A TAXONOMY AND RESEARCH OVERVIEW OF PERISHABLE-ASSET REVENUE MANAGEMENT: YIELD MANAGEMENT, OVERBOOKING, AND PRICING

LAWRENCE R. WEATHERFORD
University of Wyoming, Laramie, Wyoming

SAMUEL E. BODILY
University of Virginia, Charlottesville, Virginia
(Received February 1992; revision received July 1992; accepted July 1992)

This paper proposes the term perishable-asset revenue management to denote the field that combines the areas of yield management, overbooking, and pricing for perishable assets. After summarizing the characteristics common to problems in this field, the paper discusses the objectives and constraints faced by decision makers. Then it offers a comprehensive taxonomy with 14 different elements and reviews the research that has been done related to each element. Finally, it suggests some important areas of future research that can help bridge the gap between theory and application.

Writings on yield management, overbooking and pricing have been appearing for more than 20 years (Rothstein 1971), but the three topics have generally been considered separately. Our research on them, however, has clarified their interrelatedness, but within them, we have noted that many different types of problems are involved. Therefore, our intent here is to provide a framework for categorizing the types of problems in all three topics, to review the subtypes that have been studied, and to indicate areas of future research and application that will bring substantial benefits to practitioners.

Section 1 describes the characteristics that are common to all yield-management problems—perishability, fixed capacity, and the ability to segment. We propose to replace the term yield management with a new, more appropriate term, perishable-asset revenue management (PARM). Section 2 defines the new term.

Section 3 discuses the different objectives and constraints faced by a decision maker in a PARM situation. It also gives a brief summary of the financial measures relevant to these decisions.

Section 4 provides our comprehensive taxonomy (or classification scheme). The discussion emphasizes the wide variety of problem situations that can be considered within PARM.

Section 5 reviews the problems that have been treated in the research to date and notes that, although several problems have been solved in a theoretical manner, an enormous gap still exists between theory and practice. Therefore, Section 6 lists several important problems as candidates for future research because of the important benefits that would come from deriving implementable solutions to them.

1. COMMON CHARACTERISTICS AND TERMS

After we analyzed situations in which yield management is currently practiced, we concluded they had the following characteristics in common:

a. One date on which the product or service becomes available and after which it is either not available or it ages. The product or service cannot be stored, at least not without significant cost or aging. If the item

Subject classifications: Inventory/production, perishable/aging items: problem taxonomy yield including management and overbooking. Marketing, pricing: profit optimization from perishable assets. Marketing, segmentation: price-sensitive customers of perishable assets.
Area of review: OR FORUM

0030-364X/92/4005-0831 $01.25

were storable, one would use inventory-management approaches, which have been well studied (Silver and Peterson 1985). Examples include seats for the theater, a sporting event, or a restaurant; space on any means of transportation, in lodging, or for apartment rental; electricity and other utilities; fashion or high technology goods; services such as the barber, computing, and auto repair; broadcast advertising time periods; and traffic on fiber optic lines.

b. A fixed number of units or, more accurately, because the number of units can always be increased with some time lag, a high cost of adding an incremental unit of capacity. Usually coupled with this fixed cost is a relatively low variable cost for the product or service, which allows a wide range of prices over which selling the product or service is better than to let it be wasted. For instance, in a hotel, the variable cost for an additional occupant may be less than $20 when the room rate is $100; thus a substantial discount is preferable to letting the room be empty. The fixed capacity is generally accompanied by widely fluctuating demand, which makes balancing availability and demand difficult. (Strictly speaking, yield management can be carried out in situations without high fixed costs and high margins, but the rewards for effective yield management are much higher when high fixed costs are combined with high margins.)

The yield management problem is parallel, although opposite to, the well known news vendor problem that guides some of the thinking in inventory and production planning. The news vendor must choose how many papers (product) to order in the face of uncertain demand. In contrast, here we have a fixed amount of product—the number of hotel rooms, say—*and* uncertain demand, and we manage the demand by choosing how much to stimulate demand to use that fixed capacity.

Although fixed capacity is a common characteristic in traditional yield-management situations (e.g., airlines, hotels), its presence is *not* necessary to practice yield management. For instance, many industries (car rental companies, railroads, trucking companies) can increase capacity at a reasonable cost. Allowing for variable capacity requires additional, complex logistics planning. For example, a freight railroad may be willing to accept a freight shipment from Omaha to Chicago at a very low price because additional rail cars are needed in Chicago. On the other hand, the company may be willing to send those same cars to Los Angeles if the customer is willing to pay enough.

Even an airline that never sells out (i.e., effectively has unlimited capacity) should nevertheless stimulate demand early and then raise prices for late arrivers. A final example where capacity is not truly fixed is the hotel (or car rental) industry where there is uncertainty as to when a customer will check out (or return the car). In this case, capacity must actually be forecast! Because fixed capacity is not a defining criterion for yield-management problems, it is added to our taxonomy in Section 4.

c. The possibility of segmenting price-sensitive customers. Segmenting can be harder than it looks because arbitrary price discrimination is not allowed; the variable or characteristic used to segment the market must truly differentiate the product. If offering discount prices and limiting the number who can purchase at the discount price is impossible, legally or practically, the situation may not be amenable to yield management.

The common mechanism used to segment customers in yield-management situations is the *time of purchase*; that is, the less price-sensitive customer generally waits until the last minute to make reservations. On the other hand, people who make their reservations early are generally more price sensitive; they are willing to trade away some flexibility for a reduced price. The discount customer is thus buying a product that is truly differentiated; it has less flexibility, thus it has less value. (Another way to segment customers is by how much business a customer generates; the larger discounts go to the larger customers. This practice is often called *quantity discounting* and is a marketing tool separate from yield management.) In this paper, the scope of yield-management problems is strictly limited to those in which the time of purchase is used to segment potential customers.

2. A NEW TERM AND DEFINITION

Several definitions of yield management have been put forward, but to date no agreement exists on its meaning. For instance, American Airlines (one of the leading users of yield-management tools) defined the objective of yield management as "to maximize passenger revenue by selling the *right* seats to the *right* customers at the *right* time" (1987, emphasis added). Of course, this description reflects the goals of the airline industry. Pfeifer (1989) defined yield management for the airlines as the "process by which discount fares are allocated to scheduled flights for the purposes of balancing demand and increasing revenues." Two experts from the hotel industry give similar definitions from their perspective: "yield management is charging a *different* rate for the *same* service to a *different* individual" (Nykiel 1989), and yield management involves "controlling the *tradeoff* between average rate and occupancy" (Orkin 1989). Cross (1989), president of Aeronomics, said yield management involves "using price incentives and inventory controls to maximize the value of existing processes."

Belobaba (1987) defined yield from the airline perspective as the revenue per passenger mile and stated that yield is a function of two things: the price the airline charges for differentiated service options, and the number of seats sold at each price. Therefore, he reasoned that yield management involves two major components, pricing and seat-inventory control, which are distinct strategies that compose revenue management. Seat-inventory control is "the process of limiting the number of seats to be made available with different price levels on a future flight departure;" inventory management is "the process of *balancing* the number of seats sold at each of the fare levels so as to maximize total passenger revenues."

In light of the preceding discussion, we conclude that the term yield management has limitations in describing the broad class of revenue-management approaches. Because of the way yield is used in the airline industry, the term may even be misleading: If the goal were strictly to maximize yield, a single full-fare passenger on a 747 would achieve the objective. The proper term must allow for finding the optimal tradeoff between average price paid and capacity utilization, and be meaningful in all applicable industries. We propose Perishable-Asset Revenue Management (PARM) and define it to be the optimal revenue management of perishable assets through price segmentation.

The key questions that PARM attempts to answer on a unit-by-unit basis are:

1. How many units should be made available initially at various price levels (or, alternatively, for a given allocation scheme, what are the optimal pricing levels)?
2. How should this availability of units change over time as the time of actual availability approaches; that is, when should certain price levels be closed out (made unavailable) or opened up?

PARM problems often interact with others in operations research (e.g., demand forecasting, quantity discounts), in marketing, and in organizational behavior, which we do not cover here. Readers interested in demand forecasting for airlines may wish to examine Lee's work (1990). Readers interested in quantity discounts may wish to examine research by Crowther (1964), Lal and Staelin (1984), Monahan (1984), or Dolan (1987).

3. MANAGERIAL ASPECTS: OBJECTIVES, CONSTRAINTS, AND COSTS

This section analyzes three different managerial aspects of the problems that arise in PARM: management objectives, constraints, and costs.

3.1. Management Objectives

Among the possible objectives are:

a. Maximize profit/contribution. Throughout this paper, we use the terms profit and contribution almost synonymously. Contribution toward fixed costs is defined as revenue minus variable cost, while profit obviously must subtract fixed costs from the contribution and take taxes into account. In most situations, however, if contribution is maximized, then profit is maximized.

b. Maximize capacity utilization. This goal focuses on selling every available unit, even to the possible detriment of the price obtained. Salespeople may use this approach if they are rewarded according to the number of units sold, a simple but myopic measure.

c. Maximize average revenue/customer. This focus produces results opposite to objective b. If this objective were pursued too vigorously, an airline company, for instance, would end up with only one full-fare passenger on a 747.

d. Maximize revenue. With this objective, the cost side is ignored, perhaps because costs are negligible or essentially fixed and are not a relevant issue for the present decision.

e. Minimize lost customer good will. As an example of this objective, a company might decide that it will never re-open discount price classes after it has closed them. It is hard to see how this aim can be the sole operating objective, but it might be a secondary objective that tempers another one.

f. Maximize the net present value. None of the previous research has used this objective, presumably because the time horizon is so short that the additional accuracy of discounting the cash flows received in different time periods would not be worth the effort.

g. Extract each customer's maximum price. This goal may be the ideal, but is not really possible in most PARM situations, because it would involve a tough negotiation with every customer. Something close to this kind of approach happens in the field of broadcast advertising, where salespeople negotiate a price with each customer for equivalent 30-second time slots.

The vast majority of the operations research literature on PARM has used the criterion of maximizing expected profit. In other words, the assumption is that the decision maker is risk neutral, and therefore, the goal is to find the optimal balance between objectives b and c, where objective a defines the best balance. None of the published research looks at the effect of discounting the cash flows, and this paper, under the assumption that it will not noticeably change the results, will also ignore this effect.

3.2. Constraints

As far as the constraints that the PAR manager may face, the following possibilities exist:

a. Operational constraints

• *a fixed capacity of units to allocate*

In the airlines, as Boeing introduces the new 757 with different cabin lengths (e.g., 150, 200, 250), any of which can be used for a given flight and changed at the last minute, plane size is becoming less of a constraint than in the past. The substitution of plane sizes is possible because the Federal Aviation Administration requires only an advance commitment to the plane type. The fixed-capacity constraint also applies to hotels; in fact, it is more constraining to them than to the airlines.

• *other operational constraints*

In an airline, this constraint includes the currently scheduled routes, the number of planes, and the frequency of flights.

b. Marketing constraints

• *minimum tolerable customer-service levels*

Obviously, the level of customer service can be measured in a variety of ways, for example, the percentage of customers who are bumped because of an overbooked plane or hotel, the percentage of customers who complain about the person next to them getting a cheaper price for the same flight, etc.

c. Strategic constraints

• *long-term vision for the company*

A strategic aim of the board of directors and chief executive officer to be the price leader may define additional constraint(s).

• *competitors' actions*

Examples of this constraint for an airline include competitors' current prices, routes, and flight schedules, as well as any changes to these factors.

3.3. Costs

Two kinds of costs are relevant to PARM decisions: variable costs for the unit of product or service, and costs associated with the event for which a customer is bumped or denied a reservation. The latter is more difficult to measure; lost good will cannot be observed directly. The cost of bumping in a hotel, for example, might be estimated to be the cost of putting the customer in a neighboring hotel, one free night's stay the next time, and a box of chocolates.

4. A COMPREHENSIVE TAXONOMY

Even when we limit PARM problems to those that involve perishability and the ability to segment price-sensitive customers, a variety of interesting problems remains. Table I lists fourteen distinguishing elements. The first eleven characterize the nature of the problem or situation, and the last three are related to choices made, either by the decision maker (L), or the researcher framing the problem (M, N). The order of the descriptors for each element is such that the descriptor generally progresses from the simplest assumption to more complex ones. The numbering scheme is introduced to label problems concisely. The following subsections explain the meanings of the different elements and descriptors in more detail.

4.1. Nature of the Resource

This element is meant to specify the nature of the units of the perishable asset. The units can either be discrete (A1) or continuous (A2). For example, an airline has a discrete number of seats available for sale on a particular airplane (e.g., 23 or 24 but not 23.7), whereas an electric power plant has a continuous range of power available to sell (e.g., 0 to 100 megawatts).

4.2. Capacity

As discussed in Section 1, fixed capacity (B1) has traditionally been associated with a PARM problem, but it is not necessary. PARM can still be practiced profitably with unlimited capacity (B2) if consumers can be segmented and the good or service is perishable.

With unlimited capacity, the issues of stimulation, diversion, and displacement are still active concerns but the bumping procedure is not (Section 4.12).

Table I
Comprehensive Taxonomy

	Elements	Descriptors
A	Resource	Discrete/Continuous
B	Capacity	Fixed/Nonfixed
C	Prices	Predetermined/Set optimally/Set jointly
D	Willingness to Pay	Buildup/Drawdown
E	Discount Price Classes	1/2/3/. . ./I
F	Reservation Demand	Deterministic/Mixed/Random-independent/Random-correlated
G	Show-Up of Discount Reservation	Certain/Uncertain without cancellation/ Uncertain with cancellation
H	Show-Up of Full-Price Reservations	Certain/Uncertain without cancellation/ Uncertain with cancellation
I	Group Reservations	No/Yes
J	Diversion	No/Yes
K	Displacement	No/Yes
L	Bumping Procedure	None/Full-price/Discount/FCFS/ Auction
M	Asset Control Mechanism	Distinct/Nested
N	Decision Rule	Simple Static/Advanced static/Dynamic

4.3. Prices

In typical PARM situations, prices are predetermined (C1) by a pricing group that is separate from any decisions about the number of discount units to sell. One could, however, do exactly the opposite, that is, predetermine the number of units to be sold at each price level and then set the prices in an optimal manner (C2). Weatherford (1991) examined this extension of traditional yield management. Lastly, it is possible to make the pricing and allocation decisions jointly (C3) in the hope of improving profit.

4.4. Willingness to Pay

In airline or hotel situations, a customer's willingness to pay tends to build up (D1) as the availability date draws closer. Discount prices are, therefore, made available to those who reserve early. Some applications work in opposite fashion, however. For example, consider the case of a fashion item or high technology good, such as a computer or an electronic chip-testing instrument. The reasonable approach is to sell at a high price on the date of first availability and to discount later in the product's life cycle. Thus, the price draws down (D2) over time from the first availability of the item; the threshold curve for limiting the number of discount items made available may look like that of Figure 1, just as if the build-up curve were rotated about an axis.

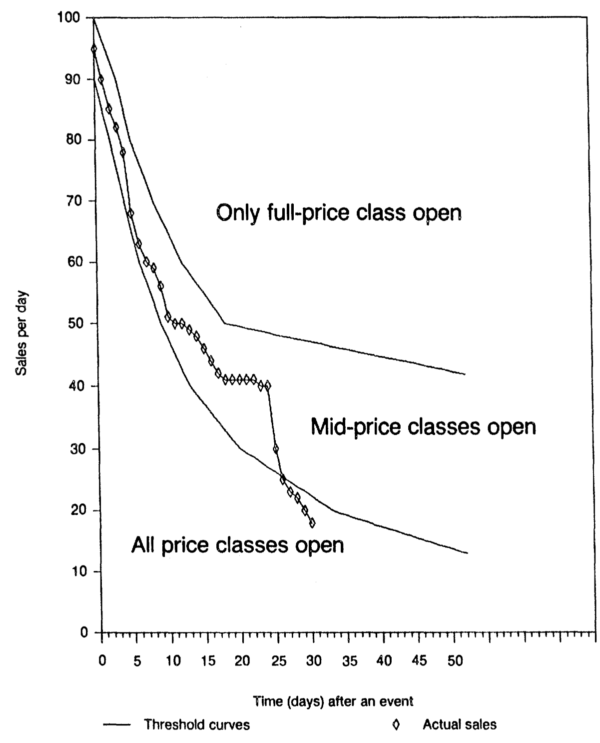

Figure 1. Threshold curves for drawdown pricing.

4.5. Discount Price Classes

The number of discount price classes refers to the number of distinct groups of customers who are willing to buy units at discount prices, which also equals the number of discount price levels that can be offered. We will let I represent the maximum number of discount price classes. Different restrictions are generally attached to the different discount prices so that the products are truly differentiated. For example, if an airline offers a Y fare (i.e., full price) and four discount fares for the same physical seats in the coach cabin (e.g., Q, M, V, and L), the airline would be classified as having an E4 PARM problem.

4.6. Reservation Demand

Each distinct class of discount or full-price customers may have demand that is deterministic, or random and independent of other classes, or random and correlated with other classes. In the deterministic heading, we include a class that has essentially unlimited demand or random demand that is so great that it is always greater than any possible allocation of resource to that price level. For two prices classes, this approach creates five possible combinations and an unwieldy number if I is large. To simplify, all possible combinations are identified by one of four categories: all classes have deterministic demand (F1), or a mix of demand types between deterministic and random-independent (F2), or all classes have random and independent (F3) demand, or all classes have random and correlated demand (F4).

The reason all deterministic is listed as a possible descriptor for this element is that many of the solved problems that address the issue of displacement effects make the assumption of deterministic demand for both full and discount units in order to obtain tractable solutions.

4.7. Show-Up of Discount Reservations

Implicit in the generic case is the assumption that the *discount* customers will show up with certainty (G1) to claim the goods/services or that they pay for it (whether or not they take it) with the provision that the item cannot be re-sold to another party. This situation occurs for rooms in a hotel sold on contract to airline crews, for example. The next level of complexity allows discount customers not to show up at the time of availability (G2) but does not allow any cancellations; that is, the goods/services are reserved, but if the discount customers do not show up to claim them, they do not need to pay. The highest level allows cancellation by discount customers at any time up to and including the time of availability (G3). This situation faces most businesses that have cancellations throughout the process. Because so many reservations are made by travelers who are not serious, airlines often set time limits on reservations: if they are not paid for by a certain date, they automatically cancel. Even though G3 is more complex to model than the other situations, it does make it easier for the business to know ahead of time that a customer is not going to show up.

This element is applicable to *each* discount price class specified in Section 4.5, but it is generally assumed to apply to *all* discount price classes simultaneously. A possible extension of these descriptors would be to include the effects of cancellation penalties.

4.8. Show-Up of Full-Price Reservations

The simplest assumption is that the full-price customers with reservations will show up with certainty (H1) to claim the good/service. This assumption can also be interpreted to mean the customers will pay whether or not they take the good/service and that the item cannot be re-sold to another party. This situation can also model the case in which reservations are not taken for full-price customers; rather, these customers just show up at the time of availability (e.g., the case of an electric utility). The next two levels of complexity are the same as in Section 4.7.

4.9. Group Reservations

All the work to date has ignored this phenomenon, and just treated the group as a collection of individual customers (I1). The possible existence of group reservations (I2) is broken out as a separate element instead of including it under element E (discount price classes) because, in practice, a yes/no decision must be made on accepting the entire group, as opposed to accepting the incremental person in a particular discount price class. Evaluating the value of a group is very much a part of PARM for airlines and hotels.

Other reasons that group bookings are evaluated separately include: the highly variable probability of cancellation, groups often negotiate significant discounts off the regular price, and groups must be treated as a block (e.g., all 80 people in the group are accepted or none at all because they want to stay together).

4.10. Diversion

So far, most work on PARM has assumed (although not very realistically) that the customers who buy at full price are separate from the ones who buy at discount prices. In other words, if an individual is willing to pay full price, he or she will not take a unit at discount (J1). In reality, many business travelers, for example, who are willing to pay full price would also take a discount unit if it were available (J2). By assuming separate segments, one ignores this dilution effect. The only published approaches that deal with diversion (Belobaba 1987, Pfeifer 1989, and Brumelle et al. 1990) assume just two types of consumers (e.g., shoppers and nonshoppers).

Certainly, one can conceive of further refinement of J2 that depends on the descriptor chosen for the number of discount classes. As mentioned, writers (with the exception of Bodily and Weatherford 1991) have assumed only one discount class in solving problems that allowed diversion. If there are two or more discount classes, there are many possible standards for defining the diversionary behavior. Let us state one reasonable standard: *for I discount price classes, the standard defines I + 1 groups of customers*. The first group will buy only the cheapest discount price unit; the second group is willing to buy either of the two least expensive units; and so on, until reaching the $I + 1$st group that will buy the cheapest available of any of the I discount prices *or* the full-price unit. Each class would have some maximum price that it is willing to pay. One can imagine that, in the limiting scenario, if the number of discount price classes approached infinity (each individual were a class), we can classify customers by their reservation price.

4.11. Displacement

To consider displacement effects in the airline industry, one has to consider the overall origin-destination combination of each potential traveler in order to maximize contribution over the entire system. As an example, an airline would have to decide whether it is better to let a discount passenger fly from Charlotte to Baltimore to Boston, or reserve that seat for a different customer who is likely to pay full fare but only fly from Baltimore to Boston.

As a base assumption, displacement is assumed to be nonexistent (K1). That is, for example, we ignore any connecting flights. Displacement can apply only when a package of products or services is sold together (K2). In the hotel industry, displacement considers multiple-night stays where there are different demand levels for each day.

The K2 category may be extended to indicate the number of links of displacement evaluated. In the airlines, a link can be a flight segment, and in the hotels, it can be the number of nights in a stay beyond one.

4.12. Bumping Procedure

What happens when demand exceeds capacity (i.e., the good/service is oversold)? Some businesses curtail reservations beyond the capacity limit (L1) and therefore never have to bump a customer. When overbooking occurs (L2 to L5), an appropriate penalty is applied for lost customer good will. This idea of overbooking has been discussed extensively in the literature, especially in relation to airlines and hotels (Rothstein 1971, 1974, 1985, Ladany 1976). Within the PARM decision, the business may deny service to full-price customers (L2) or to discount customers (L3), or some businesses may prefer to use other popular strategies such as first-come, first-serve (FCFS) (L4), or an auction to get people to give up their reservations (L5).

4.13. Asset Control Mechanism

For this element, we look at a control mechanism for the inventory of units of the perishable asset. At issue is whether a business can ever turn down a request for a full-price unit if discount units are available. In practice, a business would never want to let this happen, but the taxonomy allows for the simple treatment of these two segments as separate with a distinct (M1) limit on the number to sell to each. Because of the stochastic nature of the process, the limit may be reached for the full-price segment *before* the limit is reached for the discount-price segment. Most theoretical approaches that deal with displacement effects via networks assume a distinct control mechanism, while the stopping rules common for other problems assume a nested inventory (M2). This approach sets inventory limits that are nested or transparent from above (e.g., the highest price class—class 0—has an inventory limit equal to the capacity q_0). Figure 2 demonstrates such nested inventory, where q_i is the number of units available for sale at price class i. Belobaba (1987) is one of the main authors to look at the nested issue in detail.

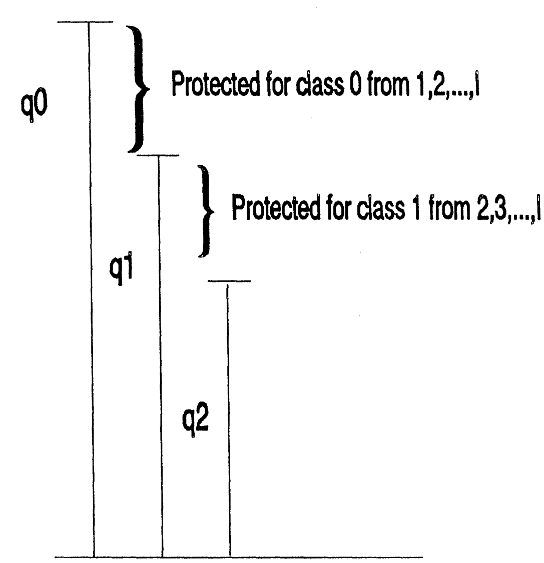

Figure 2. Nested buckets.

4.14. Decision Rules

This element has to do with the way in which the researcher attempts to solve the problem and with the asset control mechanism, rather than the problem itself (to which the first 11 elements relate). Several different decision rules might be used for the two price-class problem, including:

a. pick a fixed allocation q^* prior to reserving the first customer, knowing that it will not be changed later for any reason;

b. pick a fixed time t^* prior to reserving the first customer and accept all requests prior to t^*, knowing that it will not be changed later for any reason;

c. select a (q, t) decision rule (i.e., threshold curve) prior to the first reservation, knowing that this rule will not be changed for any reason;

d. monitor everything continuously (e.g., reservations, competition, and weather) and decide when to curtail reservations.

The first two rules are the simplest and are examples of the simple static rule (N1). By static, we mean the rule does not change for given parameters of the model, and by simple, we mean it only depends on one of the parameters (q or t). Rule c, the advanced static rule (N2), depends on both q and t, a more complex static level. Pfeifer (1989) derived a rule like a and then, through a reinterpretation of the parameters, derived a rule like d. Weatherford (1991) developed an arrival model that implements Pfeifer's rule with the form of either a or b under prior analysis, or c under a preposterior decision analysis. The last descriptor, a dynamic rule (N3) that is dependent on q, t, the path, and previous decisions, is consistent with a dynamic programming approach. Certainly, the implementation of the decision rule must consider the quality of forecasts of passenger behavior that is available.

4.15. Links

It is easy to see the variety in the number of possible PARM problems; according to our taxonomy, there are 124,416 $\cdot$ I different types. With this diversity, a logical approach to analysis is to search for some overlap among the problems. On the other hand, if the taxonomy is mutually exclusive and collectively exhaustive, we will find no overlap. In the taxonomy presented here, there is only one implied relationship between the different elements: A problem classified B2 (i.e., nonfixed capacity) has no need to specify the L element (bumping).

We also mention that the last two elements (asset control mechanism and decision rule) are not totally independent when it comes to implementation. The coordination of an asset control mechanism and a decision rule is very important. Some decision rules work better in terms of capturing more profit with certain asset control mechanisms.

5. CURRENTLY SOLVED PARM PROBLEMS

In this section, we survey the problems for which known research has been conducted. The subsections cover seven broad areas in which solutions have been found. The first is the simple, generic allocation PARM problem; the other six are extensions of the generic problem (usually attempts to move a step closer to more realistic problems). The notation used is the same as the author's original notation. At the end of each subsection, a summary of the optimal rule(s) and their references is provided.

5.1. Simple, Generic Allocation Problem

[A1-B1-C1-D1-E1-F2-G1-H1-I1-J1-K1-L1-M2-N1]

These conditions lead to the second easiest PARM problem that allows for a simple marginal-analysis approach. The easiest problem is classified as F1 (i.e., deterministic full-price demand), which can be solved by inspection. The generic problem was first solved for the airlines by Littlewood and then by Belobaba and others. Bodily and Weatherford stripped away the airline slant and developed a generic problem that can be used by other industries. The rule accepts another discount unit as long as the probability of spill (i.e., the probability that the current discount demand plus the random full-price demand will exceed capacity) is less than the ratio of the reduced price contribution R_1 to the full-price contribution R_0.

Rule: Accept another discount customer ($q_1 \rightarrow q_1 + 1$) if $\Pr_{\text{spill}} < R_1/R_0$, where q_1 is the decision variable that represents the number of discounts to accept.

References: Littlewood (1972), Bhatia and Parekh (1973), Richter (1982), Belobaba (1987), Bodily and Weatherford (1989).

5.2 Optimally Set Pricing

(Same as Section 5.1 except A2-B2-C2-F1)

Weatherford (1991) looks at several pricing problems under the assumption of a fixed-allocation decision. The model incorporates the elasticity of demand and allows diversion (J2). Results are available for demand distributions that are either normal, asymmetric triangular, or uniform.

Rule: See the reference for several variations.

Reference: Weatherford (1991).

5.3. Stochastic Discount Demand

(Generally the same as Section 5.1 except either F3 or F4)

One step toward a realistic allocation problem with predetermined prices is to allow discount demand to be a random variable and allow the possibility of it being correlated with full-price demand. The conditions of one model (Brumelle et al. 1990) allow stochastic dependence between demands in the two fare classes—discount customers (B) and full-fare customers (Y), which is expressed in terms of a correlation value (e.g., demand is assumed to be bivariate normal). The solved model does *not*, however, allow stochastic dependence between Y and the discount seat limit, which is a weakness. In other words, the level choice for the booking limit of discount seats (q) is assumed to have no effect on full-fare demand because this model ignores diversion. (Solved problems with diversion are presented in Section 5.5.)

A generic stopping rule is developed (see the rule) under the assumption that there is a specific and strict sequence to arrivals for demand, that is, discount customers arrive first, followed by full-fare customers. The model does not allow any time overlap between segments, although the authors, citing a simulation study by Titze and Griesshaber (1983), say this assumption can be relaxed somewhat without changing the results. To us, however, the results of the simulation show that the optimal seat allocation decision does change according to the arrival process model (e.g., a strict versus nonstrict sequence) that is chosen. Another restrictive assumption of this formulation is that it requires a positive correlation between the random variables, B and Y, for a bivariate normal distribution. Before adopting the model, one would want to analyze real data to see if that assumption is plausible.

Rule: $q^* = \max\{0 \leqslant q \leqslant Q\colon \Pr[Y > Q - q \mid B \geqslant q] < P_B/P_Y\}$, where Q is the capacity and P is the price subscripted for a given fare class.

Reference: Brumelle et al. (1990).

The authors derived a simple extension wherein the model relaxed the assumption about what happens in the event that demand exceeds capacity (L5 instead of L1). In other words, instead of assuming that the business will just turn away further requests once current reservations reach capacity, a cost (lost good will) is attached to this action.

Rule: $q^* = \max\{0 \leqslant q \leqslant Q\colon \Pr[Y > Q - q \mid B \geqslant q] < P_B/(P_Y + P_G)\}$, where P_G is the value placed on lost good will.

Reference: Brumelle et al. (1990).

5.4. Uncertain Customer Show-Up (Overbooking)

(Generally the same as Section 5.1 except G2 or H2 or both)

Two models (Brumelle et al. 1990, Bodily and Pfeifer 1992) greatly simplified the solution by restricting the analysis to one price class (i.e., no discounts, E0). Under the assumption of a Bernoulli cancellation process, an optimal decision rule can be derived (see the rules). Other models, such as Belobaba's, allow discount price classes (i.e., E1 or higher). Belobaba's model defines a set of terms called the overbooking factors, OV_i ($\geqslant 1.0$), which represents the extent to which expected revenue is reduced in fare class i because of this uncertainty. This new decision rule is also shown here. The easiest assumption, of course, is that $OV_1 = OV_2 = \ldots = OV_k = OV$, in which case the overbooking limit on total capacity (C) becomes $C^* = BL^* = (C) \cdot (OV)$.

Another model (Alstrup et al. 1986) treats the overbooking decision as a Markov decision process and is solved as a two-dimensional stochastic dynamic program, which also allows a dynamic solution (N3). This model is similar to one of Ladany and Bedi (1977) except that they allow penalties for upgrades and downgrades. In Alstrup's model, the plane has a movable divider that separates the C class (full fare) and the M class (discount). The model is also claimed to be a generalization of Rothstein (1971) in that it allows two types of passengers (E1). The *states* in this Markov process are the number of passengers already booked in the C and the M classes. The decision variables are the number of incremental reservations that can be accepted, given the state and the time to departure. One drawback is that the model does not allow diversion (J1). The model came out of Alstrup's experience in Scandinavian airlines.

Rules: Only one price class:

Bodily and Pfeifer:

$$\Phi\left(\frac{N-\mu}{\sigma}\right) = \frac{C}{R+C} + \frac{\phi\sqrt{1-p}}{2\sqrt{Bp}},$$

where N is the capacity, B is the number to book, μ is the mean number of customers who survive given B bookings, σ is the standard deviation of the number of customers who survive given B bookings, C is the overbooking cost penalty, R is the contribution from a customer who shows, and p is the probability that a booked customer actually shows up (survives).

Brumelle et al.:

$$\eta^*(Q) = \max\{\eta\colon \Pr[N(\eta) > Q] < P_Y/(P_Y + P_O)\},$$

where $N(\eta)$ is the number who show up if η customers are booked, and P_O is the penalty for bumping a confirmed customer.

Two or more price classes:

Belobaba:

$$S_j^i\colon (1/OV_i)\cdot \mathrm{EMSR}_i[S_j^i] = f_i \cdot P_i(S_j^i)\cdot (1/OV_i) = f_j \cdot (1/OV_j)$$

Alstrup et al.: No closed form.

References: Shlifer and Vardi (1975), Alstrup et al. (1986), Belobaba (1989), Brumelle et al. (1990), Bodily and Pfeifer (1992).

5.5. Problems With Diversion

(Generally the same as Section 5.1 except they have F4-J2)

The main contribution of these models is the relaxation of the assumption that the consumers neatly divide into two distinct, separate classes, each willing to pay one and only one price. In this section, the models make a more realistic assumption, namely, that people divide into two different kinds of groups (e.g., shoppers and nonshoppers), but the observable distinction is blurry. Pfeifer defined *shoppers* as those who will purchase a seat at discount fare P_D, but will be lost as customers if the discount fare is not available. All other potential customers are called *nonshoppers*, and their behavior is such that they will purchase the item at P_D but are willing to pay full fare P_F if no discount-fare seats are available. His model provides a decision rule that tells when to stop offering discount fares. Once this decision is made, the airline cannot consider offering additional discount fares at a later time. The rule is defined in terms of two probabilities, p_1 = the probability that the q + 1st customer is a shopper, and p_2 = the probability that $Q - q - 1$ full-fare seats will satisfy all subsequent nonshopper demand. One major drawback of this approach is the estimation of p_1 and p_2 as functions of q. This task is not easy, although the observation is made that both p_1 and p_2 decrease as q increases.

A different model by Brumelle et al. looks at the effect of diversion by taking into account the probability γ of an upgrade if a customer is denied a discount seat. The final decision rule is identical to Pfeifer's rule. Both of these rules, as mentioned, are difficult to implement in practice because of the difficulty in estimating the probabilities required.

Belobaba's model defines a probability that the person may upgrade vertically to fare class $i - 1$ (i.e., accept a higher fare for the same flight). This probability is designated $P_i(v)$, and the expected revenue from the shift is $P_i(v)\cdot f_{i-1}$. In this case, the original optimal rule is modified by an amount called the incremental protection (V_2^1) and is provided here. Like the other two rules in this section, a major drawback of implementing this rule would be the difficulty in estimating $P_i(v)$.

Rules:

Pfeifer: Book another as long as $p_1\cdot p_2 < (P_F - P_D)/P_F$.

Brumelle et al.:

$$q^* = \max\{0 \leqslant q \leqslant Q\colon \Pr[Y + U(q) > Q - q \mid B \geqslant q] < (\mathrm{P}_B - \gamma\cdot \mathrm{P}_Y)/((1-\gamma)\cdot \mathrm{P}_Y)\},$$

where $U(q)$ is the number who upgrade if discounts are cut off at q.

Belobaba:

$$\mathrm{EMSR}_1(S_2^1 + V_2^1)\cdot[1 - P_2(v)] + f_1\cdot P_2(v) = f_2.$$

References: Belobaba (1989), Pfeifer (1989), Brumelle et al. (1990).

5.6. Problems With Displacement Effects (Solved Using the Network Approach)

(Generally the same as Section 5.1 except they have K2)

A large number of network solutions has been outlined in theory, and others have actually been implemented. Each one has a different taxonomy description, but they can all be classified as a K2 problem. The solution methods to these models can be separated into two broad categories—optimal or heuristic. Each category is discussed separately here. Note that most of these approaches assume a distinct asset control mechanism (M1).

Optimal Methods. The two major shortcomings of all the optimal mathematical programming and network formulations are the size of the formulations required to model stochastic demand accurately, and the assumed independence of the stochastic demand (F3) for the different price classes, if stochastic demand is allowed. One model that used a network approach to look at the seat allocation decision, including the effects of displacement, took into account diversion (J2), but its major drawback was that all demand estimates were assumed to be deterministic (F1). This model came from work actually implemented at an airline and was designed to be interactive and run on a minicomputer.

Rule: No closed form.

Reference: Glover et al. (1982).

Heuristic Methods. To extend Glover et al.'s model to include probabilistic demand, D'Sylva (1982) used a piecewise linear approximation for the expected revenue curve in a linear programming (LP) formulation. Another model by Wollmer (1985) addresses the multileg case with no assumptions made on the demand patterns for the different classes (i.e., not restricted to F1 any more). The algorithm of the paper is specialized for the single-leg case. An LP, which becomes the network flow model, is then formulated for the single-leg case (K1) with two fare classes (E1). The optimal solution is obviously the same expected marginal revenue solution that was presented in Section 5.1. When the model is extended to allow multiple discount-fare classes (E2 or higher), the model uses binary decision variables in a linear integer programming formulation. An algorithm is presented that is near optimal, and an LP formulation is shown to yield the same near-optimal results. This LP can then be extended to the multileg case (K2); it turns out to be fairly efficient, but the result is not optimal. Another disadvantage is that the parameters have to be updated and the algorithm rerun every time a new reservation is made.

A model by Wang (1983) looks at a multiple origin-destination (O-D) model where all feasible O-D and fare combinations are ranked in terms of their expected marginal revenue. The model allows stochastic independent demand (F3). The concept of expected marginal revenue is similar to Belobaba's (1987). The model examines the expected revenue (i.e., the average yield multiplied by the probability that the seat will be occupied if allocated to that particular combination) for each fare and O-D combination. As an example, assume an airline problem with 2 flight legs (A-B-C) and 3 fares for each O-D. This problem would give 12 O-D/fare combinations (3 from A-C and 9 combinations of A-B and B-C). The model allocates the first seat to the fare/O-D combination with the highest expected revenue and repeats the process for all seats in the plane.

Another heuristic approach to the overall O-D picture that is used by one major airline is called virtual nesting, and it is reported by Williamson and Belobaba. This model attempts to put a value on each fare class that depends on the overall ticket revenue. For example, a discount fare from Atlanta to Washington, D.C. to New York might be ranked similarly in the virtual nest as a full fare from Washington, D.C. to New York. It is basically a greedy approach and does not result in optimal system-wide revenue, but it does appear to be more sophisticated than the single-flight-leg approach. A major obstacle to implementation is that other airline reservation systems work only on a flight-leg basis, so this new system can control only bookings received through the company's own central reservation system.

Rule: No closed form.

References: Buhr (1982), D'Sylva (1982), Wang (1983), Wollmer (1985), Williamson and Belobaba (1988), Dror, Trudeau and Ladany (1988).

5.7. Dynamic Problems

(Generally the same as Section 5.1 except they have N3 or N1/N2 with repeated updating)

These solutions take a step toward realism by allowing the probabilities and other model parameters to be updated and changed over time. A model by Alstrup et al. (1986) that is truly dynamic (N3) was presented in Section 5.4 because it also allows uncertain show-up. The authors model the decision as a Markov decision process and solve it as a two-dimensional stochastic dynamic program. The availability of discount units in given price categories may be open or closed out on any day and subject to daily review. What makes the analysis of this problem more complex than analysis with the other decision rules is that a price class can be closed on one date and opened up later. Examples of future problems that need dynamic solutions are the sequential decisions necessary for three or more price classes with diversion, uncertain show-up throughout the reservation process, and any problem in general that requires sequential decisions to be made.

As discussed earlier in Section 5.1, Littlewood (1972) and others developed the rule of thumb for the simple static problem (N1) that a lower fare class should be shut down when the certain revenue from the lower class is less than the expected revenue from reserving a seat for a potential higher fare customer. Note that such rules (N1) can be applied on an ongoing basis (repeated application). Two different effects need to be distinguished, both of which relate to the nature of the decision rules. The first effect is the type of rule that is being used, which is covered adequately by the descriptors (simple static, advanced static, dynamic) in the taxonomy. The second effect is the periodicity with which the decision maker chooses to implement the decision rule. For instance, a simple static rule may say to keep accepting discounts until the probability of spill is less than 0.73. The decision maker can do one of several things: make a prior forecast on demand, determine how many units to sell at the discount price based on today's forecast and $p_{\text{spill}} = 0.73$, and then never revise the forecast nor the number of units to sell [SINGLE APPLICATION]; or make an initial forecast, determine the initial number of discount units to sell, and then update the forecast half-way through the process using the same $p_{\text{spill}} = 0.73$ to determine a revised number of discount units to sell [DOUBLE APPLICATION]; or make an initial forecast, determine the initial number of discount units to sell, and then update the forecast every time a potential customer arrives [NUMEROUS REPEATED APPLICATIONS].

Static rules that are applied repetitively through time with updated forecasts should not be confused with a dynamic rule (such as a dynamic programming solution), in which p_{spill} may change over time. Even though the repeated application static rule requires constantly updated forecasts, it is based each time on the same p_{spill} (e.g., 0.73). A dominance effect exists such that an advanced static rule is better than a simple static rule (regardless of the number of repeated applications); similarly, a dynamic rule dominates an advanced or simple static rule.

Belobaba's work is an example of repeated applications using his static Expected Marginal Seat Revenue model (several researchers—Curry 1990, Robinson 1990, Brumelle and McGill 1993—have pointed out that Belobaba's result is only optimal for two nested fare classes and it is suboptimal for three or more). As seen below in the rules section, the static rule is a function of time t that shows how the rule is applied repeatedly through time to constantly update the allocation levels. The obvious advantage to repeated applications is that actual bookings are used to reduce the uncertainty in the forecast of expected demand. That is, repeated application of the static decision rule allows the *forecast* of demand to be updated over time as more data become available on the build-up of reservations for the discount units (and full-price units in the more sophisticated problem).

Rules: Dynamic decision rule: No closed form.

Repeated application of static decision rules: $S_2^1(t)$ such that $\text{EMSR}_1^t[S_2^1(t)] = f_1 \cdot P_1^t(S_2^1) = f_2$.

References: Littlewood (1972), Alstrup et al. (1986), Belobaba (1989), Bodily and Weatherford (1989), Curry (1990), Robinson (1990), Brumelle and McGill (1993).

6. CONCLUSIONS AND PROBLEMS FOR FUTURE RESEARCH

This paper has developed a 14-element comprehensive taxonomy and has classified published work on PARM problems using this taxonomy. The common threads of all PARM problems were culled out of present applications, and they can now be used to extend the use of perishable-asset revenue management to numerous other general business settings. The most realistic PARM problem, according to the taxonomy, is *five* discount price classes, *random-correlated* reservation demand, *uncertain* show-up of discount reservations *with cancellation*, *uncertain* show-up of full-price reservations *with cancellation*, group reservations, *diversion*, *displacement*, *nested* asset control mechanism, and a *dynamic* decision rule (E5-F4-G3-H3-I2-J2-K2-M2-N3). The other elements could be any of their possible descriptors, although the most common ones would be *discrete* resource, fixed capacity, *predetermined* prices, *build-up* willingness to pay, and an *auction* bumping procedure (A1-B1-C1-D1-L5). Given this most probable and realistic problem, we can see a serious gap between what needs to be solved and what has been solved already. In general, there is a strong need for solutions to problems that allow multiple realistic assumptions (e.g., uncertain show-up (G2 or H2) and multiple price classes (E2 or higher), or diversion effects (J2) and dynamic decision rules (N3), or optimally set pricing (C2) and random demand (F3 or higher)), and coordination of decision rules with actual forecasting methods and asset control mechanisms. Five of the most critical needs seem to be:

1. provide for coordination of PARM decision rules with forecasting methods and asset control mechanisms actually in use by practitioners;
2. work on the joint generic allocation (C1) and overbooking (G2) problem;

3. work on three (or more, up to general case of I) price classes (E2+) with diversion (J2);
4. an optimal pricing model (C2) that allows uncertain full-price demand (F2 or higher) and other optimal pricing models; and
5. a model that implements the rules in Section 5.5.

The intent of this paper was to bring some consistency to the discussion of the important topic of perishable-asset revenue management and to provide a broad framework from which all interested researchers could work to extend PARM solutions to practical applications.

NOTE IN PRESS

Kimes (1989) reviews the concept of yield management and different solution approaches, and then presents directions for future research.

ACKNOWLEDGMENT

The authors wish to thank an anonymous referee for his insightful comments, particularly as they relate to the actual practice of PARM.

REFERENCES

ALSTRUP, J., S. BOAS, O. B. G. MADSEN AND R. V. V. VIDAL. 1986. Booking Policy for Flights With Two Types of Passengers. *Eur. J. Opnl. Res.* **27,** 274–288.

American Airlines Annual Report. 1987. The Art of Managing Yield, 22–25.

BELOBABA, P. P. 1987. Air Travel Demand and Airline Seat Inventory Management. Ph.D. Dissertation, Flight Transportation Lab Report R 87-7, Massachusetts Institute of Technology, Cambridge, Mass.

BELOBABA, P. P. 1989. Application of a Probabilistic Decision Model to Airline Seat Inventory Control. *Opns. Res.* **37,** 183–197.

BHATIA, A. V., AND S. C. PAREKH. 1973. Optimal Allocation of Seats by Fare. Presentation by TWA to AGIFORS Reservations Study Group.

BODILY, S. E., AND P. E. PFEIFER. 1992. Overbooking Decision Rules. *Omega* **20,** 129–133.

BODILY, S. E., AND L. R. WEATHERFORD. 1989. Yield Management: Achieving Full Profit Potential From Your Market. Working Paper No. 89-30, Darden Graduate School of Business Administration, University of Virginia, Charlottesville, Va.

BODILY, S. E., AND L. R. WEATHERFORD. 1991. Perishable Asset Revenue Management: Yield Management and Pricing. Working Paper, Darden Graduate School of Business Administration, University of Virginia, Charlottesville, Va.

BRUMELLE, S. L., AND J. I. MCGILL. 1993. Airline Seat Allocation With Multiple Nested Fare Classes. *Opns. Res.* **41** (to appear).

BRUMELLE, S. L., J. I. MCGILL, T. H. OUM, K. SAWAKI AND M. W. TRETHEWAY. 1990. Allocation of Airline Seats Between Stochastically Dependent Demands. *Trans. Sci.* **24,** 183–192.

BUHR, J. 1982. Optimal Sales Limits for 2-Sector Flights. *AGIFORS 22nd Annual Symposium Proceedings,* 291–303.

CROSS, R. 1989. Yield Management: New Horizons for a Dark Science. Presented at Yield Management Multi-Industry Conference, Charlotte, N.C.

CROWTHER, J. 1964. Rationale for Quantity Discounts. *Harvard Bus. Rev.*, March-April, 121–127.

CURRY, R. E. 1990. Optimum Airline Seat Allocation With Fare Classes Nested By Origins and Destinations. *Trans. Sci.* **24,** 193–204.

DOLAN, R. J. 1987. Quantity Discounts: Managerial Issues and Research Opportunities. *Market. Sci.* **6,** 1–22.

DROR, M., P. TRUDEAU AND S. P. LADANY. 1988. Network Models for Seat Allocation of Flights. *Trans. Res.-B.* **22B**(4), 239–250.

D'SYLVA, E. 1982. O and D Seat Assignment to Maximize Expected Revenue. Sales and Technology Department, Boeing Commercial Airplane Company, Seattle, Wash.

GLOVER, F., R. GLOVER, J. LORENZO AND C. MCMILLAN. 1982. The Passenger-Mix Problem in the Scheduled Airlines. *Interfaces* **12,** 73–79.

KIMES, S. E. 1989. Yield Management: A Tool for Capacity-Constrained Service Firms. *J. Opns. Mgmt.* **8,** 348–363.

LADANY, S. P. 1976. Dynamic Operating Rules for Motel Reservations. *Dec. Sci.* **7,** 829–840.

LADANY, S. P., AND D. N. BEDI. 1977. Dynamic Booking Rules for Flights With an Intermediate Stop. *Omega* **5,** 721–730.

LAL, R., AND R. STAELIN. 1984. An Approach for Developing an Optimal Quantity Discount Pricing Policy. *Mgmt. Sci.* **30,** 1524–1539.

LEE, A. O. 1990. Airline Reservations Forecasting: Probabilistic and Statistical Models of the Booking Process. Ph.D. Dissertation, Flight Transportation Lab Report R 90-5, Massachusetts Institute of Technology, Cambridge, Mass.

LITTLEWOOD, K. 1972. Forecasting and Control of Passenger Bookings. *AGIFORS 12th Annual Symposium Proceedings*, 95–128.

MONAHAN, J. P. 1984. A Quantity Discount Pricing Model to Increase Vendor Profits. *Mgmt. Sci.* **30,** 720–726.

NYKIEL, R. 1989. 1990's: Decade of Yield Management. Presented at Yield Management Multi-Industry Conference, Charlotte, N.C.

ORKIN, E. 1989. How Does Yield Management Maximize Profits for the Hotel Industry? Presented at Yield Management Multi-Industry Conference, Charlotte, N.C.

PFEIFER, P. E. 1989. The Airline Discount Fare Allocation Problem. *Dec. Sci.* **20,** 149–157.

RICHTER, H. 1982. The Differential Revenue Method to Determine Optimal Seat Allotments by Fare Type. *AGIFORS 22nd Annual Symposium Proceedings,* 339–362.

ROBINSON, L. W. 1990. A Note on Belobaba's 'Application of a Probabilistic Model to Airline Seat Inventory Control.' Working Paper No. 90-03, Johnson Graduate School of Management, Cornell University, Ithaca, N.Y.

ROTHSTEIN, M. 1971. An Airline Overbooking Model. *Trans. Sci.* **5,** 180–192.

ROTHSTEIN, M. 1974. Hotel Overbooking as a Markovian Sequential Decision Process. *Dec. Sci.* **5,** 389–404.

ROTHSTEIN, M. 1985. OR and the Airline Overbooking Problem. *Opns. Res.* **33,** 237–248.

SHLIFER, E., AND Y. VARDI. 1975. An Airline Overbooking Policy. *Trans. Sci.* **9,** 101–114.

SILVER, E. A. 1981. Operations Research in Inventory Management: A Review and Critique. *Opns. Res.* **29,** 628–645.

SILVER, E. A., AND R. PETERSON. 1985. *Decision Systems for Inventory Management and Production Planning.* John Wiley, New York.

TITZE, B., AND R. GRIESSHABER. 1983. Realistic Passenger Booking Behaviour and the Simple Low-Fare/High-Fare Seat Allotment Model. *AGIFORS 23rd Annual Symposium Proceedings,* 197–223.

WANG, K. 1983. Optimum Seat Allocation for Multi-Leg Flights With Multiple Fare Types. *AGIFORS 23rd Annual Symposium Proceedings,* 225–246.

WEATHERFORD, L. R. 1991. Perishable Asset Revenue Management in General Business Situations. Ph.D. Dissertation, Darden Graduate School of Business Administration, University of Virginia, Charlottesville, Va.

WILLIAMSON, E. L., AND P. P. BELOBABA. 1988. Optimization Techniques for Seat Inventory Control. *AGIFORS 28th Annual Symposium Proceedings,* 153–170.

WOLLMER, R. D. 1985. An Airline Reservation Model for Opening and Closing Fare Classes. Unpublished Internal Report, McDonnell-Douglas, St. Louis, Mo.