis of unique goods is merely a special case, with $N_2 = 0$. The firm's maximization can be written as

(25)

$$\max_{R_1, R_2, \tilde{R}_2} R_1(1 - F(R_1)) + (N_2 + 1)R_2(1 - F_2(R_2))F(R_1) + N_2\tilde{R}_2(1 - \tilde{F}_2(R_2))(1 - F(R_1))$$

It is especially revealing to treat this as a dynamic program and to examine what happens if the good sells in period 1.

The second-period problem is

(26)
$$\max_{\tilde{R}_2} N_2\tilde{R}_2(1 - \tilde{F}_2(R_2)) .$$

Again, using Bayes' Theorem

$$\tilde{f}_2(V) = 0 \qquad \text{for } V \leq R_1$$
$$= \frac{f(V)}{1 - F(R_1)} \qquad \text{for } V > R_1$$

so

$$\tilde{F}_2(V) = 0 \qquad \text{for } V \leq R_1$$
$$= \frac{F(V) - F(R_1)}{1 - F(R_1)} \qquad \text{for } V > R_1 .$$

The maximization in (26) becomes

(27a) $$\max_{\tilde{R}_2} N_2\tilde{R}_2 \qquad \text{if} \quad \tilde{R}_2 < R_1$$

(27b) $$\max_{\tilde{R}_2} \frac{N_2\tilde{R}_2(1 - F(\tilde{R}_2))}{1 - F(R_1)} \qquad \text{if} \quad \tilde{R}_2 \geq R_1 .$$

Branch (27a) is always increasing in $\tilde{R}_2$ so if the solution is on this branch, it is at $\tilde{R}_2 = R_1$. If $\tilde{R}_2 \geq R_1$, then the first-order condition for (27b) is relevant:

$$\frac{d}{d\tilde{R}_2} = \frac{N_2}{1 - F(R_1)} (1 - F(\tilde{R}_2) - \tilde{R}_2 f(\tilde{R}_2)) = 0$$

or

$$\tilde{R}_2 = \frac{1 - F(\tilde{R}_2)}{f(\tilde{R}_2)} .$$

This solution is identical to that of the one-period problem. But the optimal price in the one-period problem can never exceed R_1, so the corner is relevant here, too. Thus, the solution is $\tilde{R}_2 = R_1$.

This implies that price in the second period never rises, even if the good sells during the first period. The reason is that the part of the distribution below R_1 is irrelevant anyway, so knowing that $V > R_1$ does not change the decision on the optimal price.

Given that $\tilde{R}_2 = R_1$, and using the definition of $\tilde{F}_2(V)$, eq. (26) can be rewritten as

$$\max_{R_1, R_2} R_1(1 - F(R_1)) + (N_2 + 1)R_2(1 - F_2(R_2))F(R_1) + N_2R_1(1 - F(R_1))$$

or

$$\max_{R_1, R_2} (N_2 + 1)R_1(1 - F(R_1)) + (N_2 + 1)R_2(1 - F_2(R_2))F(R_1) .$$

Since the scalar $(N_2 + 1)$ is irrelevant, this problem is identical to the one in (2), where goods were assumed to be unique. Thus, all results already derived hold even in the case of non-unique goods. In this example,

$$R_1 = 2/3$$
$$R_2 = 1/3$$
$$\tilde{R}_2 = 2/3 .$$

If some sales are made in period 1, then price is held at R_1. If no sales are made, then price is halved.[15]

E. Spoiling the Market and Nondurable Goods

The result that the price following a successful period 1 never falls hinges on the assumption that demanders who find the price too high in period 1, return for another look in period 2. (If no sale occurs, there are $N_2 + 1$ buyers in period 2.) This is an inappropriate assumption in two obvious cases. The first is that buyers lose interest when they find that $R_1 > V$. This may be rational when buyers do not know the firm's prior so that they cannot forecast its price cutting behavior. The second is that the good is nondurable. For example, a hotel room that was vacant on Saturday night cannot be stored and sold again on Sunday.

Under these circumstances, the maximization problem is

$$\max_{R_1, R_2, \tilde{R}_2} \; R_1(1 - F(R_1)) + R_2 N_2(1 - F_2(R_2))F(R_1) + \tilde{R}_2 N_2(1 - \tilde{F}_2(R_2))(1 - F(R_1))$$

Using the results of the previous section, this can be written as

$$\max_{R_1, R_2} \; (N_2 + 1)R_1(1 - F(R_1)) + R_2 N_2(1 - F_2(R_2))F(R_1) .$$

The first-order conditions imply that at the optimum,

$$\tilde{R}_2 = R_1$$
$$R_2 = R_1/2$$
$$R_1 = \frac{2N_2 + 2}{3N_2 + 4} .$$

If $N_2 = 1$ so that demand is constant over time, $R_1 = 4/7$, instead of 2/3 as obtained before. The reason is that losing a sale in period 1 is now more costly since the market is spoiled for that buyer, so the firm selects a lower first-period price. The learning effect, which is still present, is offset to some extent by the desire to avoid having first-period buyers walk away without buying. It is not offset completely because $R_1 > 1/2$, so this is not the same as consecutive one-period problems. Even though the good is not storable, the information derived from period 1 is, so sellers of nondurables do not behave myopically. As N_2 gets large, R_1 approaches 2/3 because the lost sales in period 1 are trivial, relative to revenue generated in period 2. The information effect dominates.

IV. Consumer Behavior

There are a number of aspects of consumer behavior that are worth considering. We start by analyzing strategic play by purchasers and link this analysis to the auction literature.

A. Strategic Considerations

Consumers may know the firm's pricing policy and in particular, that $R_2 < R_1$. Does it pay for a consumer in period 1 to wait for period 2, knowing that by doing so he may be able to purchase the good at a lower price? The decision depends on the number of rival customers.

Suppose that a buyer has located a gown in period 1 that she values at $V > R_1$. If she buys the dress, she earns rent $= V - R_1$. Alternatively, she can wait until period 2 hoping that no others will get there first. The more potential customers there are in the market, the lower is the probability that the gown will remain on the rack into the next period. If the buyer passes up the gown this time, there are $N - 1$ other customers who might beat her to it next period. Therefore, the expected rent from waiting is

$$(V - R_2)/N \ .$$

She waits if[16]

$$V - R_1 < (V - R_2)/N$$

or if

$$R_1(1 - 1/2N) > V(1 - 1/N)$$

since $R_1 = 2R_2$.

As N goes to infinity, the left side goes to R_1 and the right to V. So the consumer waits if $R_1 > V$. But $R_1 > V$ precludes buying the good in period 1 anyway, so for sufficiently large N, strategic behavior is not an issue. There are too many others around who can beat this customer to it. She buys it when she finds it.[17] The argument is reinforced if some of this period's buyers might obtain the good first.

On the other hand, as N goes to 1 it is certain that the consumer behaves strategically because $R_1/2 > 0$. The consumer is sure to get it next period since she has no competition so he might as well wait for a lower price. Thus, strategic behavior is not an issue when there is a large number of potential buyers, but may be important when only a few individuals are even potentially interested in the good.[18]

If goods are not unique, then without time preference, strategic behavior results in an equilibrium that is identical to that of the one-period problem. The reason is that all buyers gain if no sales are made in period 1. Since the good is not unique, all are satisfied in period 2 at a lower price ($R_2 < R_1$). No buyer has any incentive to purchase in period 1. Of course, if the seller knows this, then he can infer nothing from the fact that no one bought in period 1. As such, his problem is like the one-period problem so the solution is $R_t = 1/2$. Given that solution, no strategic waiting occurs.

The original solution (2/3, 1/3, 2/3) is restored if some buyers have high time preferences for the good. (A calculator that produces services over time provides more utility, the sooner it is acquired.) Then at least some buyers have an incentive to deviate from the waiting strategy. The initial "no waiting" equilibrium is restored when some buyers have sufficiently high time preference to buy in period 1. This requires $V - R_1 > (V - R_2)/(1 + r)$ where r is the discount rate. For r sufficiently large, a sale in period 1 is guaranteed when $V > R_1$. Then the solution reverts to setting $\tilde{R}_2 = R_1$ so that strategic waiting is not an issue. Since price does not fall over time, nothing is gained by waiting.

B. Auctions and Stochastic Arrival of Customers

One way to deal with uncertainty about consumer demand is to hold an auction.[19] In fact, the solution to the basic problem is in many respects simply a Dutch auction, where the price begins high and continues to fall until a purchaser declares that he is willing to buy at that price. In the case where the time horizon is long, so that the reduction in price is small at each period, and where N is small so that consumers may behave strategically with respect to waiting time, the analysis is that of a traditional Dutch auction.

There are two major differences between this analysis and the one that pertains to the standard Dutch auction. The first is one of emphasis. This analysis for the most part assumes that N is large and consequently ignores most strategic behavior by consumers.[20] It focuses instead on the rule that the seller uses to choose the optimal size of the step as a function of the number of bidders and their types (shoppers or buyers), and of the number of periods in the horizon. Recall that the number of periods depends on the cost of changing price because a period is defined as that time during which price does not change. If there were no cost to changing price, then the period notion would be dispensed with and the seller would change price each time a customer examined the good. Then R would move with N only and time would be irrelevant.

The second important difference is related. In a Dutch auction, there is no reason to alter the size of the step once the process is in motion. That is, once the seller selects an amount by which to lower price each time, no new information appears until a buyer agrees to purchase the good, at which point it is too late to use that information. That is not true in the retailer context, nor is it true in this model. Although we have assumed throughout that N is fixed, there is nothing in the setup of the problem that precludes a stochastic N. In fact, for a given choice of R_1, the optimal R_2 is given by equation (9), reproduced here:

$$R_2 = \frac{1}{2}\left[R_1(1 - P^N) + P^N\right] \quad (9)$$

If N here is interpreted as the realization of N in period 1, then (9) still holds as the optimum R_2 (because the second period N does not enter). Thus, the retailer can alter his choice of R_2, i.e., change the size of the step in the Dutch auction, after having observed something from the first period. Of course, the ability to do so changes the choice of R_1 because that problem now involves an integral over all possible realizations of N.[21] But the point is that the retail pricing policy has an additional instrument that is useful in all cases where N is stochastic. That instrument is the ability to select the size of the step after obtaining some information.

There are at least two related reasons why a seller might choose a strict retail pricing rule over some form of auction or haggling. The first, already mentioned, is that an auction may encourage strategic behavior on the part of consumers that is absent when retail pricing is used. For example, consider confronting every potential buyer with the maximum possible V and lowering price very rapidly until the consumer agrees to purchase. Obviously, the consumer would wait until $P = 0$, knowing that the considerations that prevented strategic waiting in the last section are not relevant with this type of pricing behavior. The retail pricing method, where R_1 is chosen and fixed in advance, discourages strategic behavior by consumers when the number of customers is sufficiently large. This suggests that retailing is more likely to be used when there are a large number of anonymous buyers.

The second, and perhaps more compelling reason why large, impersonal stores might prefer retailing to some form of haggling has to do with delegation of authority. Even if no agency problems exist, it is not unreasonable to believe that the management of a department store would not trust price setting to low-paid retail clerks. Since buyers can always refuse to buy if the price is too high, but can purchase if the price is too low, bad price setting can result in losses to the firm even if those prices are only randomly too high or too low. To avoid this adverse selection problem, the firm may decide to have its experts announce a rigid price (or price rule) that is posted. No haggling is permitted because the clerk who represents the store may not be good at it. Agency problems reinforce this result.

This suggests that haggling is more likely to occur when the owner (who is presumably the high quality price setter) is also the sales agent. Mom-and-pop stores are more likely to bargain with their customers over price than are large department stores, which use retail pricing almost exclusively.

There are some more subtle elements, special to this setup, that are not generally part of the Dutch auction. First, with a Dutch auction, all bidders are present at the same time. Equivalently,[22] each may submit a binding sealed bid. In the case of the former, it is necessary that customers examine the good at the same time. The story in this paper allows individuals to arrive at different times during each period. The first one to arrive who will pay price R_1 in period 1 (or R_2 if it goes to period 2) gets the good. The Dutch auction imposes the cost that a common meeting time must be found. Retail pricing does not impose that cost because consumers can choose their own shopping time.

Sealed bids do not force a common meeting time, but they do create a waiting period between the time that the bid is made and the winner is determined. This, too, is costly. For example, a woman bids on a dress for the Ball and then sees another one before she learns whether she made the winning bid on the first. Buying the second dress might leave her with two, but failure to purchase might result in her having none.

V. Empirical Thoughts

The purpose of this theory is to provide some empirical implications on pricing and transactions behavior as a function of some observable parameters. There are a substantial number of predictions about pricing and time on the shelf as a function of the number of customers per period, the type of customers (shoppers or buyers), the time horizon, interest rates, the durability of the good, and the shape of the prior.

With the exception perhaps of the last, all of these have observable analogues. The number of customers and general thinness of the market can be proxied by the turnover rate for the good. E.g., houses that turn over more rapidly are sold in markets with higher N. The type of customer, P, can be measured by the proportion of individuals who examine a good relative to the number that actually buy it. The time horizon relates to the maximum number of times that a price is changed before the good is taken off the floor. E.g., bargain basements eventually give up on some goods. After how many price reductions does this occur?

Additionally, some state contingent behavior has been predicted. For example, when goods are non-unique, the price in period two that follows a successful period 1 differs in a specific way from the price that follows an unsuccessful period 1. This relationship is observable and can be tested. Similarly, the time distribution of transactions is related to the initial price and to the speed with which price falls over time. Again, both are observable. Finally, more uncertainty in the prior implies a higher initial price with more rapid decline. It also implies that at the optimal prices, more goods are left unsold when the prior is diffuse. (Fewer men's suits are left unsold than women's dresses.) A relationship between initial price, rate of price decline, and proportion of unsold goods is predicted, and all are observable.

VI. Summary and Conclusion

Sellers must gauge the market any time they attempt to sell a new item. Their attempt to do so and to learn from experience leads to pricing and selling behavior that varies in predictable ways with some observable characteristics of the market.

The major theme is that prices start high and fall as a function of time on the shelf. The speed of that fall and the initial price itself increase as the number of customers per unit of time increases, as the proportion of customers who are "genuine buyers" as opposed to "window shoppers" increases, and as prior uncertainty about the value of the good increases. The optimum price path implies that in the case considered, the probability of making a sale is constant over time.

A number of additional predictions are obtained. First, diffuse priors may result in higher initial prices and more rapid fall, but also in higher average price and more goods left unsold. This might explain why women's clothes carry a higher average price than men's, but are of lower average "quality," even in a competitive market.

Second, goods that become obsolete more rapidly or are more susceptible to fashion exhibit lower initial prices as well as prices that fall more rapidly with time on the shelf. Positive discount rates have a similar effect on initial price, but not the same effect on the rate of price fall.

Third, non-uniqueness of the good does not alter the solution. A successful first period is followed by no change in the price, whereas an unsuccessful first period results in the same price reduction as is warranted when goods are unique.

Fourth, for nonstorable goods, or when spoiling the market is an issue, the price reduction policy is the same but the initial price is lower than for storable goods. It is higher than the price that a myopic seller of non-durables would charge because even though the good is not storable, the information derived from period 1 is of value.

Fifth, rigid price reduction policies used by bargain basements are predicted under certain circumstances.

Finally, the paper examines strategic behavior by consumers and the relationship between a retailer's pricing policy and Dutch auctions.

References

Bass, Frank. "A New Product Growth Model for Consumer Durables." Management Science 15 (January 1969):

__________ and Bultzen, Alain. "A Note on Optimal Strategic Pricing of Technological Innovations." Marketing Science 1 (Fall 1982): 371-438.

Butters, Gerald. "Equilibrium Price Distributions and the Economics of Information," Unpublished Ph.D. Dissertation, University of Chicago, 1975.

Clarke, Darral G., and Dolan, Robert J. "A Simulation Analysis of Alternative Pricing Strategies for Dynamic Environments." Journal of Business 57 (January 1984): S179-S199.

DeGroot, Morris H. Optimal Statistical Decisions, New York: McGraw-Hill.

Engelbrecht-Wiggans, R. "Auctions and Bidding Models: A Survey." Management Science 26 (February 1980): 119-42.

Grossman, Sanford, Kilstrom, Richard, and Mirman, Leonard. "A Bayesian Approach to the Production of Information and Learning by Doing." Review of Economic Studies 44 (3): 533-48.

Harris, Milton, and Raviv, Artur. "A Theory of Monopoly Pricing Schemes with Demand Uncertainty." American Economic Review (June 1981):

__________ and __________. "Allocative Mechanisms and the Design of Auctions." Econometrica 49 (1981b): 1477-99.

Lazear, Edward P. "Raids and Offer-Matching." NBER Working Paper No. XXX, August 1984.

Milgrom, P., and Roberts, J. "Price and Advertising Signals of Product Quality." Mimeographed, Yale and Stanford Universities, May 1984.

Milgrom, Paul, and Weber, R. J. "A Theory of Auctions and Competitive Bidding." Econometrica 50 (September 1982): 1089-1122.

Myerson, Roger. "Optimal Auction Design." Mathematics of Operations Research (1981).

Riley, John, and Samuelson, William. "Optimal Auctions. American Economic Review 71 (June 1981): 381-92.

Spence, A. Michael. "The Learning Curve and Competition." Bell Journal of Economics 12 (Spring 1981): 49-70.

Stokey, Nancy. "Rational Expectations and Durable Goods Pricing." Bell Journal of Economics 12 (Spring 1981): 112-28.

Vickrey, William. "Counterspeculation, Auctions and Competitive Sealed Tenders." Journal of Finance 16 (March 1961): 8-37.

Wilson, Robert. "Competitive Bidding with Asymmetric Information." Management Science 13 (July 1967): 816-20.

__________. "A Bidding Model of Perfect Competition." Review of Economic Studies 138 (1977): 511-18.

Footnotes

[1]Neither the questions or methodology are entirely novel. The updating procedure used is described in DeGroot (1970) and has been used more recently by Grossman, Kilstrom, and Mirman (1975). Another line of literature, Bass (1969), Bass and Bultzen (1982), Clark and Dolan (1984), and Spence (1982) examines some of these problems, but the focus is on cost conditions that change over time, generally tied to some learning by doing. It is my view that the essence of the marketing problem that faces a firm that introduces a new product is selecting a strategy in the face of uncertainty about the demand for its product. The evolution of prices and transactions over time is more likely to reflect learning about the market than learning about producing the product. Both models give declining prices over time, but in the case of cost changes, myopic sellers charge prices that are too high (Bass and Bultzen) whereas in the case of learning about demand, myopic sellers choose prices that are too low.

[2]It is useful to recognize that $F(V)$ is determined after the retailer has seen the good himself. For example, retail sellers of dresses know that they vary in price from \$50 to \$10,000. After having examined the good, the seller may know that a particular dress has a V between \$500 and \$1,000, but he does not know the exact value within that range.

[3]This is different from the usual price discrimination problem where the demand curve of the market is known, but no buyer will reveal where he is on that demand curve. That problem is treated by Stokey (1981).

[4]The same solution is obtained if (2) is maximized simultaneously choosing R_1 and R_2 because of the time-consistent nature of the problem.

[5]This pricing pattern resembles a Dutch auction. This is discussed in greater detail below.

[6]This is an inaccurate assumption in some cases. For example, the demand for telephones depends on the number of others who can be called.

[7]This result is obtained by Grossman, Kilstrom and Mirman (1975). Although most of their work focuses on consumer learning, they show that an experimenting monopolist always starts with a higher price than a non-experimenting one.

[8]Harris and Raviv (1981a) consider the use of different types of pricing mechanisms when demand conditions are uncertain. Their "priority pricing" scheme resembles an intertemporal price decline. They show that such a scheme is optimal when capacity falls short of potential demand. That is the situation that is implicit in this setup because the unique good can be sold to only one of many potential buyers.

[9]This abstracts from shopping for the pure pleasure of it and from any information that might be useful in making future purchases.

[10]Note that an identical mechanism is at work in the labor market context when trying to infer a worker's product/wage ratio from past transactions. This is the subject of Lazear (1984).

[11]One might ask whether advertising the price could bring about an ex post competitive equilibrium. The answer is no. Consider a store that said that it offered paintings for $300. The consumer would not know whether that is a low price for a prior distribution of V that lies, say, between $300 and $400 or a high price for a prior that lies between $200 and $300. This is why the point of footnote 2, that the seller draws $F(V)$ from some larger distribution, and sees the narrower distribution before pricing, is important.

Recently, Milgrom and Roberts (1984) have shown that a seller should use price and the level of advertising to signal the quality of the good to the consumer. That will not solve the problem here for two reasons: First, repeat purchases play a crucial role in their story. Second, the distribution that the seller sees before pricing differs from the one that the buyer sees before entering the store. The buyer must engage in some search before he knows even the prior on V on which the seller bases his pricing. Note, also, that by definition, V is the price at which the buyer purchases the good and this takes into account the option of walking out and examining the paintings in the next gallery.

[12]The derivative is messy and without obvious intuitive appeal. However, it can be shown that expected revenue takes on the following values:

pN	Revenue
.00	.333
.05	.322
.20	.290
.50	.200
.75	.111
.95	.024
1.00	.000

[13]P may or may not vary with N. If it does, then stability conditions must be checked and equilibrium, if it exists, need not be unique.

[14]Second-hand markets are ignored.

[15]The solution that price never rises after a successful period depends critically on two assumptions. First, there are no contagion or network effects that shift demand in period 2 relative to period 1. Second, the group of buyers is homogeneous in the assessment of V.

[16]This ignores one-period bargaining considerations.

[17]The argument here is a special case of the more general one made by Wilson (1977). Wilson shows that as the number of bidders gets large, a sealed bid auction results in bids that are almost certain to be equal to the

true reservation value. As will be pointed out below, the declining price retail policy is like a Dutch auction, which is equivalent to a sealed bid auction in fundamental respects. As such, the Wilson result is relevant in this context.

[18]This assumes that the seller ignores the strategic behavior of consumers. This assumption is troublesome, and without it, a pure strategy equilibrium may not exist.

[19]There is a large literature on auctions starting with Vickrey (1961). He explores the allocative and profit implications of a number of different kinds of auctions, including the second price and Dutch auction. A number of recent papers have characterized the conditions under which various types of auctions are efficient and profit-maximizing. Among those are Butters (1975), Engelbrecht-Wiggans (1980), Harris and Raviv (1981b), Myerson (1981), and Milgrom and Weber (1982).

[20]Most of the auction literature focuses on strategic behavior by consumers in selecting a bid. An early example of this kind of analysis is Wilson (1967), who examines what happens in an auction when one and only one party is informed about the value of the good.

[21]The problem in (7) then becomes

$$\max_{R_1, R_2} E\{R_1[1 - P^{\tilde{N}_1} - (1 - P^{\tilde{N}_1})R_1] + R_2[P^{\tilde{N}_1} + (1 - P^{\tilde{N}_1})R_1 - R_2](1 - P^{\tilde{N}_2})\}$$

where $\tilde{N}_t$ is the stochastic number of buyers in period t.

[22]See Riley and Samuelson (1981).