Decision Biases in Revenue Management: Some Behavioral Evidence

J. Neil Bearden
Departments of Management and Organizations and Systems and Industrial Engineering, University of Arizona, Tucson, Arizona, 85721

Ryan O. Murphy
Center for the Decision Sciences, Columbia University, New York, New York, 10027

Amnon Rapoport
Department of Management and Organizations, University of Arizona, Tucson, Arizona, 85721 and Department of Economics, Hong Kong University of Science and Technology, Hong Kong

jneilb@gmail.com • *rom2102@columbia.edu* • *amnon@u.arizona.edu*

June 2, 2006

We study a problem of selling a fixed number of perishable goods over a finite and known horizon. After presenting a procedure for computing optimal decision policies and some numerical results on a simple heuristic policy for the problem, we describe results from two experiments involving financially motivated subjects. The experiments reveal that decision makers employ policies of the same form of the optimal policy. However, they show systematic biases to demand too much when they have many units to sell and too little when they have few to sell, resulting in between 8 and 11% revenue losses. We also propose some new ways in which behavioral studies can complement theoretical work in pricing and revenue management.
(Behavioral Operations; Revenue Management; Dynamic Pricing; Knapsack Problems; Decision Bias; Heuristics)

---

1. Introduction

Firms often face the problem of deciding how to best price and control the inventory of perishable products for which demand is stochastic and price sensitive. Airlines must do so for seats on particular flights; hotels must do so for rooms on particular nights; and fashion retailers must do so for seasonal goods. Keeping prices too high may result in unsold items, and keeping them too low may have significant opportunity costs because costumers would have been willing to pay more. There has been considerable theoretical work in the operations literature on methods for optimally solving dynamic pricing and revenue management problems, but, as far as we know, there has been no direct experimental work on how well actual decision makers perform. Since managers—who are not necessarily perfectly rational decision makers nor extensively trained in optimization methods—are generally responsible for revenue management (RM) decisions in most firms, investigating how their decisions are biased may prove valuable. In the current paper, we use laboratory experiments to investigate decision behavior in a stylized RM problem that captures some of the important features of the problems faced by managers in the "real world."

Suppose a firm has a finite number of periods—a *season*—in which to sell a fixed number of units of a product. Bids to buy a unit of the product at a particular price arrive sequentially and stochastically in time. Each time it receives a bid, the firm must choose either to accept or reject it on the spot. When it accepts one, it sells a unit of the product to the bidder at the bid price. Otherwise, it irrevocably rejects the bid and must wait for another one, which may or may not arrive before the end of the season. There are a number of ways to interpret this general problem. One is to think of it as the one faced by airlines, hotels, and travel agencies that sell their goods on Priceline.com. Visitors to the site make offers to purchase goods (e.g., a single one-way ticket from Tucson to New York on July 5, 2006) at particular prices, and their offers are either accepted or rejected. The visitor's credit card is automatically charged the bid price if the bid is at least as high as the current reservation price for the good, which is determined by Priceline, and, of course, unknown to the visitor; otherwise, she pays nothing and receives nothing. Another interpretation is that goods in different product classes (e.g., fare classes) are priced by the seller and posted. Then, when a buyer attempts to buy a good in a particular class at the posted price, the seller decides whether to make a unit of the good available. Whatever interpretation one assigns, this general problem possesses the fundamental features of problems faced by firms in many industries, namely, a fixed stock of items, a finite selling horizon, and uncertain demand, and is the archetypal problem in RM. It is also the type of problem we study in the current paper.

There are a number of excellent reviews that can be consulted for introductions to and overviews of research in dynamic pricing and RM (e.g., Bitran and Caldentey, 2003; McGill and van Ryzin, 1999; Talluri and van Ryzin, 2004; Weatherford and Bodily, 1992). Therefore, we have chosen to focus on some of the prior experimental studies of decision behavior that are most relevant to the RM problem we consider here. These studies all share a common framework: They investigate behavior in sequential decision problems with known optimal decision policies. The current paper employs this same approach.

Optimal stopping is central to many operations management decision problems such as when to hire a job applicant and when to adopt a new technology, and is at the core of many RM problems (Brumell and McGill, 1993). The general theory of optimal stopping has received considerable theoretical (e.g., Chow, Robbins, and Siegmund, 1971; Gilbert and Mosteller, 1966) and experimental attention. Rapoport and Tversky (1970), for example, examined decision behavior in the classical full-information optimal stopping problem in which the decision maker (DM) sequentially observes up to N random draws from a distribution with a known density $f(x)$, must irrevocably either accept or reject each draw when it is observed, and receives a payoff equal to the value x of the single selected observation. Rapoport and Tversky found that DMs tended to stop sooner than was dictated by the optimal—expected payoff maximizing—policy. Similar findings regarding behavior in full-information optimal stopping problems were reported by Cox and Oaxaca (1989) and Schotter and Braunstein (1981).

Seale and Rapoport (1997) studied decision behavior in the classical secretary problem (CSP). In the CSP, the DM observes up to N alternatives, learns only the quality of each alternative relative to those previously observed and rejected, must irrevocably accept or reject each alternative when it is encountered, and earns a positive payoff if and only if she selects the single best alternative (among all N of them). This problem has received considerable attention in the applied probability literature primarily because of a special feature of its optimal policy, which is a threshold rule: Reject the first $n-1$ applicants, and thereafter select the next encountered *candidate*, that is, the next applicant who is better than all of those previously observed and rejected. What is notable is that as $N \to \infty$, $n \to e^{-1}N$; that is, for large N, the DM should skip the first roughly 37% of the applicants, and then choose the next encountered candidate. (For general reviews of theoretical work on secretary problems, consult Ferguson, 1989, and Freeman, 1983.) Seale and Rapoport reported evidence that subjects employed a threshold policy of the same form as the optimal policy but set their threshold n beneath the optimal threshold; hence, the subjects stopped searching too soon. Seale and Rapoport (2000) studied an extension of the CSP (due to Presman and Sonin, 1972) in which the DM does not know the exact value of N but does know the distribution from which N is drawn. For example, in one condition the subjects knew that N was drawn with uniform probability from the set of integers $\{80, 81, \ldots, 120\}$. Again, subjects used policies of the same form as the optimal policy but displayed a tendency to stop searching too soon.

One restrictive feature of the CSP is that the DM only earns a positive payoff for selecting the single best applicant in the pool. It is difficult to think of real-world decision problems that have this payoff structure. Bearden, Rapoport, and Murphy (in press a) studied a generalization of the CSP in which the DM earns payoffs that are increasing in the quality of the selected applicant (relative to all N applicants). They also found strong evidence that subjects tend to select alternatives earlier than is dictated by the appropriate optimal policy, which is a sequence of thresholds that progressively relax as the DM moves closer to the hiring deadline. However, they did find that the subjects employed policies of the same form as the optimal policy. Thus, though the subjects behaved suboptimally, their decision policies were rather sophisticated. Bearden, Murphy, and Rapoport (2005) reported similar findings in a multi-attribute extension of the same problem.

The conclusion drawn from this stream of experimental studies of optimal stopping problems is that people have a propensity to search inadequately, but employ rather complex decision policies that have the same structural form as the optimal policies. Assuming that the early stopping bias generalizes to situations beyond the laboratory, one might predict that people do not search enough before making purchase decisions, that they may not visit enough sites before purchasing books or airline tickets online, for example. Indeed, Johnson et al. (2004) show that online search is quite limited. For real-world search problems, such as online shopping, it would be very difficult, if not impossible, to determine what truly optimal search would require. This is one reason why laboratory studies are so powerful. We can confront actual DMs with decision problems for which we know the optimal policies. By comparing actual to optimal decision behavior, we may gain insights into where decision making is done well and where it breaks down.

Note that in all of the optimal stopping experiments that we have described the DM had a single unit to sell (e.g., a single position to fill). (There also has been some theoretical work on secretary problems with multiple selections in which the DM may select several of the sequentially encountered alternatives (e.g., Praeter, 1994), but there has been no experimental work on these problems.) Thus, it is unclear how these behavioral results on optimal stopping inform the more general RM problem in which the DM has multiple units of a perishable good to sell over some period of time. For example, do the findings suggest that DMs are likely to set their selling prices too low?

Bearden, Rapoport, and Murphy (in press b) experimentally studied a sequential assignment problem in which the DM has to assign N jobs to N machines with the objective of minimizing assignment costs. The DM's objective is to assign the best jobs to the best machines. The values of the machines are known up front. The problem is rendered difficult by the fact that the DM only learns the values of the jobs sequentially and relative to those previously seen and assigned, but her payoff for an assignment is a function of the overall value of the assigned job (relative to all N jobs). This problem is a secretary-problem-like extension of Derman et al.'s (1972) *Sequential Stochastic Assignment Problem*. Bearden et al.'s experimental results revealed that DMs had a bias to over-assign early jobs to intermediate-quality machines. The DMs' assignments in early periods then forced them to make overly costly assignments in later periods. In RM settings, this kind of error could manifest itself in a number of ways. A clothing retailer might overprice goods early in the season and then be forced to slash prices and have a "clearance sale" at the end of the season. Alternatively, the DM might sell goods for which there is great demand too cheaply early on and then forgo the opportunity to sell those goods at higher prices later.

Overall, given the dearth of experimental work on decision behavior in dynamic decision problems, it is difficult to derive specific predictions for how DMs will perform in RM problems such as the Priceline.com problem described above. We can, however, predict the following: Decision behavior will depart from optimality. Given the relative complexity of making optimal RM decisions, this prediction is obvious and not all that interesting by itself. However, by finding the *ways* in which decision behavior *systematically departs* from optimality, we can establish a basis for prescription. At the least, experimental work on these problems can be used to warn DMs about the broad ways in which their RM decisions are likely to err. This by itself, we believe, is valuable.

The rest of the paper is organized as follows. In § 2, we formally describe a simplified RM problem, and present a numerical procedure for computing its optimal decision policies. We then demonstrate in § 3 that a relatively simple decision heuristic can perform quite well in the problem. Next, in § 4, we present results from two behavioral experiments involving financially motivated subjects in which we examine actual decision behavior in the problem. Finally, § 5 contains a discussion of the experimental findings and presents several suggestions for future experimental research on pricing and revenue management.

2. The Problem and Its Solution

In this section, we describe a simplified RM problem and present a method for computing its optimal policy. The problem has appeared under various guises in the operations literature. Lee and Hersh (1993), for example, present it as a model of airline seat inventory control. Papastavrou, Rajagopalan, and Kleywegt (1996) describe it quite generally as a dynamic and stochastic knapsack problem, and relate it to transportation scheduling problems, taking reservations in restaurants, and airline booking. To provide ourselves with a useful shorthand, and with apologies for perhaps sounding too grandiose, we will refer to the problem as the *Revenue Management Problem* (RMP).

2.1 The *Revenue Management Problem*

A DM can sell up to S units over a *season* of T discrete time periods. Periods are indexed by t ($t = T, T-1, \ldots, 0$), which represents the number of *periods remaining* until the end of the season. Units cannot be sold after period $t = 1$, and the salvage value for a unit is 0. The number of available-to-be-sold units is indexed by s ($s = S, S-1, \ldots, 0$). The *state* of the system is denoted by (t, s). In each period, an offer to buy a single unit arrives with probability p, and no offers arrive with probability $1-p$. Each offer has an associated bid (revenue) r, which is a random variable taken from a distribution with density function $f(r)$. Whenever the DM has $s \geq 1$ units and receives an offer with bid r, she can sell a unit, thereby increasing her total revenue by r and leaving her with $s-1$ units, or she can reject the offer. The decision to either accept or a reject an offer cannot be delayed—it must be made on the spot. The DM's objective is to maximize her expected (total) revenue for the selling season.

2.2 Dynamic Programming Solution for the RMP

Based on Lee and Hersh (1993) and Papastavrou et al. (1996), we know that:

Theorem 1. *At stage t with s remaining units, the optimal policy is a threshold rule:*

$$\psi^\star(t,s,r) = \begin{cases} \text{accept offer if} & r \geq R_t^s, \\ \text{reject offer if} & r < R_t^s. \end{cases}$$

The *threshold* R_t^s dictates what revenue levels r the DM finds acceptable given t and s. These thresholds are computed from

$$R_t^s = \begin{cases} V_{t-1}^s - V_{t-1}^{s-1} & \text{if } s \geq 1, \\ \infty & \text{if } s < 1, \end{cases} \tag{1}$$

where

$$V_t^s = p\left[\int_0^{R_t^s} f(r) V_{t-1}^s \mathrm{d}r + \int_{R_t^s}^{\infty} f(r)\left(r + V_{t-1}^{s-1}\right)\mathrm{d}r\right] + (1-p)V_{t-1}^s, \tag{2}$$

with the boundary conditions

$$\begin{aligned} V_0^s &= 0, \ \forall s, \text{ and} \\ V_t^0 &= 0, \ \forall t. \end{aligned}$$

The value function V_t^s gives the DM's expected future revenue for following the threshold rule $\psi^\star(t,s,r)$ from period t to period 0, given that she has s units left. The optimal policy depends on t, s, and r, and not on the history prior to t. Thus, Bellman's (1957) optimality principle of dynamic programming is satisfied; and the optimal policy for the full problem from period T to period 0 can be computed by solving Eqs. 1 and 2 recursively from $t = 0$ to $t = T$. Under the optimal policy, when she receives an offer, the DM simply decides whether the expected marginal value for holding a unit for one more period exceeds the marginal revenue for selling it at the current bid value. If she is (expected to be) better off keeping the unit, she does so; otherwise, she sells it.

Observe that Eq. 2 can be expressed quite simply for uniformly distributed bids. If $r \sim \text{Uni}[a, b]$, it can be written as

$$V_t^s = p \left[\frac{R_t^s - a}{b - a} V_{t-1}^s + \frac{b - R_t^s}{b - a} \left(\frac{R_t^s + b}{2} + V_{t-1}^{s-1} \right) \right] + (1 - p) V_{t-1}^s. \tag{3}$$

We will utilize Eq. 3 in much of what follows.

From Papastavrou et al. (1996), we know that the optimal policy for the RMP has the following properties:

Theorem 2.

(i) R_t^s is nonincreasing in s for all t.

(ii) R_t^s is nondecreasing in t for all s.

Theorem 2 shows that the optimal DM is less choosy when she has more units to sell and when she has less time to sell them. The intuitions for this are clear. Since the DM is faced with a deadline beyond which she can no longer sell her units, she must be less demanding when she has a large number of them to sell; otherwise, since future demand is uncertain, she may end up with unsold units, which are worthless. The DM should become less demanding as her deadline approaches, since getting something for a unit is better than getting nothing. These properties can be discerned from the optimal threshold values exhibited in Fig. 1. The pattern of thresholds displayed in the figure also holds for other bid distributions such as the normal, exponential, and triangular; it is not peculiar to the uniform distribution.

Some numerical examples of the time course of optimal thresholds are displayed in Fig. 2. The height of the curve in each period represents the optimal threshold in that period. Each path is based on simulating the application of the optimal policy to offers generated according to the problem parameters. The spikes in the curves occur when a unit is sold (periods in which sales occur are represented by squares on the curves). At these points, since there are fewer units to sell, it behooves the DM to become more demanding. Consider the paths in the left-hand two plots when $t = 20$. For the one on the top in which four of the five units have been sold, an optimal DM is very demanding, requiring more to sell a unit than at $t = 40$. In contrast, in the one on the bottom, where the DM still has four units to sell, she is willing to virtually give them away. The important point to take from Fig. 2 is that the optimal thresholds are non-monotonic: they strongly depend on the outcomes of the random process that gives rise to the offers.

Theorems 1 and 2 provide straightforward testable predictions for our experiments. We describe below two experiments in which we examine the decision behavior of actual financially motivated DMs in the RMP. But, first, we present some numerical results on the performance of a simple, non-dynamic decision heuristic for the RMP.

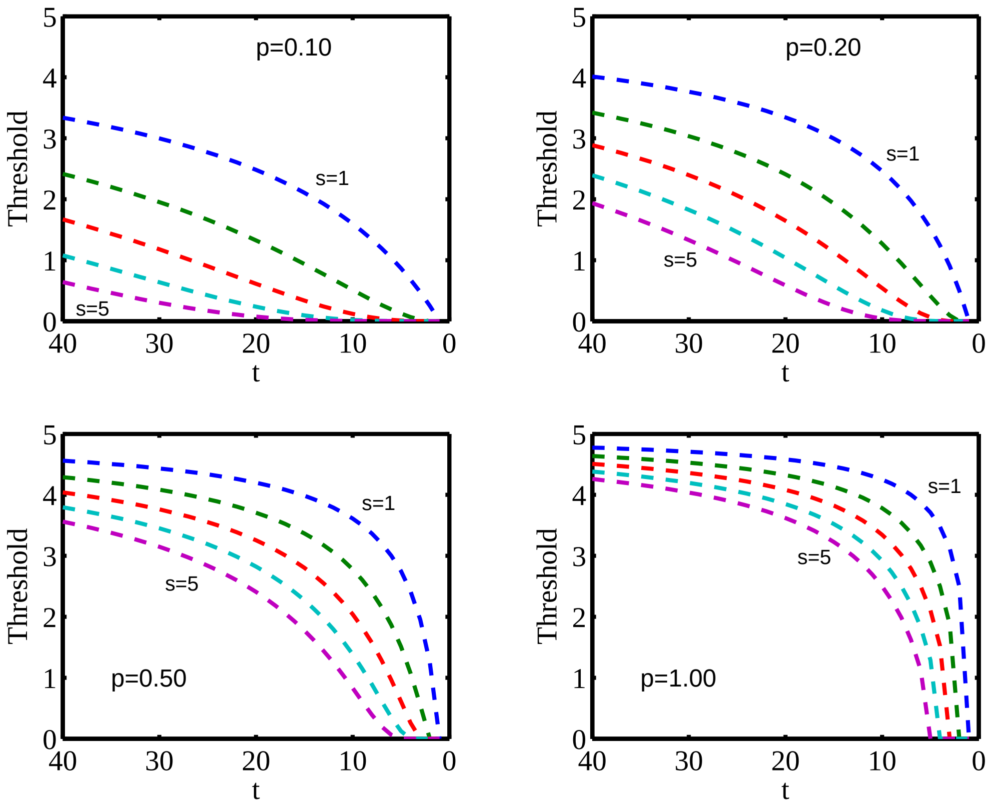

Figure 1: Optimal thresholds R_t^s for problems in which $T = 40$ and $r \sim \mathrm{Uni}[0.01, 5.00]$ for several arrival probabilities p.

3. A Simple Heuristic for the RMP

There has been considerable discussion of simple decision heuristics in the psychology literature (e.g., Gigerenzer et al., 1999). The now-relatively-conventional argument—which actually dates back to Simon (1955)—is that human DMs have limited computational capacity but are generally able to make good decisions using simple heuristics. Work along these lines often proceeds by demonstrating that simple decision heuristics *can* perform well, using Monte Carlo simulation to do so. Less frequently, researchers actually test whether people employ simple heuristics when making decisions (for some exceptions, see Bröder, 2000; Newell and Shanks, 2003; Schulte-Mecklenbeck, Willemsen, and Johnson, 2006).

We have shown that expected revenue maximization in the RMP requires the use of relatively sophisticated dynamic decision policies. Under these, the DM determines her acceptable revenue levels (i.e., her thresholds) in each period after taking into consideration both how many periods she has left to sell units and how many units she has to sell. But how might a DM fare if she employs a simple, *static* decision policy? One of the simplest heuristics a DM might employ is the *Fixed Threshold Policy*: Accept any offer for which $r > R$, where R is fixed for all t and s. If a DM wants to optimize the performance of this policy, where should she set her threshold? Further, how effective would such a policy be?

The value (expected revenue) of a fixed threshold policy V_H can be obtained by substi-

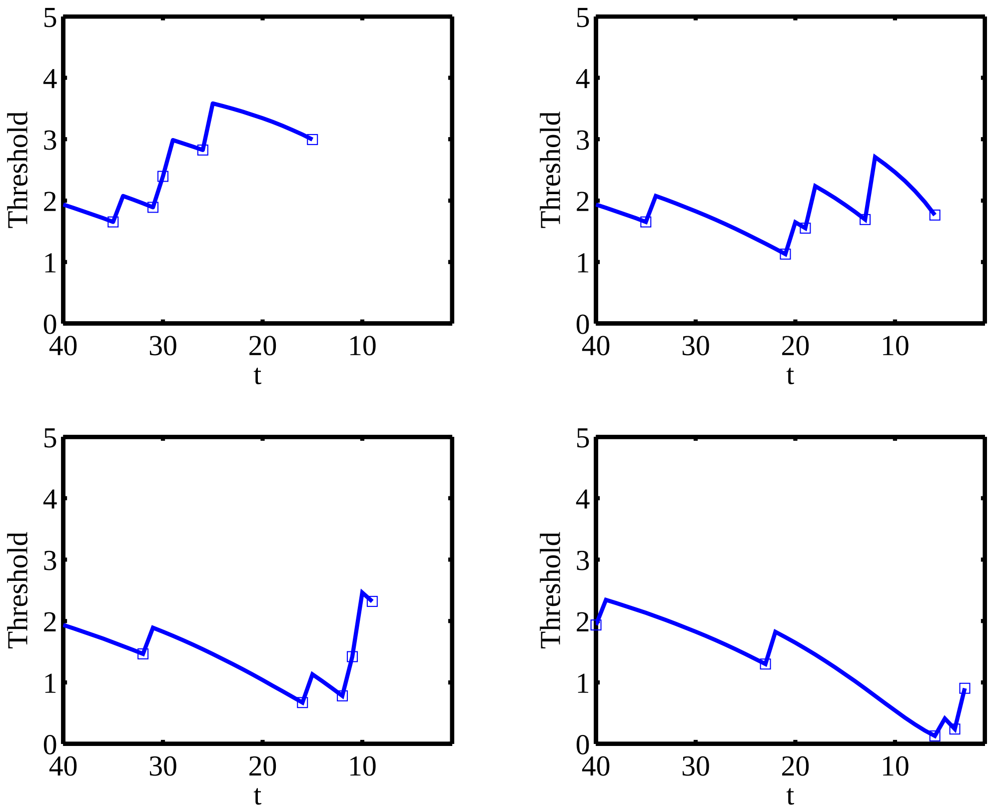

Figure 2: Some optimal threshold (price) paths for random instances of problems with $T = 40$, $S = 5$, $p = 0.20$, and $r \sim \text{Uni}[0.01, 5.00]$. The squares correspond periods during which a unit was sold. The curves terminate when all units have been sold.

tuting the single threshold R for each R_t^s in Eq. 2, performing the recursion from $t = 0$ to $t = T$, and setting $V_H = V_T^S$. The optimal heuristic threshold R_H is found by solving

$$R_H = \arg\max_R V_H, \tag{4}$$

which can be achieved using line-search methods. Some numerical results on the performance of the Fixed Threshold Policy are displayed in Table 1 for several special cases of the RMP. It turns out that very little is lost by using a simple fixed threshold heuristic in these problems. For each of them, a DM can expect to earn more than 97% of the optimal expected earnings by using a fixed threshold policy with the optimal threshold. We evaluated V_H/V_T^S for a large number of other combinations of T, S, and p, and always found that $V_H/V_T^S > 0.94$. As might be expected, our numerical experiments show that V_H/V_T^S tends to 1 as T grows.

Gallego and van Ryzin (1994) reported a similar result in a dynamic pricing problem. They showed that the expected revenue under a fixed price heuristic is always close to the optimal revenue, and, in fact, converges to the optimal revenue as the selling-horizon T grows. Based on their results, they concluded that when demand functions are well known and prices can be set freely, there is likely to be little benefit from highly dynamic pricing policies (analogous to using dynamic thresholds like those displayed in Fig. 2).

Next, we describe two behavioral experiments in which we examine actual decision be-

	R_H	V_H	V_T^S	V_H/V_T^S
Problem 1	2.42	17.05	17.66	0.97
Problem 2	2.29	9.65	10.01	0.97

Table 1: Optimal fixed thresholds R_H, expected earnings under the optimal fixed threshold policy V_H, expected earnings under the (full-blown) optimal policy V_T^S, and the efficiency of the fixed threshold policy V_H/V_T^S for some RMPs. The problems were parameterized as follows. Problem 1: $S = 5$, and $p = 0.30$. Problem 2: $S = 3$, and $p = 0.18$. Both had $T = 40$ and $r \sim \text{Uni}[0.01, 5.00]$. These are the same RMP parameters that we use in Experiments 1 and 2, respectively.

havior in the RMP.

4. Behavioral Studies of the RMP

4.1 Overview

We examined decision behavior in the RMP in two experiments. Both had the same general setup and used incentive-compatible payoffs to elicit careful decisions. The experiments differed from each other in the offer arrival rates p and the number of to-be-sold units S. Of course, one could manipulate any number of parameters of the RMP, including the length of the selling season T and bid distribution $f(r)$. Our main objective is to look for broad, replicable patterns of decision behavior in the RMP, and varying S and p allows us to cover a large region of the feasible problem space.

We fixed $T = 40$ and $r \sim \text{Uni}[0.01, 5.00]$ for both experiments. Experiment 1 had $S = 5$ and $p = 0.30$. Both the number of available units ($S = 3$) and the arrival rate ($p = 0.18$) were lowered in Experiment 2. To maintain some basis for comparison, we held the ratio of expected number of offers pT to available units constant across both experiments ($pT/S = 2.40$).

Thirty subjects participated in each experiment. All were recruited through flyers posted around the Columbia University School of Business to take part in a decision making experiment. The subjects were paid based on their performance in the experimental task and did not receive any course credit. Specifically, the subjects were paid their earnings from one randomly selected trial in Experiment 1 and from two randomly selected trials in Experiment 2. They earned an average of about \$15 for the 1 hr session.

Each subject was provided with extensive written instructions describing the task and the interface of the computer program that administered the experiment. The cover story for the task involved selling "contracts" to "bidders." The instructions described the RMP in non-technical language, and the values of the parameters of the problem (T, S, p, and $f(r)$) were all presented explicitly. To be clear, the subjects faced the RMP with perfect information about the problem parameters. Once the subjects were confident in their understanding of the task, they performed 50 (independent) trials of the RMP.

On each trial, the program automatically advanced through periods in which no offers arrived, pausing for 500 ms in each period. To emphasize that there was a selling deadline, the number of remaining periods was displayed textually (e.g., "Periods Remaining: 20")

	ℓ	$\alpha(\hat{R})$
Experiment 1	0.08 (0.02)	0.96 (0.02)
Experiment 2	0.11 (0.08)	0.96 (0.04)

Table 2: Average lost revenue (ℓ) and policy fits $\alpha(\hat{R})$ taken across subjects for Experiments 1 and 2. The values in parentheses are standard deviations.

and also by a graphical progress bar that shrank in each period. Whenever an offer arrived, the subject was shown the bid value and had to choose to either accept or reject it. No time restrictions were imposed in the accept-reject decisions. The computer program also continuously displayed the number of available contracts, the revenue from each sold contract, and the cumulative revenue (for the current trial). A trial terminated either when the deadline was reached or all contracts had been sold. The arrivals and offer values were generated randomly and independently for each subject by the experimental program according to the appropriate experimental parameters.

Given the similarity of the two experiments, and to conserve space, we will report the combined results.

4.2 Results

The most robust measure of decision behavior in the RMP is earnings (revenue). For each subject, we compared the subject's earnings π for each trial to the earnings $\pi^\star$ that *would have* resulted from the application of the optimal policy on that trial. The measure $\ell = 1 - \pi/\pi^\star$ quantifies the *revenue loss* resulting from the application of the subject's policy. The average lost revenue for each experiment is shown in Table 2. The losses are considerable and consistent across experiments. Real-world decisions resulting in $8-11\%$ revenue losses would certainly garner attention, especially in industries with narrow profit margins. Average losses did not change as a function of experimental trial (cf. Schweitzer and Cachon, 2000, who found no evidence of learning in the newsvendor problem).

For each accept-reject decision, we can determine whether it was optimal by using the optimal threshold R_t^s. We define a *miss* as a rejection of an offer r for which $r > R_t^s$. Similarly, a *false alarm* occurs when a subject accepts an offer of r for which $r < R_t^s$. Table 3 shows the median miss and false alarm rates (MR and FAR, respectively) for each value of s for each experiment.[1] We tested whether miss and false alarm rates differed within individuals across s using Wilcoxon (matched-pair) signed rank tests.[2] For each value of s, these analyses allow us to test whether the median MR and FARs differ. The z-statistics z(MR-FAR) for the tests are shown in Table 3. Note what when z(MR-FAR) < 0, false alarm rates exceed miss rates. The results show that the rates differ significantly when the subjects held a large number of units ($s = 5$, Experiment 1), and when they held a small

[1] Keep in mind that the threshold R_t^s for a given state (t, s) is optimal independent of whether the decisions made prior to t were made optimally. This is just a statement of the principle of optimality. A miss in one period cannot be compensated for by a false alarm later on, for example.

[2] The distributions of error rates are skewed due to the bounded [0,1] interval, and the assumptions of standard GLM parametric tests are not met. Therefore, we shall use non-parametric tests to evaluate the error rates.

	$s=1$	$s=2$	$s=3$	$s=4$	$s=5$
			Experiment 1		
MR	0.04	0.06	0.08	0.09	0.10
FAR	0.15	0.14	0.11	0.07	0.04
z(MR-FAR)	**-3.12**	**-2.93**	-1.00	0.28	**2.47**
ρ(MR,Earn)	0.17	0.33	0.15	-0.18	-0.23
ρ(FAR,Earn)	**-0.44**	**-0.54**	**-0.53**	**-0.59**	**-0.37**
			Experiment 2		
MR	0.05	0.06	0.09		
FAR	0.19	0.18	0.10		
z(MR-FAR)	**-3.90**	**-3.19**	-0.79		
ρ(MR,Earn)	0.02	0.00	0.21		
ρ(FAR,Earn)	**-0.50**	**-.047**	**-0.52**		

Table 3: Median miss and false rates for each value of s for Experiments 1 and 2. Wilcoxon sign-rank tests were used to compare MRs and FARs for each value of s; the z(MR-FAR)s are the z-values for the resulting test statistics. When $z(\text{MR-FAR}) < 0$ ($z(\text{MR-FAR}) > 0$) the FARs are greater (less) than MRs. Spearman (rank-order) correlations between error rates and average earnings (Earn) are denoted by ρ. Test statistics for the entries in bold are significant at the 0.05 level.

number of units ($s < 3$, Experiments 1 and 2). Specifically, when they had many units, they tended to be more demanding with respect to the optimal thresholds (since $z > 0$). In contrast, when they had only a few units to sell, they tended to be less demanding ($z < 0$).

Within each error type, we tested whether the changes in rates across s were significant using pair-wise comparisons. To compare the MR for i remaining units MR_i to the MR for j remaining units MR_j, we used Wilcoxon (matched-pair) signed rank tests. Likewise for the FARs. Table 4 shows the z-values for each pair-wise test. The general pattern is clear: Miss rates tend to increase in s and false alarm rates tend to decrease in s. The subjects tended to become more over-demanding (more under-demanding) as they held more (fewer) units.

To determine where departures from the optimal policy tended to have the greatest effects on revenue losses, using individual subjects as the unit of analysis, we correlated the MR and FAR for each value of s with average earnings. The results are striking and quite revealing. In both experiments, there is no significant (rank order) correlation between miss rates and average earnings ρ(MR,Earn) for any value of s. In contrast, for all values of s, there is a significant negative correlation between FARs and average earnings ρ(FAR,Earn). In words, those subjects who had more FARs (i.e., those who were least demanding) tended to earn less, whereas there was no clear relationship between miss rates and earnings.

To summarize, we find that the subjects displayed a general tendency to be too demanding when they held a large number of units, but insufficiently demanding when they held only a small number of them. Further, those subjects who tended to be more demanding for their units tended to earn more.[3] These findings, however, do not speak to the form of the

[3]This is not tautological: If a DM is highly over-demanding—that is, if her thresholds for accepting offers are too high—then she can end up with many unsold contracts, which thereby diminishes her revenue.

	Experiment 1				
	MR_1	MR_2	MR_3	MR_4	MR_5
MR_1	–	-1.03	**-2.51**	**-2.45**	**-3.29**
MR_2		–	**2.05**	-1.94	**-3.33**
MR_3			–	-0.48	**-2.40**
MR_4				–	**-2.22**
	FAR_1	FAR_2	FAR_3	FAR_4	FAR_5
FAR_1	–	0.73	**3.05**	**4.59**	**5.02**
FAR_2		–	**2.78**	**4.09**	**5.07**
FAR_3			–	**2.73**	**4.58**
FAR_4				–	**3.26**
	Experiment 2				
	MR_1	MR_2	MR_3		
MR_1	–	-1.42	**-2.43**		
MR_2		–	-1.41		
	FAR_1	FAR_2	FAR_3		
FAR_1	–	1.58	**3.80**		
FAR_2		–	**3.14**		

Table 4: Comparisons of MRs and FARs across different values of s. The table entries are z statistics from Wilcoxon signed rank tests comparing the row rate to the column rate. When $z > 0$ ($z < 0$), the row rate is greater than (less than) the column rate. Test statistics for the entries in bold are significant at the 0.05 level.

decision policies that the subjects actually employed in order to make their selling decisions. We examine this issue in the next section.

4.3 Estimating Decision Policies

Some questions we wish to address are: What kind of decision policies do subjects employ in the RMP? Do they use sophisticated policies or do they use simple heuristics? Based on the results of earlier experiments on sequential search, it is reasonable to assume that the subjects employ some kind of threshold when making decisions. In what follows, we will assume that each subject employs a threshold decision policy of the form:

$$\hat{\psi}(t, s, r) = \begin{cases} \text{accept offer if} & r \geq \hat{R}_t^s, \\ \text{reject offer if} & r < \hat{R}_t^s. \end{cases}$$

where $\hat{R}_t^s$ is the subject's (empirical) threshold for state (t, s). Given this assumption, we wish to determine the thresholds that best predict the subjects' decision data. *A priori*, some feasible threshold-setting heuristics are:

Sophisticated Threshold Policy. $\hat{R}_t^s = \hat{R}_t^s$.

Fixed Threshold Policy. $\hat{R}_t^s = \hat{R}$, $\forall t, s$.

Time Insensitive Threshold Policy. $\hat{R}_t^s = \hat{R}^s$, $\forall t$.

Units Remaining Insensitive Threshold Policy. $\hat{R}_t^s = \hat{R}_t$, $\forall s$.

The *Sophisticated Threshold Policy* permits the DM to adjust her threshold as a function of both t and s. The optimal policy is a special case (parameterization) of this policy. The *Fixed Threshold Policy* is the simplest reasonable policy that a DM could employ. Under it, she decides on a target marginal revenue and only accepts offers whose revenues exceed her minimum requirement regardless of how many units she has to sell and how much time she has remaining to sell them. As we showed above in § 3, this policy is not as foolish as it might appear: A DM who employs it optimally can do quite well. Another possible decision heuristic that the DM might employ is to set her threshold only on the basis of how many units she has left to sell, which we label the *Time Insensitive Threshold Policy*. Finally, deciding acceptable offer values solely on the basis of the number of periods remaining to sell units is captured by the *Units Remaining Insensitive Threshold Policy*.

The last three policies are special cases of the Sophisticated Threshold Policy (STP). Therefore, evidential support for any of them will provide support for the STP. On the other hand, in principle, the simpler policies can each be rejected based on the empirical data. Our approach to evaluating the relative success of these policies in accounting for the data is based on elimination. For each policy, we are looking for reasons to reject the hypothesis that subjects used it. Put differently, we cannot show inductively that a particular policy is the correct one, but we can show that a particular policy is an incorrect one. This problem in evaluating models of decision making in dynamic decision problems was discussed in Bearden and Rapoport (2005).

We can evaluate the subjects' actual decision policies by estimating their thresholds from their decision data. To do so, we find for each subject the set of thresholds $\hat{\mathbf{R}} = \{\hat{R}_t^s\}$ that maximize the *proportion of correctly predicted decisions* $\alpha(\hat{\mathbf{R}})$. To be clear, $\alpha(\hat{\mathbf{R}})$ measures both correctly predicted accept and reject decisions. Formally, we solve

$$\begin{array}{rll} \max_{\hat{\mathbf{R}},\Delta} & \alpha(\hat{\mathbf{R}}) & \\ \text{s.t.} & \hat{R}_{t-1}^s \leq \hat{R}_t^s & \\ & \hat{R}_t^s = \hat{R}_{t'}^s & \forall t, t' \in \{(j-1)\Delta + 1, \ldots, j\Delta\}, j = 1, \ldots, T/\Delta \end{array} \tag{5}$$

The constraints allow us to aggregate decisions made in different periods in order to estimate thresholds for a particular period. This seems reasonable: Subjects probably employ roughly the same threshold in periods 21 and 20, for example. If they adjust their thresholds at all, they presumably do so by becoming less choosy as they approach the deadline. The latter assumption was verified by the (raw) data. We can adjust the coarseness of the aggregation by manipulating Δ. In short, by imposing these structural constraints on the estimated policies we are implicitly using neighboring states to estimate thresholds that the subject never had an opportunity to employ.

The average estimated thresholds for each experiment are exhibited in Fig. 3. Note that the curves do not span the entire range of t. This is because some (t, s) states were either never encountered (e.g., holding 5 units when $t = 10$) or encountered very infrequently (e.g., holding 1 unit when $t = 30$ in Experiment 1); so estimating thresholds for these states was either impossible or likely to be overly sensitive to error (i.e., to response variability). The

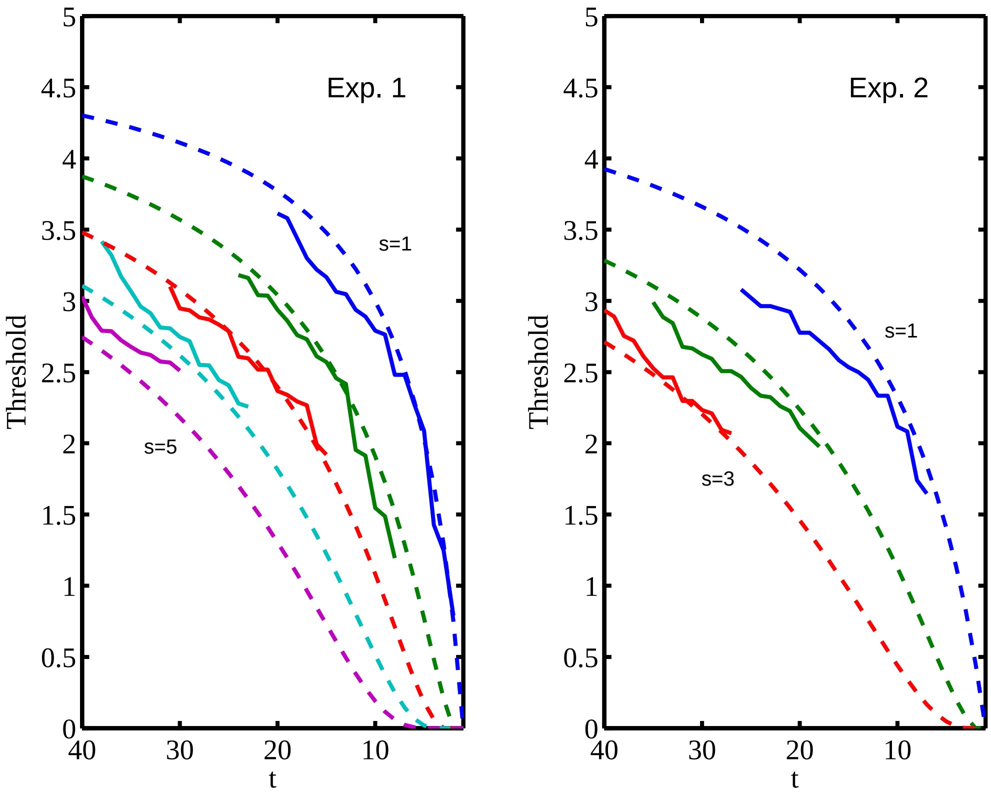

Figure 3: Optimal (dotted lines) and mean estimated (solid lines) thresholds for Experiments 1 (left) and 2 (right). For each value of s, estimated thresholds are only shown for values of t at which at least 2% of the offers for that s were encountered.

average policy fits $\alpha(\hat{\mathbf{R}})$ are presented in Table 2. The subjects' data are accounted for quite well by the threshold rule; on average, the policy predicts more than 96% of the subjects' decisions.

Based on the curves in Fig. 3, it seems that the subjects did not employ any of the three non-sophisticated policies because: 1) the curves are increasing in t, and 2) the curves are decreasing in s. Only the STP simultaneously permits both of these properties. It is important to emphasize that we did not constrain $\hat{R}_t^s \leq \hat{R}_t^{s-1}$. Thus, the analyses were not biased in favor of the sophisticated decision policy. However, we should also point out that the estimated policies in Fig. 3 are based on averaging over subjects. That is, they represent the "average subject" and not necessarily any particular one. We also examined the estimated policies for the individual subjects, and found support for the STP at the this level of analysis. Specifically, for most subjects we find downward sloping thresholds (as t decreases) that tend to decrease with more units (as s increases). We conclude that the Sophisticated Threshold Policy best accounts for the decision data from the RMP, and that the subjects' policies are in line with the qualitative predictions from Theorem 2.

Although they are of the same form as the optimal policy, the subjects' thresholds systematically depart from the optimal thresholds in several ways. First, when the number of remaining units s is large, the subjects tend to set their thresholds too high: they are too demanding. This is consistent with the high miss rates we observed for $s = 5$ in Experiment

1. On the other hand, when the subjects have only a small number of units remaining to sell, they tend to set their thresholds too low: They should be more demanding. This finding is consistent with the high false alarm rates for small s we observed in Experiments 1 and 2.

On average, the subjects could have earned more than they actually did by employing a simple fixed-threshold heuristic (see § 3). For the RMPs studied in the experiments, the fixed-threshold heuristic earns around 97% of the optimal, whereas the subjects earned (on average) around 92% and 89% of the optimal earnings in Experiments 1 and 2, respectively. Taken together, our results show that the subjects used sophisticated decision policies that were suboptimally parameterized, and that they could have done better by using a simple heuristic (if appropriately parameterized).

One prominent convention in behavioral decision research is to explain departures from optimality by appealing to the use of simple heuristics (e.g., Gigerenzer, et al., 1999). The arguments usually proceed by suggesting that DMs do not behave optimally because they use simple, easy-to-implement heuristics, which—though not optimal—tend to perform quite well. Importantly, the simple heuristics that are invoked are often structurally distinct from the optimal policy (i.e., they are of a different form). In contrast, the results from the studies of the RMP reported here (and also those from Bearden et al., in press a, among others) show that subjects behave suboptimally but employ rather sophisticated decision policies that are of the same structural form as the optimal policies. Therefore, explanations for suboptimal decision behavior need not invoke the use of simple heuristics. The question should not always be: What heuristics do people use that cause them to perform suboptimally? A clever researcher can always answer this question by a judicious choice of a heuristic or making up a new one. Instead, one might think that we should ask why the decision policies employed by experimental subjects are parameterized suboptimally. But since *a priori* we should not expect these policies to be set optimally—since, in the case of the RMP, the assumption further implies that we should expect the subjects to solve the dynamic program "in their heads"—the latter question also seems inappropriate.

Although we do not expect behavior to necessarily agree with the dictates of optimal models, finding where actual decision behavior departs from optimality—that is, where it is *biased*—is still valuable. Obviously there are *factors* that cause decisions to be biased, and searching for them is certainly an appropriate epistemic aim for experimental research on decision behavior in operations management problems. By uncovering these factors, we place ourselves in a better position to improve decision making: We can warn managers to be on guard against them. There seems to be no single "smoking gun" factor that drives the pattern of biases we observe in our studies of the RMP. Risk aversion, for example, cannot account for the subjects' tendency to be too demanding when they hold a large number of units, as it entails that they have thresholds that are lower than the risk neutral ones (Feng and Xiao, 1999). The bias to be insufficiently demanding when s is small, which is correlated with being near the deadline, may be driven by a fear of holding unsold units after the deadline. This would be consistent with the desire to minimize ex-post inventory error that Scweitzer and Cachon (2000) offered as an explanation for suboptimal behavior in the newsvendor problem. That is, near the deadline the subjects may have been highly averse to ending the season with unsold units, which then caused them just to "take what they could get." The RMP is a complex decision problem and a complete explanation of the (unobservable) factors that drive decision behavior in it will prove to be complex. Nonetheless, the observable behavioral biases are unambiguous and easy to describe.

5. Discussion and Future Directions

Generally, it would be difficult to assess the quality of RM decisions in natural environments. To determine optimal policies in these environments, one must make strong assumptions, and whether these assumptions are (precisely) met would be difficult to determine. For instance, pricing models often require that the DM knows the demand density function for all feasible prices. The optimal policies are not based on the DM having a "rough sense" of these functions or "good intuitions" about them; rather, these models assume that the DM *knows* the densities with precision. Clearly, conditions such as this are unlikely to be met in most of the real-world scenarios faced by managers. This fact illustrates one of the reasons why experimental studies are so useful. We can place DMs in contexts in which they do have all of the information that is assumed by the models, which, in turn, allows us to (more) legitimately compare decision behavior to the predictions of the appropriate optimal policies. By examining the ways in which laboratory behavior departs from optimality, we can establish some basis for making predictions about how DMs are likely to err in the wild.

Heching, Gallego, and van Ryzin (2002) compared the actual pricing policies of a women's apparel retailer to several model-based pricing schemes. Each of the models they examined required certain assumptions (e.g., knowledge of the demand function), which are unlikely to be perfectly met in reality, in order to derive pricing policies. Nonetheless, accepting these limitations, based on analyses of the company's historical pricing and sales data, Heching et al. concluded that the company's markdown prices were generally lower than those suggested by the models. They also concluded that the company would have increased its revenue significantly by employing smaller price markdowns earlier in the sales season rather than their actual practice of implementing steep markdowns late in the season. Our results on behavior in the RMP are consistent with these findings. In particular, we find that the largest driver of revenue losses in the RMP was subjects' tendencies to be insufficiently demanding when they held only a small number of units, which was correlated with nearing the end of the selling season. It is as if the subjects in the experiment employed steep markdown policies, and lost revenue for doing so.

The biases that we have found are compatible with those reported in Heching et al. (2002). Of course, this alone does not establish that our results generalize to RM decision making by actual managers. As we stated earlier, one way to gain confidence in the generality of biases observed in laboratory studies is to demonstrate that those biases occur in a range of problems. Below, we propose a few directions in which experimental studies of RM may prove useful. Ideally, this work will eventually be complemented by actual RM decision data from outside the laboratory.

Each time an offer arrives in the RMP, it is only for a single unit. Papastavrou et al. (1996) proposed a problem—of which the RMP is a special case—in which offers can arrive in batches of $k \geq 1$ units. The DM must, therefore, carefully consider both the capacity that will be consumed by an offer and the marginal revenue that she receives for accepting it. Airlines face this problem when selling blocks of seats. As another example, consider a transportation firm that has a truck scheduled to depart at a particular time.

		$c = 1$	
	$T = 10$	$T = 40$	$T = 100$
$p = .10$	2	5	17
$p = .30$	4	10	33
$p = 1.00$	9	25	81
		$c = 2$	
	$T = 10$	$T = 40$	$T = 100$
$p = .10$	1	3	12
$p = .30$	2	7	24
$p = 1.00$	6	18	60

Table 5: Optimal inventory levels $S^\star$ for several combinations of purchase cost c, selling horizon T, and arrival probability p when $r \sim \text{Uni}[0.01, 5.00]$.

Offers to transport loads that vary in size arrive stochastically. To maximize revenue, the firm must optimally set transportation fees for those loads. We intend to experimentally study this problem in the future. Doing so can help establish the generality of our results from the RMP.

In our statement of the RMP, we assumed that the bid functions $f(r)$ and arrival probabilities p were the same for all t. But there are many real-world scenarios in which demand is not constant across time. Demand for airline tickets, for instance, often increases as the departure date nears. In addition, since late-bookers are often business travellers, the buyers' reservation price distribution may change in time (Curry, 1990). Van Slyke and Young (2000) proposed a stochastic knapsack problem in which demand can be a function of time, and presented methods for computing its optimal policies. This problem would be a natural starting point for laboratory studies aimed at examining the effects of time-inhomogeneous demand on inventory control decisions.

One can easily add a stocking dimension to the RMP to bring it a bit closer to the full-blown RM problems faced by managers. Suppose that prior to the selling season the DM must decide how many units S to purchase, and that each unit costs c. Given T, p, and $f(r)$, what is the optimal number of units $S^\star$ for the DM to buy at the beginning of the season? Formally, the problem is to solve:

$$S^\star = \arg\max_S \left(V_T^S - Sc \right). \tag{6}$$

Requiring the DM to solve Eq. 6 prior to the selling season adds newsvendor element to the RMP. (Gallego and van Ryzin ,1994, presented results on a related, continuous-time pricing problem.) Solving Eq. 6 can be accomplished with little additional computational effort. Table 5 shows optimal inventory levels for several special cases of this Newsvendor-RMP. Consistent with natural expectations, the DM should start with more units when the units are cheaper (as c decreases), the selling season is longer (as T increases), and when she can expect to receive more offers (as p increases). It would be interesting to see if the newsvendor biases observed in Scweitzer and Cachon (2000) would obtain in the Newsvendor-RMP. Might DMs, for example, order too few units when c is low and too many when c is high?

So far, we only have discussed problems for which the arrival rate for offers is determined exogenously. Quite often, the DM can affect arrival rates by adjusting selling prices. Generally, demand for a good increases when prices decrease. Dynamic pricing problems, where the DM gets to set prices and influence demand, are another potentially fruitful area for experimental research. Consider the following pricing version of the RMP, which we have adapted from Bitran and Mondschein (1997). The DM has T periods to sell up to S units of a good. In each period t, the DM can set her unit price q for the good from a finite set of prices $\mathbf{q} = \{q\}$. The price she chooses determines the (stochastic) demand in that period. Let $\lambda(q) = \gamma e^{-\eta q}$, where $\gamma > 0$ and $\eta > 0$, be the demand rate function for each period. The actual demand d in a given period with price q follows a Poisson distribution with mean $\lambda(q)$; and the number of buyers is given by $\min(s, d)$. The value function for the Pricing-RMP can be expressed as:

$$V_t^s = \max_{q \in \mathbf{q}} \left\{ \sum_{d=0}^{s} \frac{e^{-\lambda(q)} \lambda(q)^d}{d!} \left(\min(s,d) q + V_{t-1}^{s-\min(s,d)} \right) \right\}, \tag{7}$$

with boundary conditions

$$\begin{aligned} V_0^s &= 0, \;\; \forall s, \text{ and} \\ V_t^0 &= 0, \;\; \forall t. \end{aligned}$$

The dynamic program is solved backwards in time to find the optimal price $q^\star$ for each state (t, s).

Gallego and van Ryzin (1994) showed that the optimal pricing policy for their continuous-time dynamic pricing problem, where prices can be chosen from an interval, had two important structural properties. First, the optimal price decreases in the number of units left in inventory. Second, for any given inventory level, the optimal price decreases as the end of the selling season approaches. Chatwin (2000) showed that the same results hold for a continuous-time pricing problem where prices are constrained to a fixed number of values, as they are in the Pricing-RMP. These general properties can be discerned from Fig. 4, which shows the optimal pricing policy for an example of the Pricing-RMP that constrains the permissible prices to a discrete set $q \in \{0.50, 0.75, 1.00\}$.

Bitran and Mondschein (1997) proposed a special case of the Pricing-RMP in which the price at each period q_t is constrained to be nondecreasing in t, reflecting some retailers' (e.g., clothing retailers) reluctance to increase prices for a good during a selling season. Zhao and Zheng (2000) present results on a related (continuous-time) pricing problem in which demand is time-inhomogeneous. Some important questions remain: How well do actual DMs solve dynamic pricing problems? Do they tend to set prices too high or too low? How well are their pricing policies adapted to time-inhomogeneous demand? A number of other pricing problems that may be suitable for laboratory investigation can be found in Talluri and van Ryzin (2004).

Behavioral experiments designed to investigate decision biases in revenue management problems are relatively inexpensive. If the findings from these studies can help actual managers avoid decision biases even by a small amount, the increased revenue for firms could be significant. The benefits of this research can be considerable relative to the costs. Simply warning managers that they may be biased to cut prices too much for seasonal goods, for example, could be quite valuable to them. And doing so may be more useful than an explanation of complicated dynamic pricing algorithms, since they are simply more likely to understand the former than the latter.

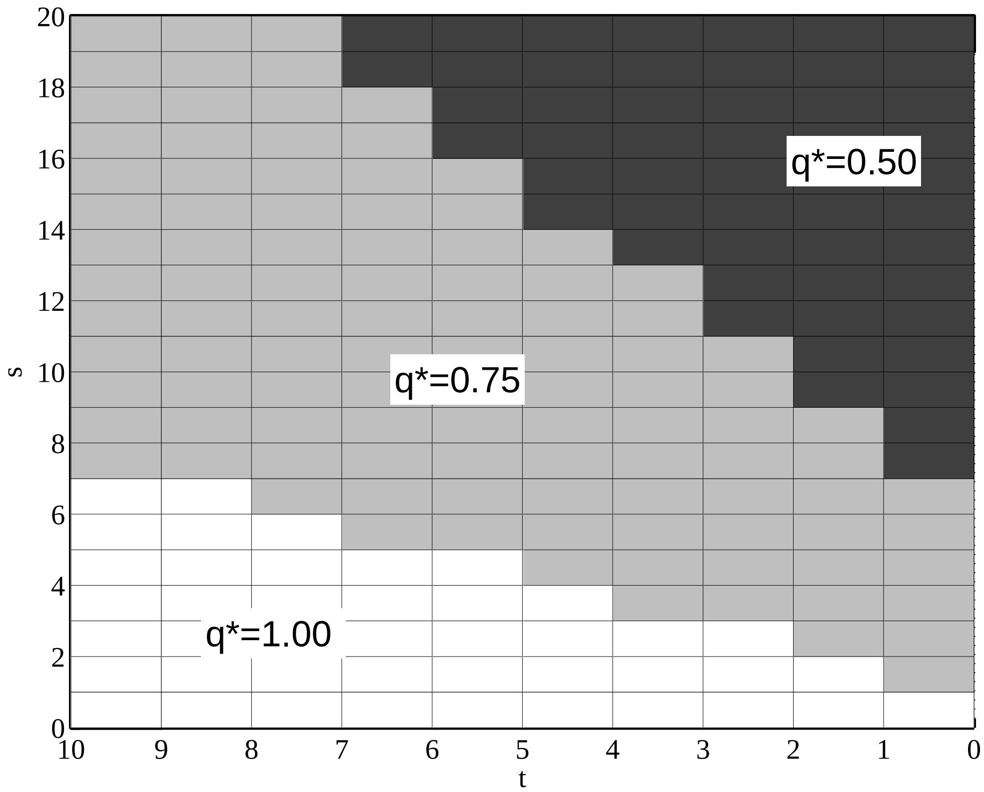

Figure 4: Optimal prices in the time-inventory plane for the Pricing-RMP with $T = 10$, $S = 20$, $\mathbf{q} = \{0.50, 0.75, 1.00\}$, $\gamma = 5$, and $\eta = 2.25$. The (t, s) states in each shaded region all have the same optimal price.

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., A. Rapoport. 2005. Operations research in experimental psychology, J. C. Smith, ed., *Tutorials in Operations Research: Emerging Theory, Methods, and Applications*. INFORMS: Hanover, MD, 213-236.

Bearden, J.N., R.O. Murphy, A. Rapoport. 2005. A multi-attribute extension of the secretary problem: Theory and experiments. *Journal of Mathematical Psychology* **49** 410–425

Bearden, J.N., A. Rapoport, R.O. Murphy. In press a. Sequential observation and selection with rank-dependent payoffs: An experimental test. *Management Science*.

Bearden, J.N., A. Rapoport, R.O. Murphy. In press b. Experimental studies of sequential selection and assignment with relative ranks. *Journal of Behavioral Decision Making*.

Bellman, R.E. 1957. *Dynamic Programming.* Princeton University Press: Princeton, NJ.

Bitran G., R. Caldentey. 2003. An overview of pricing models and revenue management. *Manufacturing & Service Operations Management* **5** 203–229.

Bitran, G.R., S.V. Mondschein. 1997. Periodic pricing of seasonal products in retailing. *Management Science* **43** 64–78.

Brumelle, S.L. McGill. 1993. Airline seat allocation with multiple nested fare classes. *Operations Research* **41** 127–137.

Bröder, A. 2000. Assessing the empirical validity of the "Take-The-Best" heuristic as a model of human probabilistic inference *Journal of Experimental Psychology: Learning, Memory, and Cognition* **26** 1332–1346.

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

Chow, Y.S., H. Robbins, D. Siegmund. 1971 *Great Expectations: The Theory of Optional Stopping.* Houghton & Mifflin: Boston, MA.

Cox, J.C., R.L. Oaxaca. 1989. Laboratory experiments with a finite-horizon job-search model. *Journal of Risk and Uncertainty* **2** 301–330.

Curry, R.E. 1990. Optimal airline seat allocation with fare classes nested by origins and destinations *Transportation Science* **24** 193–204.

Derman, C., G.J. Lieberman, S.M. Ross. 1972. A sequential stochastic assignment problem. *Management Science* **18** 349–355.

Feng, Y., B. Xiao. 1999. Maximizing revenues of perishable assets with a risk factor. *Operations Research* **47** 337–341.

Ferguson, T.S. 1989. Who solved the secretary problem. *Statistical Science* **4** 282–296.

Freeman, P.R. 1983. The secretary problem and its extensions: A review. *International Statistical Review* **51** 189–206.

Gallego, G., G.J. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science* **40** 999–1020.

Gigerenzer, G., P. Todd, The ABC Research Group 1999, *Simple Heuristics That Make Us Smart*, Oxford University Press, Oxford, New York.

Gilbert, J., F. Mosteller. 1966. Recognizing the maximum of a sequence. *Journal of the American Statistical Association* **61** 35–73.

Heching, A., G. Gallego, G.J. van Ryzin. 2002. Mark-down pricing: An empirical analysis of policies and revenue potential at one apparel retailer. *Journal of Revenue and Pricing Management* **1** 139–160.

Johnson, E., W. Moe, P. Fader, S. Bellman, G. Lohse. 2004. On the depth and dynamics of online search behavior. *Management Science* **50** 299–308.

Lee, T. C., M. Hersh. 1993. A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science* **27** 252–265.

McGill, J.I., G.J. van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Science* **33** 233–256.

Newell, B.R., D.R. Shanks. 2003, Take the best or look at the rest? Factors influencing "one-reason" decision making, *Journal of Experimental Psychology: Learning, Memory, and Cognition* **29** 53-65.

Papastavrou, J.D., S. Rajagopalan, A.J. Kleywegt. 1996. The dynamic and stochastic knapsack problem with deadlines. *Management Science* **42** 155–172.

Praeter, J. 1994. On multiple choice secretary problems. *Mathematics of Operations Research* **19** 597–602.

Presman, E.L., I.M. Sonin. 1972. The best choice problem for a random number of objects. *Theory of Probability and Its Applications* **17** 657–668.

Rapoport, A., A. Tversky. 1970. Choice behavior in an optimal stopping task. *Organizational Behavior and Human Performance* **5** 105–120.

Schotter, A., Y.M. Braunstein. 1981. Economic search: An experimental study. *Economic Inquiry* **19** 1–25.

Schulte-Mecklenbeck, M., M. Willemsen., E. Johnson. 2006. Process models and process data. Unpublished manuscript, Columbia University, New York.

Schweitzer, M. E., G.P. Cachon. 2000. Decision bias in the newsvendor problem with a known demand distribution. *Management Science* **46** 404–420.

Seale, D., A. Rapoport. 1997. Sequential decision making with relative ranks: An experimental investigation of the secretary problem. *Organizational Behavior & Human Decision Processes* **69** 221–236.

Seale, D., A. Rapoport. 2000. Optimal stopping behavior with relative ranks: The secretary problem with unknown population size. *Journal of Behavioral Decision Making* **13** 391–411.

Simon, H. 1955. A behavioral model of rational choice. *Quarterly Journal of Economics* **69** 99-118.

Talluri, K. T, G.J. van Ryzin. 2004. *The Theory and Practice of Revenue Management*, Kluwer Academic Publishers, Norwell, MA.

Weatherford, L.R., S.E. Bodily. 1992. A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking, and pricing. *Operations Research* **40** 831–844.

van Slyke, R., Y. Young. 2000. Finite horizon stochastic knapsack with applications to revenue management. *Operations Research* **42** 1706–1718.

Zhao, W., Y. Zheng. 2000. Optimal dynamic pricing for perishable assets with nonhomogenous demand. *Management Science* **46** 375–388.