Authors: *Becker, Bearden, & Rapoport*
Article submitted to *Management Science; manuscript no. MS-*

Perishable Asset Dynamic Pricing in the Laboratory

William J. Becker

Department of Management and Organizations, University of Arizona

J. Neil Bearden

Decision Sciences Area, INSEAD, Singapore, and
Department of Management and Organizations, University of Arizona

Amnon Rapoport

Department of Management and Organizations, University of Arizona, and
Department of Marketing, Hong Kong University of Science and Technology

Abstract

Individual decision behavior in a computerized, discrete-time, perishable asset dynamic pricing problem is examined experimentally under three different demand conditions. The conditions differ from one another in the relationship between the stochastic distribution of demand and termination time of the selling season. In a between-subjects design, subjects are faced with either constant, increasing, or decreasing expected demand. In all three conditions, we find that they improve their pricing policies (and thereby their revenue) with experience. Yet, even after substantial experience with the problem, several systematic pricing biases persist. The general pattern of results suggests that, in comparison to the optimal pricing policy, subjects tend to overreact to relative demand: they price the asset too high when demand is relatively high and too low when demand is relatively low. The departures from optimal pricing result in average revenue losses of about 8%. These results demonstrate that optimal dynamic pricing models may help to significantly increase the performance of unaided decision makers.

Key words: dynamic pricing; behavioral experiment; decision biases

JEL Classification: D4, D8, L1

History: This paper was first submitted July 25, 2007

1. Introduction

There has been considerable theoretical work in operations research (OR) on *optimal* decision behavior in revenue management settings, but little is known about the actual behavior of flesh-and-blood decision makers in the problems that have been modeled. Consequently, it is not known whether or by how much sophisticated OR techniques could help increasing the revenue of firms that currently do not employ them. If, for instance, unaided decision makers (e.g., retailers, hotel and motel managers) perform as well as the optimal algorithms, then firms may see no compelling reason to pay for expensive software that implements them. On the other hand, if one could demonstrate that the algorithms would produce significant increase in revenues, then firms might be persuaded to pay more attention to theoretical work in OR. Primarily intended to examine decision behavior in a perishable asset dynamic pricing problem, the current paper reports evidence that OR methods may, indeed, help a great deal.

Dynamic pricing is one of a few basic categories of demand-management decisions addressed by the discipline of revenue management (see, e.g., Bitran & Caldentrey, 2003; McGill & Van Ryzin, 1999; and Talluri & van Ryzin, 2004, for recent and excellent reviews). One of the fastest growing areas in OR, revenue management has been phrased succinctly as “selling the right product to the right customer at the right time” (Kimes, 1989). Briefly, firms would like to sell their products to those customers who value them highly so that higher margins may be achieved. However, if sellers wait too long for the arrival of high valuation customers, then they might end the selling season with unsold units that could have been sold to the low valuation customers. For this trade-off to be non-trivial, both stochastic demand and perishable inventory have to hold (Bitran & Caldentrey). We present and experimentally test under controlled conditions and alternative assumptions about changes in the stochastic demand over time a model of dynamic pricing that balances utilization and profitability of the available capacity.

Advances in science and technology have instigated a rapid growth of theoretical and practical activity in the area of dynamic pricing. Numerous theoretical models have been proposed and their implications carefully studied that differ from one another on one or several dimensions including, among others, the seller's capacity, number of sellers, types of buyers, number of products, characteristics of the demand distribution, and pricing strategy. Dynamic pricing has been implemented in the airline, hotel, apparel, grocery, and electronics industries (Talluri & van Ryzin, 2004). However, there have been very few attempts to test the implications of these models under the controlled conditions of the laboratory. As ours is one of the first studies to test optimal pricing models in the laboratory, we begin our program of research with a simple model whose assumptions are described next.

Major Assumptions

Quantity-based vs. Price-based Decisions. Firms may commit themselves to certain quantity or price decisions by deploying capacity in advance, or by advertising prices in advance and thereby limiting their ability to adjust quantities or prices over the selling season. Two major approaches to revenue management have been identified. In the first, the industry (e.g., airline industry) categorizes customers into different classes and focuses the analysis on the allocation of capacity among these classes. The second approach (e.g., retailing industry) focuses on the dynamic pricing policy for one or several products and homogenous customers, who are price- or time-sensitive. Models of revenue management typically treat capacity decisions and pricing decisions separately. In this paper, we focus on price-based revenue management.

The Customers. A major assumption of any model of dynamic pricing concerns the level of sophistication of the customers. We assume that customers are *myopic*, namely, that they purchase the product the first time the price drops below their maximum willingness to pay. This is in contrast to dynamic pricing models that incorporate *strategic* customers, who optimize their purchase behavior in response to the pricing strategies of the sellers. A direct implication of this assumption is that the demand can be treated as *exogenously* given. Customers are assumed to be price-takers; they observe the price listed by the seller and react instantaneously by either buying or not buying the product.

Strategic customers respond to the dynamic pricing policy of the seller by *timing* their purchase. They realize that since the price of a perishable product may change over time, it is possible to wait and purchase the same product at a lower price. However, associated with this delay is the possibility that the product may be sold out or be rendered un-affordable. Although models that assume strategic buyers are more realistic, there are several reasons for assuming that customers are myopic. First, myopic-customer models are much more tractable and, therefore, more widely used. With strategic customers, the dynamic pricing problem becomes a non-cooperative game between the customers and the sellers; with myopic customers it becomes a simpler dynamic optimization problem. Second, in many situations customers do not have the time or information to behave strategically. This is particularly the case with low-priced goods. In fact, customers who shop on-line may now use shopbots that check prices at various websites more or less at the same time (Montgomery et al., 2004; Zhou, Fan, & Cho, 2005). These shopbots typically are not programmed to anticipate changes in future prices at each site. Thirdly (Talluri & van Ryzin, 2004), forecasting models that use observations of past customer behavior and assume a stationary environment already reflect the effects of strategic behavior by the customers.

Monopoly. Another key assumption of the model that we examine in this paper is that the demand faced by the seller only depends on the seller's price and not on the prices of its competitors. Consequently, in this paper we only consider the case of a single seller. Primarily to gain tractability, most of the dynamic multi-period pricing models have focused on the monopoly setting (e.g., Besanko & Winston, 1990; Bitran & Mondschein, 1997; Gallego & Van Ryzin, 1994; Zhao & Zheng, 2000). Models that consider dynamic pricing under competition among several sellers include Granot, Granot, and Mantin (2005, 2006), Netessine and Shumsky (2005), and Perkins and Sood (2004). Talluri and van Ryzin (2004) have provided partial justification to the monopoly model, similar to the one briefly outlined above for the myopic-customer model, by noticing that an observed historical price response often embeds in it the effects of the responses of the competitors to the seller's pricing strategy.

Perishability. Another assumption of our model is that the product is perishable, thereby rendering capacity a time-dependent quantity. We further assume a single product and no replenishment of the inventory during the selling season.

Pricing Policy. We impose the constraint, which is characteristic of many applications of dynamic pricing, that the seller can select prices from a finite list of admissible prices (e.g., Chatwin, 2000; Feng & Gallego, 2000; Feng & Xiao, 2000a, 2000b). The reasons for making this assumption range from marketing considerations such as customers' perception of prices to managerial considerations of ease of implementation. Our model does not restrict the number of price changes; prices may be updated on each period. Nor does it enforce a markdown or markup pricing policy.

Risk-neutrality. Because it is by far the most commonly used formulation in the literature, and because of its mathematical tractability, we assume risk-neutrality in modeling the seller's preferences.

In summary, we study in the laboratory dynamic pricing of a single perishable product over a finite and discrete-time selling horizon, given a fixed inventory at the start of the season that is not replenishable. We assume for simplicity that the salvage value of the product at the end of each season is zero. We assume that the seller is a monopolist, customers are myopic, prices are selected from a fixed list of admissible prices, and the demand is stochastic and time-and price-sensitive. Simplistic as it may be, this model is representative of the type used in modeling the revenue management of style and seasonable goods, where production and ordering cycles are much larger than the sales season. Our independent variable is the change in the demand distribution over the selling horizon. Using a between-subjects design, we compare to one another three conditions. In one of them, for a given price selected by the seller the expected value of the demand is *constant* over all time periods. In a second experimental condition, for a given price selected by the seller the expected value of the demand *decreases* over the selling horizon. In the third condition, the expected demand *increases* over the same selling horizon. In all three conditions, the seller knows in advance the time- and price-sensitive stochastic demand distributions.

Given the widespread use of dynamic pricing software in a variety of industries, and the apparent success of dynamic pricing models in significantly enhancing revenue, what is the motivation for and justification of conducting experiments designed to test the descriptive power of optimal pricing models? Our answer is that in large measure only major national industries (e.g., the airline industry, car rental industry, hotel industry) are implementing (sophisticated, model-based) dynamic pricing methods. We *are not* attempting to address potential pricing biases in those industries that already base their pricing decisions on sophisticated revenue management principles. Instead, we are interested in how firms that do not have access to such software are likely to err in their pricing decisions and how their revenues might be affected by the adoption of revenue management technology. It is instructive to find, document, and subsequently study the implications of any biases, if at all, that characterize the behavior of firms—i.e., the managers in those firms who are responsible for pricing decisions—that do not have access to or do not wish to rely on revenue management technology when making their pricing decisions. As a start, this is best accomplished under the controlled conditions of the laboratory, where the assumptions of the model under investigation are implemented with precision and financially motivated subjects are assigned the roles of sellers and paid contingent on their performance.

There has been very little previous work on revenue management decision making in laboratory settings. One of the few studies is Bearden, Murphy, and Rapoport's (in press) examination of behavior in a capacity control revenue management problem. In the problem they considered, sellers received bids sequentially for units of their (perishable) inventory and had to decide whether or not to exchange their units for the bid amounts. (Although it was not framed in this way to the experimental subjects, the problem was essentially the same as the one faced by airlines deciding when to make available certain fare classes.) The experimental results showed that subjects without access to the optimal decision policy incurred significant revenue losses. The losses resulted from a tendency (or "bias") to be too demanding when holding higher levels of inventory and insufficiently demanding when holding lower levels, with the greatest revenue losses resulting from the latter bias. Importantly, their experimental subjects could *not* affect demand by manipulating price. As price manipulation is one of the most basic means by which managers can affect their revenues, we believe that it is important to examine carefully the price setting behavior of sellers in a controlled setting.[1] We do so in the current paper.

The rest of the paper is organized as follows. Section 2 presents the model and the optimal dynamic pricing policy for each of our three experimental conditions. Section 3 begins with a description of the experimental procedure, and Section 4 continues with the presentation and analysis of the results. Section 5 concludes.

2. The Model

We assume a single decision maker (DM) with a fixed number of units N to sell over a *season* of T (t=1, …, T) discrete time periods. We do not allow for replenishment of inventory, there are no costs associated with holding (i.e., carrying) inventory, and the units have a salvage value of 0, that is, they are worthless at the conclusion of the selling horizon. The number of units remaining in inventory at any particular time is denoted by n ($0 \le n \le N$). At the beginning of each time period t, the DM can set her unit price p_t for the good from a fixed, finite set of prices $\boldsymbol{P}$. The price she chooses then determines the (stochastic) demand in that period. Let $\lambda(p,t) = \alpha_t e^{-\beta p}$, where $\alpha, \beta > 0$, be the *demand rate function* for each period. The realized demand d in a given period with price p follows a Poisson distribution with mean $\lambda(p,t)$, and the number of sales is given by min(n,d). The DM's objective is to maximize her expected revenue by dynamically setting the selling price for her units in each period.

The DM can determine her optimal pricing policy for each feasible *state* (t,n) by solving the following dynamic program for her (optimal) expected revenue-to-go:

$$V_t\ (n) = \max_{p \in P} \left\{ \sum_{d=0}^{\infty} \frac{e^{-\lambda(p,t)} \lambda(p,t)^d}{d!} \left(\min(n,d)\, p + V_{t+1}(n - \min(n,d)) \right) \right\}, \qquad (1)$$

with boundary conditions

[1] There is a superficially related literature on price competition in experimental economics. The experiments have focused on simplified pricing games such as the Bertrand model in which sellers make a single pricing decision, and there is no dynamic element (for a review of this literature, see Holt, 1995). Since it is precisely the dynamic ele-

$$V_{N+1}(n) = 0, \forall n, \text{ and}$$
$$V_t\ (0) = 0, \forall t.$$

The dynamic program is solved backwards in time to find the optimal price p^* for each state (t,n). This model is a straightforward Poisson extension of the Bernoulli demand dynamic pricing model presented in Talluri and van Ryzin (2004). Bitran and Mondschein (1997) examined a special case of the Poisson demand model with a discrete set of prices in which the DM is constrained to employ prices that are non-increasing in time (i.e., to use markdown pricing). We do not impose this restriction.

In the experiment described below, we examined actual, in contrast to optimal, pricing behavior in the dynamic pricing problem under three different experimental conditions. The conditions varied from one another only with respect to the time-sensitivity of the demand rate function. We examined situations in which *ceteris paribus* expected demand was *constant* over time, *increasing* in time, and *decreasing* in time. Recall that the demand rate function is given by $\lambda\left(p,t\right) = \alpha_t e^{-\beta p}$. We manipulated the time-sensitivity of the demand distribution by adjusting (or not) the α term over time. Specifically, the demand manipulations were operationalized as follows:

♦ In Condition CON (for "constant"), we assume that $\lambda(t, p)$ is constant in t for any value of $p \in \mathbf{P}$. Put differently, $\alpha_t = \alpha_0$ for all t.

♦ In Condition DEC (for "decreasing"), for any value of $p \in \mathbf{P}$ we assume that $\lambda(t, p)$ is decreasing in t. Specifically, the demand is influenced by the following dynamics:

$$\alpha_t = \left(1 - \frac{t}{T+1}\right)\alpha_0 .$$

♦ In Condition INC (for "increasing"), we assume that $\lambda(t, p)$ is increasing in t for any $p \in \boldsymbol{P}$. The increase in the demand function is governed by

ment of the decision problem that makes dynamic pricing so important, these previous studies do little to inform the more general issue of decision behavior in multi-period perishable asset pricing problems.

$$\alpha_t = \frac{t}{T+1}\alpha_0 .$$

Condition CON is mainly a control condition where, for a specified price set by the seller, the stochastic demand distribution is time invariant. In the other two conditions, the demand distribution changes over the planning horizon. Condition DEC is designed to model cases in the fashion industry where, for a given price, there are more early adopters than those who purchase the good toward the end of the season. Condition INC is more typical of the airline industry where, for a given price of a ticket demand may be increasing toward the scheduled date of the flight.

Optimal pricing decisions for these three instances of the dynamic pricing problem are presented in Tables A1, A2, and A3 of Appendix A, corresponding to problems with constant, decreasing, and increasing demand, respectively. The parameter values on which the solutions are based are presented in Table 1. Each table presents the optimal price that the seller should charge as a function of the period number (1-10) and level of inventory (1-21). Thus, when the season starts with 10 periods to go and initial capacity of 21 units, the optimal price for Conditions CON, DEC, and INC is 15, 14, and 14, respectively (see lower left-hand corner of each table). For a second example, if there are 6 more periods before the season terminates (including the current one) and the seller's inventory is 14, then she should price the good at 15, 12, and 17 in Conditions CON, DEC, and INC, respectively. Note that for the same level of inventory, the optimal DM should never increase her price as she nears the end of the season. Similarly, within any particular period, the optimal DM should not increase her price as her inventory level increases. These two properties of the optimal policy provide straightforward, demand-free, qualitative predictions that are experimentally testable. Next, we describe a laboratory experiment in which we studied the behavior of actual decision makers in the dynamic pricing problem.

3. A Dynamic Pricing Experiment: Method

Experimental Design. There were three experimental conditions corresponding to increasing (INC), decreasing (DEC), and constant (CON) expected demand. Each subject participated in only a single experimental condition. The parameter values for each experimental condition are shown in Table 1. In each condition, there were T=10 time *periods* in each of 40 identical *seasons*. The *initial inventory* for each season was set at N=21 units. The subjects were charged 10 *cost units* (experimental dollars) for each unit of inventory at the beginning of each season. This resulted in a fixed inventory cost of 210 units that was subtracted from the subject's revenue at the end of each season to determine that season's *profit*. Consequently, the subject could conclude a season with either a positive or negative payoff. In each period, the subject could choose a price from the *price set* $\boldsymbol{P}$={1,2,…,20}. The unit costs were sunk and did not affect the optimal pricing policy: The expected profit maximizing policy is the same as the expected revenue maximizing policy.

--Insert Table 1 about here--

Subjects. Sixty-nine subjects participated in the experiment, 23 in each of the three conditions. The subjects were recruited through a computerized recruitment system requesting volunteers to participate in an individual decision making task on revenue management for payoff contingent on performance. Both male and female subjects, almost all of them business undergraduate students, responded in nearly equal numbers. Individual sessions lasted between 40 and 90 minutes.

Procedure. The subjects were run individually in a computer-controlled experiment. After being seated in separate rooms, they were handed written instructions and then proceeded to read them at their own pace. Questions were answered individually by the experimenter.

The instructions describing the experimental task (see Online Appendix) were detailed. They first presented the subjects with an overview of the "Pricing Game" that provided a non-technical, global description of the dynamic pricing problem. This was followed by a non-technical explanation of a representative Poisson distribution of the demand. Then, the various screens presented by the software were explained one at a time in considerable detail. Particular emphasis was placed on presentation and explanation of the price- and time-sensitive demand distributions. For any feasible price level, the subject could select and examine the demand distribution that would be operative in the current and also in each of the subsequent time periods. The demand distributions for each combination of price and period were displayed graphically (as bar charts) one at a time by the experimental software. To reinforce the demand dynamics, each subject was presented with a hard (paper) copy of the visual display of all the demand distributions for selected price values of p=1, p=10, and p=20 for each period t=1 ,…, 10. Subjects could keep the hard copy for the entire experiment and refer to it when needed.

Complete outcome information was provided to the subjects throughout each season. In each period, they were shown the price, demand, and revenue history for each observed period in the current season. The dynamic pricing game was iterated with the same parameter values (but with different random draws of the stochastic demand) for each of the 40 seasons. Prices were stated in US dollars. At the end of the session, each subject was paid 20 percent of the average of her cumulative profits across the 40 selling seasons plus a $5 show up fee. Mean payoff for Conditions CON, DEC, and INC was $12.33, $7.86, and $9.40, respectively.

4. A Dynamic Pricing Experiment: Results

There are three independent variables in the experiment that jointly affect the subject's pricing decisions, namely, the demand distribution, period within a season, and seller's inventory. The first two variables are determined exogenously, whereas the level of inventory on period t is determined endogenously and jointly by the seller's pricing decision and the stochastic demand on period t-1. We focus the analysis in this section on three major dependent variables, namely, the pricing decision, the price error (defined below with respect to the optimal prices listed in Appendix A), and the implied revenue loss (see below). We begin this section with presentation of aggregate results computed across periods to uncover general trends across seasons, continue with the statistical comparison of the three conditions to one another, and later proceed with a more detailed analysis of within-season behavior. In all the analyses that we report, the subject is the statistical unit of analysis. When we report means for the different conditions, they correspond to the means of the subjects' dependent measures (e.g., means of regression coefficients computed separately for each subject). Similarly, measures of dispersion are computed across the subjects' dependent measures.

In all three conditions, we opted not to provide the subjects with practice seasons but to let them learn with experience the mechanics of the game, the effects of their pricing decisions on revenue and inventory level, and, in particular, comprehend the stochastic nature of the demand. Inexperienced students facing the pricing game for the first time require considerable experience in order to acquire a good grasp of the problem, as are flesh and blood sellers in the real world. There is no simple criterion to determine when the learning phase is over. Our interest in the present study, in part dictated by our goal to draw practical implications of our findings, is in steady-state behavior. As we show below, the subjects' behavior stabilized after about 15 seasons. Therefore, in what follows we shall consider the first 15 seasons as "practice" seasons and focus the analysis on the pricing decisions in seasons 16-40.[2]

Overall Performance across Seasons. The top left panel of Fig. 1 exhibits the mean price for each of the three experimental conditions as a function of the season number. In all three conditions, there is a clear increase in prices during the early seasons. Mean prices tend to stabilize around season 15.

To obtain a comprehensive description of the outcomes of these decisions, we computed the mean revenue (across the 23 subjects in each condition) for each condition and each season separately. Because the means tended to fluctuate quite drastically due to the stochastic problem instances, thereby obscuring the general trend over sessions, we computed the running means of the revenues in steps of 5 (seasons 1-5, 2-6, … , 36-40). The right left panel of Fig. 1 displays these running means for each of the three conditions. Examining the mean revenues across seasons is not exceedingly informative. In general the curves look rather flat, suggesting no improvement. And, further, based on the moving averages, it looks as if performance in the CON condition decreases in early trials and increases in later trials; this, however, is misleading: we found that the quadratic term was not significant in a model that regressed mean revenues

[2] The qualitative characteristics of the summary statistics we report below are unaffected if we use all of the data, rather than only the last 25 seasons. However, including the early seasons -- in which subjects were familiarizing themselves with the task -- does increase substantially the dispersion in the data.

onto season and squared season. (A small number of periods in middle seasons with unusually low earnings were sufficient to affect the curves of the moving averages.) As we show below, a much clearer picture emerges when one examines the revenue losses that result from *each* individual decision.

--Insert Fig. 1 about here--

Price error in a pricing decision is defined as the difference between the observed and optimal price (price error = empirical price – optimal price). We computed the price error for each period, and then averaged these ten measures for each subject. The bottom left panel of Fig. 1 shows that the mean price error starts out quite negative in all three conditions and moves toward 0 as the subjects gain experience with the pricing problem, reaching a roughly steady state around season 15. Regressing the mean price error onto the season number for the first 15 seasons resulted in significant and positive slope coefficients in each condition (β = 0.19, 0.08, and 0.18 for conditions CON, DEC, and INC, respectively; $p<0.05$ in each case). However, repeating the same analysis for seasons 16 through 40 could not reject the null hypothesis of zero regression coefficients (β= 0.01, 0.02, and 0.01, $p>0.05$ in each case). This is an additional evidence for focusing the analysis on the pricing decisions in seasons 16-40.

For a given period t, consider the *optimal* expected revenue-to-go, which is obtained by solving the dynamic program described above, when the price is set at p^*. Next, consider the "*actual*" expected revenue-to-go, which is obtained by the same dynamic programming procedure, when the price is set at the observed price p, conditional on pricing optimally for the rest of the season. Define the *implied revenue loss* for a given price p as the difference between the optimal and actual revenue-to-go. The implied revenue loss is a conservative measure of revenue loss due to possible differences between optimal and actual prices, as it assumes optimal pricing on the part of the subject in all future periods. The primary virtue of this measure is that it allows the estimation of revenue consequences of each individual decision. The bottom right panel of Fig. 1 displays the mean implied revenue losses for each of the three experimental conditions as a function of season. Similarly to the other dependent measures depicted in Fig. 1, we observe that the implied revenue loss measure tends to stabilize around season 15.

We next present the results of statistical tests on these measures that compare the three conditions to one another. Table 2 presents the major statistics by condition.

--Insert Table 2 about here--

Price. The mean prices for seasons 16-40 were subjected to a one-way ANOVA to test the null hypothesis of equality of mean prices between conditions. The ANOVA yielded a significant condition effect, $F(2,66)=3.38$, $p=0.04$. Post-hoc comparisons revealed that the significant condition effect was due to the difference between Conditions CON and INC. Specifically, mean prices were significantly higher in Condition CON than INC. No other pair-wise comparisons were significant. It should be noted that the mean price measure ignores the relationship between the price on one hand and the period number and inventory levels on the other hand. We will examine these relationships more closely below when we study within-season behavior.

Season Revenue. Expected revenue in our experiment is not very sensitive to the differences between the conditions. The optimal expected season revenue is 267 in the Conditions CON and 250 in Conditions DEC and INC. These theoretical values are shown in brackets in Table 2. Computing the difference between actual mean revenues and the appropriate optimal expected revenue, we find no difference in mean (total) revenue losses across the experimental conditions, $F(2,66)=1.97$, $p=0.15$. Importantly, the mean seasonal revenue is significantly smaller than the optimal expected revenue in *all* three conditions (see upper right panel in Fig. 1). On average, the subjects in all three conditions earned significantly less than they would have under the application of the optimal policy.

Price Error. The price errors were also subjected to a one-way ANOVA to test the null hypothesis of equality between the means in the three conditions. The ANOVA revealed a significant condition effect on mean price error, $F(2,66)=3.43$, $p=0.04$. Post-hoc comparisons revealed that the mean price error was significantly smaller in Condition DEC than in Condition INC. However, as we show below, simply averaging price error ignores important information. A much richer picture emerges when the relationship between price error and period number is examined.

Implied Revenue Loss. There was no significant difference between the three conditions in mean implied revenue losses, F(2,66)=1.02, *p*=0.37. However, in all three experimental conditions, the mean implied revenue losses were significantly greater than 0 (see Table 2).

End of Season Inventory. Another dependent variable that we examined is the inventory (number of unsold units) at the end of the season. We simulated optimal play in 100,000 iterations of the pricing problem to obtain accurate estimates of the end-of-season inventory. The expected end-of-season inventory under the optimal play was found to be 2.76, 4.62, and 2.77 for Conditions CON, DEC, and INC, respectively (Table 2). Using one-way ANOVA once again, we could not reject the null hypothesis that the difference between optimal and observed end-of-the season mean inventory was the same for all conditions, F(2,66)=0.90, *p*=0.41.

Within Season Behavior. Our primary interest is with within-season behavior. We wish to determine what factors affect pricing decisions, departures of prices from optimal (revenue maximizing) prices, and the losses resulting from these departures. To address these issues, we employ multiple regression methods to isolate the effects of period and inventory levels on pricing decisions. A systematic comparison of several models that included a number of independent variables (e.g., cumulative revenue, number of units sold in previous period, whether profit was positive) lead us to conclude that a simple model including only period, inventory, and a period by inventory interaction term as independent variables best characterizes the data.

We utilized the random regression approach to modeling the data (see, e.g., Lorch & Myers, 1990). Under this approach, one fits a regression model to each subject separately and uses the resulting estimated regression coefficients in subsequent analyses. Thus, for example, one may compare regression coefficients within a condition to some hypothesized value (e.g., using *t*-tests), or compare the regression coefficients across different conditions (e.g., using ANOVAs). Using the random regression procedure alleviates the problem of dependency between individual subjects' responses. It also provides measures of subject heterogeneity (from the dispersion of the estimated coefficients).[3]

There is considerable subject heterogeneity in pricing policies. Subjects vary a great deal in their mean prices, and individual subjects show (seemingly non-systematic) variability in their mean prices across seasons. Because our main interest is in how subjects' pricing decisions respond to *changes* in period number and inventory levels – and not necessarily in the absolute relationship between prices and these variables – we normalized each subject's responses with respect to his or her overall pricing behavior. Not doing so significantly inflates standard errors and thereby significantly reduces the statistical power of the tests. We explain the normalization procedure below. Summary results from the regressions that we describe below are presented in Table 3.

--Insert Table 3 about here--

Price. The left panel of Fig. 2 displays the mean price for each condition as a function of period, whereas the right panel displays the mean price as a function of level of inventory. At the aggregate level – i.e., averaging over all subjects in each condition – it is impossible to discern (properly) how price relates to period number and inventory level at the individual subject level. Possible selection biases may be introduced due to different pricing policies which, in turn, affect the seller's inventory on each period. Consider the mean prices for Condition INC as a function of period. It appears that mean prices are increasing in period, but the inference one draws from this may be misleading. To show that, suppose that there are just two types of subjects, low price takers and high price takers in Condition INC, and that each type uses only a single price throughout the season. Since low price takers will tend to sell out sooner (and therefore have fewer observations in later periods), at the aggregate level one might observe that mean prices are increasing in period even when no individual subject increases her price as the end of the season approaches. Furthermore, since inventory and period are so strongly (negatively) correlated -- in early

[3] We also used hierarchical linear modeling (HLM) methods to analyze the data, and reached the same general conclusions.

periods one tends to have more inventory than in later periods -- it is not evident from the figure whether subjects are affected by *both* period number and inventory level or only by one of them. This scenario demonstrates the difficulty of only considering aggregate data in the pricing problem. By employing the random regression approach, we can obtain *independent estimates* of the effects of period and inventory on pricing decisions for each individual subject.

--Insert Fig. 2 about here--

To examine the effects of period and inventory on *price,* we fit the following model to each subject:

$$\left(p_{i,t,s} - \overline{p}_{i,s}\right) = \beta_{i,0} + \beta_{i,t}t + \beta_{i,n}n + \beta_{i,tn}\left(t \cdot n\right) + \varepsilon_i ,$$

where $p_{i,t,s}$ is the price set by subject i in period t of season s, and $\overline{p}_{i.s}$ is the mean price of subject i in season s. Subtracting $\overline{p}_{i.s}$ from each price effectively normalizes responses for the subjects. Doing so partially factors out individual subjects' general tendencies to price higher or lower (with respect to one another), and within subjects it accounts for season to season fluctuations in pricing policies.[4]

As shown earlier, the optimal policy implies that the subjects should decrease their prices as the period (t) increases, and also as the inventory level (n) increases. Table 3 presents the mean period and inventory slope coefficients for each experimental condition. Consistent with the predictions of the optimal model, each mean period and inventory coefficient is significantly negative ($p<0.05$ in each case, based on testing the null hypotheses of zero regression coefficients by t-tests). Thus, in all conditions, subjects tended to set low prices as they moved toward the end of the sales season. They also tended to lower their prices when they had lower levels of inventory. A one-way ANOVA on the period coefficients reveals a significant condition effect, $F(2,66)=5.82$, $p=0.004$. Post-hoc comparisons revealed that the means for

[4] We realized the necessity to adjust for seasonal fluctuations after examining the residuals of the regressions for each individual subject. There are clear seasonal fluctuations in some subject's pricing policies: In some seasons they tend to price higher on average than in others, resulting in autocorrelated residuals when season is not adjusted for. One alternative way of accounting for these seasonal factors is to include a dummy variable for each season; however, this approach introduces computational difficulties in estimating coefficients because of the high degree of collinearity among the dummy variables. Adding a single seasonal term (which would suppose a linear relationship across seasons) does not tend to alleviate the seasonal problem.

conditions DEC and INC differed significantly from each other. No other pair-wise comparisons were significant. On average, compared to subjects in Condition INC, the subjects in Condition DEC tended to lower their prices more as the period increased. This is to be expected since, *ceteris paribus*, expected demand was decreasing in period in the latter and increasing in the former condition.

We find a similar pattern of results for the inventory level coefficients. The effect of condition is significant, $F(2,66)=3.14$, $p<0.05$, and the only significant pair-wise difference is between Conditions INC and DEC. However, we cannot conclude from these results whether the differences in responsiveness to period and inventory were more or less than what we would expect, given the differences between the demand distributions in these two conditions and their respective optimal policies. To do so, we consider next the price error, which implicitly allows us to take the difference in the pricing decisions into consideration when evaluating the subjects' policies.

Price Error. Figure 3 displays the aggregate price error results. The most salient feature of this figure (left panel), and a major finding of this research, is that price errors tend to proceed from positive (negative) to negative (positive) in Conditions DEC (INC) as the end of the season approaches. In other words, on average, the subjects tended to overprice (under-price) early on and to under-price (over-price) later on in Condition DEC (INC). The aggregate results for the price error are much more ambiguous when considered as a function of inventory level rather than period (right panel of Fig. 3). Because they are potentially misleading, the propensities of the individual subjects must be examined separately in order to construct a complete picture of behavior in the pricing problem.

--Insert Fig. 3 about here--

As we did earlier when examining price setting, we factored in the seasonal variability in each subject's prices when examining the effects of period and inventory on price errors. The random regression model that we tested was:

$$\left[\left(p_{i,t,s} - p^{*}\right) - \overline{p}_{i,s}\right] = \beta_{i,0} + \beta_{i,t}t + \beta_{i,tn}\left(t \cdot n\right) + \varepsilon_{i} ,$$

where $p*$ is the appropriate optimal price. Table 3 presents the mean coefficients for the model. There are several notable findings. First (column 3), all of the mean coefficients are significantly different from 0. Second, in Conditions CON and INC the coefficients on period are positive, whereas in Condition DEC it is negative (column 4). A one-way ANOVA on the period coefficient revealed a significant condition effect, $F(2,66)=32.35$, $p<0.001$; all the pair-wise post-hoc comparisons were significant. Taken together, these results imply that there is a general tendency to over-price early on (i.e., in early periods) and to under-price later on in Condition DEC. This pattern is reversed in Condition INC. There is a tendency to under-price throughout in Condition CON, but this tendency decreases somewhat with period.

The qualitative effect of inventory on price error is the same in all three conditions: the difference between the actual and optimal price tends to increase as the inventory levels increase (column 5). Comparing the three conditions to one another, a one-way ANOVA revealed a significant condition effect, $F(2,66)=14.70$, $p<0.001$, all the pair-wise post-hoc comparisons were significant. Our results show that the effect of inventory on price errors tends to increase as one moves from Condition DEC to CON and then to Condition INC.

Implied Revenue Loss. Figure 4 exhibits the mean implied revenue loss as a function of the period (left panel) and inventory level (right panel). To examine the implied revenue losses at the individual level – and with some abuse of notation – we fit the following random regression model:

$$\left(V\left(p^*\right)-V\left(p\right)\right)=\beta_{i,0}+\beta_{i,t}t+\beta_{i,n}n+\beta_{i,tn}\left(t\cdot n\right)+\varepsilon_i ,$$

where $V(p^*)$ is the optimal expected revenue-to-go, and $V(p)$ is the expected revenue-to-go for pricing at (a possibility non-optimal price) p, conditional on optimal pricing thereafter. The summary results for each experimental condition are shown in Table 3 (column 6). None of the coefficients of the model are significantly different from 0, and there is no difference between the conditions (all $F<1$ in each case). We conclude that the implied revenue losses resulting from the pricing errors are not systematically related to the period and inventory level.

--Insert Fig. 4 about here--

Inventory Levels. Figure 5 portrays the mean end-of-period inventory levels for each experimental condition as well as the expected inventory levels under the optimal decision policy. The curves are quite close in Condition CON. In Condition DEC, we observe that the subjects tend to carry more inventory across periods than expected under optimal play. This pattern is reversed in Condition INC with the subjects carrying less inventory than they would under the optimal policy. These behavioral patterns are, of course, consistent with (as they are a consequence of) the pricing errors across the three conditions.

--Insert Fig. 5 about here--

4. Conclusions

The participants in our experimental study of the dynamic pricing problem employed pricing policies that were qualitatively consistent with the optimal dynamic pricing policy. They tended to decrease their prices as they approached the end of the season and also when they held more inventory. This pattern obtained in the cases of constant, decreasing, and increasing expected demand. However, even after the learning phase in seasons 1-15 was more or less completed, there were systematic departures from the optimal pricing policies that resulted in significant losses of revenue. The general pattern of results is consistent with the pricing bias to respond in an *exaggerated* manner to *relative demand*. Subjects in Condition DEC, where the expected demand decreased over time, tended to overprice their assets when the demand was relatively high (early on) and under-price them when it was relatively low (later on). This pattern was reversed in Condition INC, where the expected demand increased over time. Subjects priced their assets too high when the stochastic demand was relatively high (later on) and too low when it was relatively low (early on). In the case of constant demand there was a slight bias to under-price *across* all periods, but this bias tended to diminish as the subject neared the end of the season.

What is the significance of these findings? One may consider the experimental results from at least two different perspectives. Experimental economists tend to emphasize how behavior does or does not correspond to the predictions of economic theory. They are mostly concerned with documenting how behavior is biased with respect to normative theory. From this perspective, one might focus on how subjects' decision behavior departs from optimality. Alternatively, OR researchers might prefer to focus on how much an optimal dynamic pricing model *could* help unaided decision makers increase their revenues. Under this more positive frame, one emphasizes the benefit of the model rather than the cognitive biases and bounded rationality of the decision maker. We regard these two perspectives as complementary. Although both are present in our study, our focus here is on the latter perspective. It clarifies the potential utility of behavioral studies of decision making for OR theorists. Laboratory experiments can provide a relatively clean (i.e., confound free) and inexpensive means of demonstrating the potential worth of sophisticated OR algorithms. What use would a complicated algorithm be if it performed not significantly better than a manager using her intuition?

The line of research presented here can be extended in a number of ways to examine the generality of our conclusions. One pressing question that we are currently examining is whether systematic pricing biases persist in the presence of competition. To address this question, we have been studying dynamic pricing decisions in a laboratory setting in which sellers compete with each other for (simulated) myopic customers. Other possible extensions include studying the strategic timing of purchases of costumers (i.e., experimental subjects) who are faced with sellers who price their product dynamically. Perhaps by gaining a better understanding of actual buyer behavior, modelers can include more behaviorally plausible assumptions in their theories and practitioners in their revenue management algorithms. There is no shortage of opportunities for research that combines theoretical and behavioral approaches to the study of decision behavior in management settings.

Acknowledgments

We gratefully acknowledge financial support by a contract F49620-03-1-0377 from the AFOSR-MURI to the Department of Systems and Industrial Engineering and the Department of Management and Organizations at the University of Arizona.

References

Bearden, J. N., Murphy, R. O., and Rapoport, A. (in press). Decision biases in revenue management: Some behavioral evidence. *Manufacturing & Service Operations Management*.

Besanko, D., and Winston, W. (1990). Price optimal skimming by a monopolist facing rational customers. *Management Science, 36,* 555-567.

Bitran, G. R. and Caldentrey, R. (2003). An overview of pricing models for revenue management. *Manufacturing & Service Operations Management, 5,* 203-229.

Bitran, G. R., and Mondschein, S. V. (1997). Periodic pricing of seasonal products in retailing. *Management Science, 43,* 64-79.

Chatwin, R. E. (2000). Optimal dynamic pricing of perishable products with stochastic demand and a finite set of prices. *European Journal of Operations Research, 125,* 149-174.

Feng, Y., and Gallego, G. (2000). Perishable asset revenue management with Markovian time dependent demand intensities. *Management Science, 46,* 941-956.

Feng, Y., and Xiao, B. (2000a). A continuous-time yield management model with multiple prices and reversible time changes. *Management Science, 48,* 644-657.

Feng, Y., and Xiao, B. (2000b). Optimal policies of yield management with multiple predetermined prices. *Management Science, 49,* 938-949.

Gallego, G., and Van Ryzin, G. J. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science, 40,* 999-1020.

Granot, D., Granot, F., and Mantin, B. (2005). A dynamic pricing model under duopoly competition. University of British Columbia, Sauder School of business, unpublished manuscript.

Granot, D., Granot, F., and Mantin, B. (2006). Revenue management of perishable products under competition. University of British Columbia, Sauder School of business, unpublished manuscript.

Holt, C. (1995). Industrial organization: A survey of laboratory research. In *Handbook of Experimental Economics,* J. Kagel and A. Roth (Eds.). Princeton University Press.

Kimes, S. E. (1989). Yield management: A tool for capacity-constrained service firm*. Journal of Operations Management, 8,* 348-363.

Lorch, R. F., and Meyers, J. L. (1990). Regression analyses of repeated measures data in cognitive research. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 149-157.

McGill, J. I., and Van Ryzin, G. J. (1999). Revenue management: Research overview and prospects. *Transportation Science, 33,* 233-256.

Montgomery, A. Hosanagar, K, Krishnan, R., and Clay, K. (2004). Designing a better shopbot. *Management Science, 50,* 189-206.

Netessine, S., and Shumsky, R. A. (2005). Revenue management games: Horizontal and vertical competition. *Management Science, 51,* 813-831.

Perakis, G., and Sood, A. (2003). Competitive multi-period pricing with fixed inventories. Massachusetts Institute of Technology, Operations Research Center, unpublished manuscript.

Tallluri, K., and Van Ryzin, G. J. (2004). *The Theory and Practice of Revenue Management*. New York: Kluwer.

Zhao, W., and Zheng, Y. S. (2000). Optimal dynamic pricing for perishable assets with nonhomogenous demand. *Management Science, 46*, 375-388.

Zhou, Y,-P., Fan, M., and Cho, M. (2005). On the threshold purchasing behavior of customers facing dynamically priced perishable products. University of Washington, Dept. of Management Science, unpublished manuscript.

Table 1. Parameters employed in each condition of the dynamic pricing experiment

	T	***N***	***C***	**P**	α_0	β
COND	10	21	10	{1, 2,…, 20}	6.50	0.085
DEC/INC	10	21	10	{1, 2,…, 20}	12	0.085

.

Table 2. Mean price, season revenue, price error, implied revenue loss, and end-of-season inventory for each condition for seasons 16-40

		Condition	
	CON	**INC**	**DEC**
Price	14.73 (0.25)	13.97 (0.26)	13.61 (0.40)
Season Revenue	246.65 (2.48) [267]	224.28 (3.11) [250]	231.49 (2.36) [250]
Price Error	-0.40 (0.40)	-0.43 (0.35)	-1.21 (0.55)
Implied Revenue Loss	1.27 (0.13)	1.62 (0.16)	1.53 (0.23)
Ending Inventory	2.97 (0.39) [2.76]	4.97 (0.27) [4.62]	3.61 (0.35) [2.77]

Table 3. Mean regression coefficients for each experimental condition and each dependent variable

	Condition	**Mean $\beta_{i,0}$**	**Mean $\beta_{i,t}$**	**Mean $\beta_{i,n}$**	**Mean $\beta_{i,tn}$**
Price	CON	8.50* (1.54)	-0.77* (0.18)	-.0.37* (0.07)	0.004 (0.006)
Price	DEC	8.75* (1.17)	-1.26* (0.19)	-0.23* (0.05)	0.013 (0.008)
Price	INC	6.80* (1.52)	-0.35* (0.20)	-0.44* (0.07)	0.025* (0.008)
Error	CON	-24.02* (1.44)	0.68* (0.17)	0.40* (0.06)	0.002 (0.005)
Error	DEC	-14.20* (1.53)	-0.44* (0.20)	0.16* (0.07)	0.026* (0.007)
Error	INC	-31.68* (1.60)	1.70* (0.19)	0.66* (0.07)	-0.034* (0.008)
Implied Revenue Loss	CON	0.041 (1.56)	0.17 (0.16)	0.03 (0.07)	-0.003 (0.005)
	DEC	1.96 (1.06)	-0.03 (0.11)	0.05 (0.05)	-0.020* (0.006)
	INC	1.51 (1.54)	0.06 (0.16)	0.03 (0.06)	0.005 (0.006)

Entries marked with (*) are significant at the 0.05 level. The values in parentheses represent one standard error of the mean.

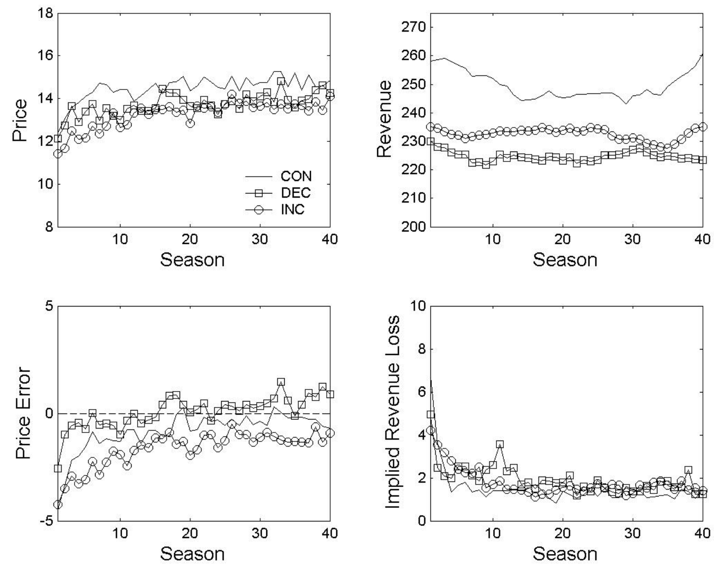

Figure 1. Mean price, season revenue (5 season moving average), price error, and implied revenue loss across the 40 experimental seasons for each experimental condition.

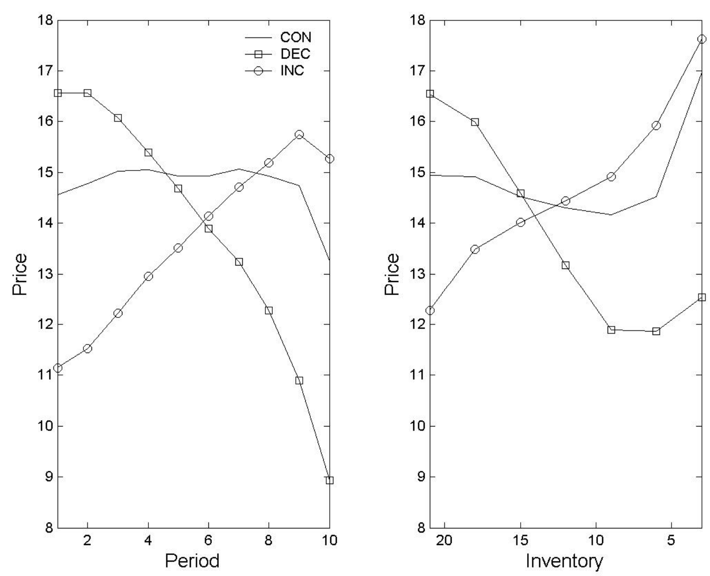

Figure 2. Mean price as a function of period and inventory level for each experimental condition.

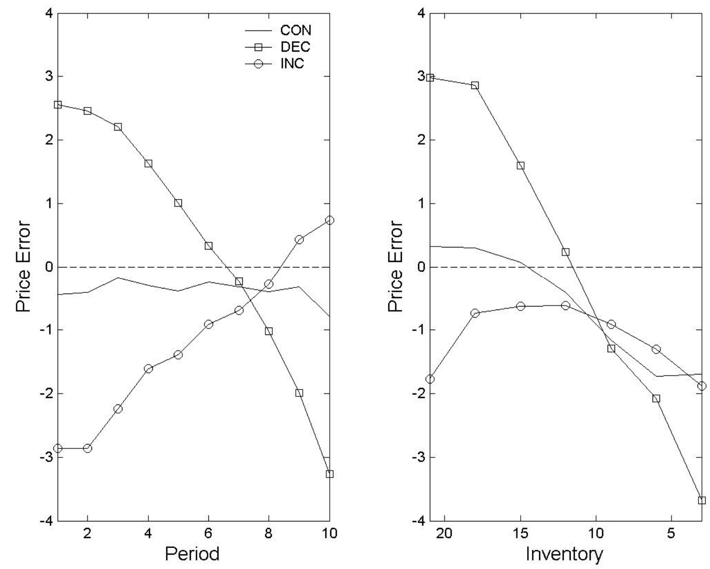

Figure 3. Mean price error as a function of period and inventory for each experimental condition

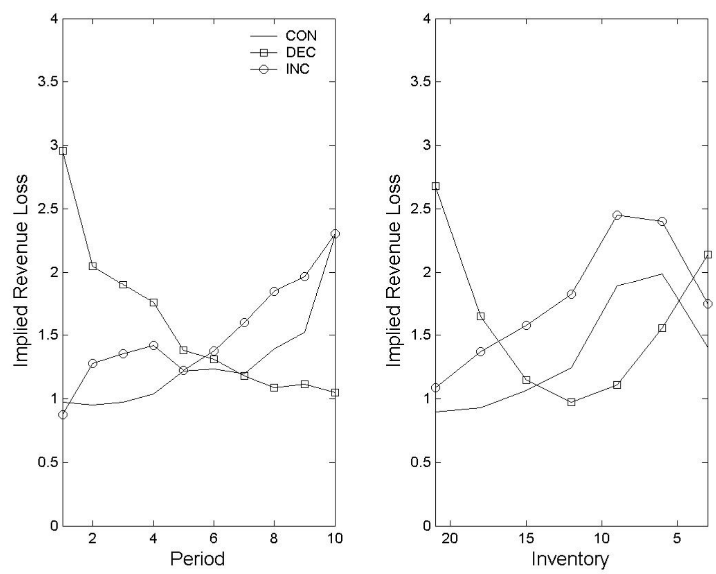

Figure 4. Mean implied revenue loss as a function of period and inventory for each experimental condition

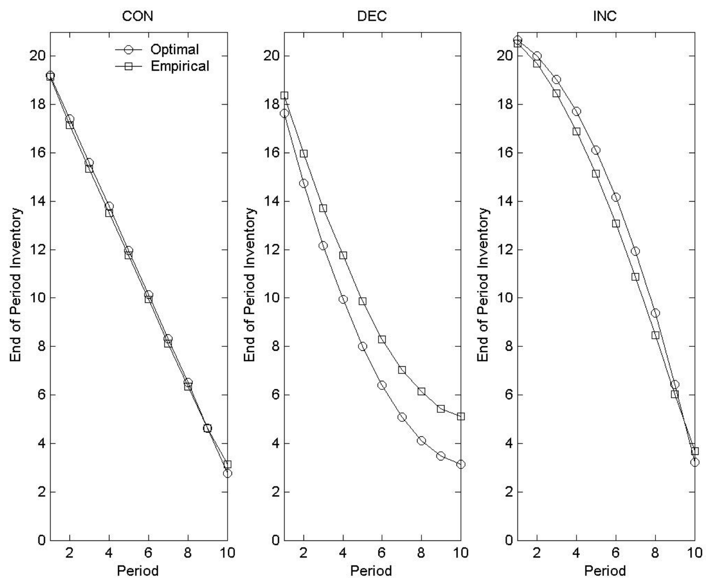

Figure 5. Mean end-of-season inventory under the optimal policy for each experimental condition

Appendix A

Table A1. Optimal pricing policies for Condition CON for the parameter values shown in Table 1.

						Period					
		1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
	1	20	20	20	20	20	20	20	20	20	20
Inventory	**2**	20	20	20	20	20	20	20	20	20	*17*
	3	20	20	20	20	20	20	20	20	20	*15*
	4	20	20	20	20	20	20	20	20	*18*	*13*
	5	20	20	20	20	20	20	20	19	*16*	**13**
	6	20	20	20	20	20	20	20	*17*	**14**	*12*
	7	20	20	20	20	20	20	*18*	*16*	*13*	*12*
	8	20	20	20	20	20	19	*17*	**15**	*13*	12
	9	20	20	20	20	19	*17*	*16*	*14*	*12*	12
	10	20	20	20	20	18	*16*	**15**	*13*	12	12
	11	20	20	20	19	*17*	*15*	*14*	*13*	12	12
	12	20	20	19	18	*16*	**15**	*13*	12	12	12
	13	20	20	18	*17*	*15*	*14*	*13*	12	12	12
	14	20	19	17	*16*	**15**	*13*	12	12	12	12
	15	19	18	16	*15*	*14*	13	12	12	12	12
	16	18	17	*16*	**15**	*14*	13	12	12	12	12
	17	17	16	**15**	*14*	13	12	12	12	12	12
	18	17	*16*	*15*	*14*	13	12	12	12	12	12
	19	16	**15**	*14*	13	12	12	12	12	12	12
	20	15	*15*	14	13	12	12	12	12	12	12
	21	**15**	*14*	13	13	12	12	12	12	12	12

Inventory levels marked in bold in each period represent the mean inventory levels from a 5000 trial simulation employing the optimal pricing policy. The italicized inventory levels represent one standard deviation above and below the mean.

Table A2. Optimal pricing policies for Condition DEC for the parameter values shown in Table 1.

						Period					
		1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
	1	20	20	20	20	20	20	20	20	*19*	*14*
Inventory	**2**	20	20	20	20	20	20	20	*19*	*15*	*12*
	3	20	20	20	20	20	20	19	*16*	*13*	**12**
	4	20	20	20	20	20	20	*17*	*14*	**12**	*12*
	5	20	20	20	20	20	*18*	*15*	**13**	*12*	*12*
	6	20	20	20	20	19	*16*	**13**	*12*	*12*	*12*
	7	20	20	20	20	*17*	*14*	*13*	*12*	*12*	12
	8	20	20	20	18	*16*	**14**	*12*	*12*	12	12
	9	20	20	20	17	*15*	*13*	*12*	12	12	12
	10	20	20	18	*16*	**14**	*12*	12	12	12	12
	11	20	20	17	*15*	*13*	*12*	12	12	12	12
	12	20	18	16	**14**	*13*	12	12	12	12	12
	13	20	17	*15*	*14*	*12*	12	12	12	12	12
	14	19	17	*15*	*13*	12	12	12	12	12	12
	15	18	16	**14**	*13*	12	12	12	12	12	12
	16	17	*15*	*14*	12	12	12	12	12	12	12
	17	16	*15*	*13*	12	12	12	12	12	12	12
	18	16	**14**	13	12	12	12	12	12	12	12
	19	15	*14*	12	12	12	12	12	12	12	12
	20	15	*13*	12	12	12	12	12	12	12	12
	21	**14**	13	12	12	12	12	12	12	12	12

Inventory levels marked in bold in each period represent the mean inventory levels from a 5000 trial simulation employing the optimal pricing policy. The italicized inventory levels represent one standard deviation above and below the mean.

Table A3. Optimal pricing policies for Condition INC for the parameter values shown in Table 1.

						Period					
		1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
	1	20	20	20	20	20	20	20	20	20	20
Inventory	**2**	20	20	20	20	20	20	20	20	20	20
	3	20	20	20	20	20	20	20	20	20	18
	4	20	20	20	20	20	20	20	20	20	*16*
	5	20	20	20	20	20	20	20	20	20	*14*
	6	20	20	20	20	20	20	20	20	18	**13**
	7	20	20	20	20	20	20	20	19	*16*	*13*
	8	20	20	20	20	20	20	20	18	*15*	*12*
	9	20	20	20	20	20	20	19	*17*	**14**	*12*
	10	20	20	20	20	20	19	18	*16*	*13*	12
	11	20	20	20	20	20	18	17	*15*	*13*	12
	12	20	20	20	20	18	17	*16*	**14**	*12*	12
	13	20	20	19	18	17	16	*15*	*13*	12	12
	14	19	19	18	17	17	*16*	**14**	*13*	12	12
	15	18	18	17	17	16	*15*	*14*	13	12	12
	16	17	17	17	16	*15*	**14**	*13*	12	12	12
	17	16	16	16	15	*15*	*14*	*13*	12	12	12
	18	16	16	15	*15*	**14**	*13*	13	12	12	12
	19	15	15	*15*	**14**	*14*	13	12	12	12	12
	20	15	*15*	**14**	*14*	*13*	13	12	12	12	12
	21	**14**	**14**	*14*	13	13	12	12	12	12	12

Inventory levels marked in bold in each period represent the mean inventory levels from a 5000 trial simulation employing the optimal pricing policy. The italicized inventory levels represent one standard deviation above and below the mean.

Online Appendix

The Pricing Game: Instructions to Subjects

Introduction

Thank you for volunteering to participate in today's "Pricing Game" experiment. We appreciate your time, and will compensate you for your efforts. A research foundation has contributed the funds to support this research. The amount of money you earn will depend on your performance. Understanding the task is the most important factor in determining your earnings.

This set of instructions first explains the task and then demonstrates the software that you will use to perform it. Please read the instructions carefully. At any point during the experiment, feel free to refer to the instructions or ask questions.

Overview

Your role in the "Pricing Game" is that of a manager who must decide how to price identical units of a certain seasonal good. As in the real world, the lower the price you set for the good, the greater the demand for the good. Conversely, the higher the price you set, the lower the demand. The revenue you receive for selling a unit of the good increases with the price you charge. Therefore, in this "Pricing Game", you face a trade-off between setting the price low and possibly selling many units at a low profit per unit, or setting it high and possibly selling fewer units at a higher profit per unit.

In the game, you have a fixed period of time—called a season—during which you can sell your units. Each season is divided into 10 periods. At the beginning of the selling season, you will have a fixed number of units to sell over the course of the season. The number of units that you have at the beginning of the season is referred to as your starting inventory. You will not be able to obtain any more units to sell once the season starts. For instance, if you run out of units in period 7 of the selling season, then you will have no units to sell in periods 8 through 10. However, if you terminate period 10 with some units still in inventory, then you will not have any additional chance to sell them. They will be worth nothing to you.

You will start each season with 21 units to sell. Each unit costs you $10 to purchase. Your total costs are equal to the number of units in your starting inventory multiplied by the unit cost of the good. Therefore, your starting costs in each season are $210. Your profits will depend on your total costs and on the revenue you accrue during the selling season. Specifically, your profits will equal to your total revenue minus your total costs. Exactly how your payments are calculated will be described below.

As you decide on a price, you can select a price and observe the demand for that period and subsequent periods. Demand depends on the price you set in a probabilistic fashion, as we explain below. The lower you set the price, the greater the expected demand; and conversely, the higher you set the price, the lower the expected demand. (see Appendix in the last page.)

You must select your selling price from a fixed and pre-determined set of 20 prices. You cannot set the price to any price not in the list. Demand for the good will remain constant each period.

You will sell goods over a total of 40 identical seasons.

Your Revenue

The number of units you sell in a period is equal to the number of units demanded or the number of units you have left in inventory, whichever is smaller. For instance, if you have 10 units in inventory and the demand is for 4 units, then you will sell 4 units. However, if you have 2 units in inventory and the demand is for 4 units, then you will only sell 2 units. You cannot sell more units than you have in your inventory. The revenue you earn in a particular period is equal to the price you charge times the number of units you sell.

Summary

Before getting into the technical details of the game, here is a brief summary of the basic structure of the game:

- You have 10 periods to sell 21 units of a certain seasonal good (your starting inventory).
- Each unit has an associated unit cost of $10, which is the amount it costs to obtain each unit. Therefore, your total costs are equal to the number of units in your starting inventory times the unit cost, namely, 21 X $10 = $210.
- At the start of each period, you must select your unit selling price for the good for that period. The prices you can charge are restricted to a fixed set of 20 prices determined by the software.
- Over the course of the selling season, you can change the prices you charge in each period. The price set in each period will be the price you receive for all units sold in that period.
- Demand in each period is a probabilistic function of the price you charge (see below). The demand for the good will be constant each period.
- The number of units you sell in a period is equal to the number of units demanded or the number of units you have in inventory, whichever is smaller.
- The revenue you earn in a particular period is equal to the price you charge times the number of units you sell.
- Your profit for a selling season will be equal to your total revenue minus your total costs of purchasing your starting inventory.
- Your payment today will be based on your profits. The higher your profits, the more money you will make. (Details regarding how payments are determined are given below.)

How the Demand is Determined [Demand Profile Screenshot]

It was stated above that the demand is a probabilistic function of the price you set. We will now be more specific and explain exactly what this statement means. In each period you can set your selling price to any value between $1 and $20 in multiples of $1. The price you set will determine the demand distribution for that period. The graph below (which is a screenshot from the software) shows the probabilities of demand for a period when the price is set at $17. The height of the bar at each demand value represents the probability of demand equaling that value given the price you set. For instance, the probability that the demand will equal 2 units if the price is $17 is about 23%; the probability that demand will equal 0 units is about 7%, etc.

The Pricing Game Software

In this section, we illustrate the game and lead you through the software used to implement it. Prior to each season, all of the relevant information for the season will be displayed on the following screen:

This screen shows:

- Periods in Season: The number of periods in which you can sell your units.
- Salvage Value: How much unsold goods are worth.
- Unit Cost: How much each unit cost you to purchase.
- Starting Inventory: How many units you begin the selling season with.
- Total Costs: Your total costs for purchasing your inventory (Starting Inventory X Unit Cost).

Once you are comfortable with the season parameters, you can click the ‘Start Selling Season’ button at the bottom of the screen. Doing so will take you to the main screen:

[Game interface screenshot]

On the screen above, you are in the first period of the selling season. You can tell this by looking in the Sales Period Information Panel and also in the Sales History table in the lower left corner of the screen. The current period will always be highlighted in the Sales History table.

The Sales Period Panel contains a lot of important information. In it, you can find your revenue for the current period, your total revenue for the current season, and also your current profit. Remember that profit is equal to your total revenue minus your total costs. At this point, you have not made any revenue and, therefore, your profit is -$210 (that is a $210 loss). Right now, the software is waiting for you to set your price for the first period.

The screen shows that you started with 21 units of inventory (Starting Inventory) and that you currently have 21 units (Inventory). The screen reminds you how much you paid for each unit (Unit Cost), how much unsold units are worth (Salvage Value), and your starting costs (Total Costs).
At this stage, you have to decide how to price your units for the first period. As shown, the price is now set to $10. Before deciding how much to price each unit, please examine the demand distributions for the permissible prices. You can do so by moving the slider below the demand distribution.

These screens show the demand at a price of $10 and $15. You will notice that the mean demand for $15 is less than for $10. Once you have decided on a price, you must click on the current period button to be able to post your price for the current period by clicking on the post price button.
[Demand profile screenshots]

For the sake of demonstration, suppose that you set the price in the first period to $12. To do so, you move the slider to the $12 price position, and press the Post Price button. After doing so, you learn that you have sold 2 units and that your revenue for period 1 is $24. This information can be found in the Sales Period Information panel and also in the Sales History Table. See the figure above. Also, your inventory has been reduced by 2 units. Thus, you will have 19 units in inventory when you begin period 2. You can continue to period 2 by pressing the Go To Next Selling Period button.

Skipping ahead a bit, suppose that you continue to set your prices in the remaining periods of the season. See screen below. Your total (cumulative) revenue in this example is $332, and your profit is $122 (total revenue minus total costs).

[Season 10 screenshot]

After the final period, you'll be presented with a summary of your earnings information. You can view this information by pressing the Examine Earnings button. After doing so, you will see a screen like the one below.

[Summary screenshot]

The interpretation of the entries should be obvious. The one that you are most interested in is the final row of the table: Profit. In this example, you profit is $122.

Once you have examined your earnings, you can click on the 'Continue to Next Season' button to proceed to the next season. After doing so, you will observe the parameter screen for the next season, and then continue to the selling season where you must set your selling prices again.

After you have complete playing all 40 seasons, you will be shown a summary of all of your earnings in the experiment and receive payment for your participation.

Calculation of Experimental Payments

You can make a considerable amount of money in this experiment. You will be paid the 20% of the average of your profits over the 40 selling seasons. Your payment will be given to you in cash at the conclusion of the experiment.

Some Advice for Increasing Earnings

To maximize your earnings, you should make sure that you have a good grasp of the Pricing Game. In particular, before you begin, please consider carefully the relationship between the price and demand. You should always keep in mind how many units you have in inventory and how many periods you have left in the season to sell them.

Please let the experimenter know if you have any questions.

Appendix

The next page exhibits the demand distributions for periods 1-10. To save space, we have only included the demand distributions for the two extreme costs, namely cost=$1, and cost=$20, and for the middle value cost=$10. Feel free to consult these diagrams at any stage during the experiment.