

Introduction to the Theory and Practice of Yield Management¹

Serguei Netessine
The Wharton School
University of Pennsylvania

netessine@wharton.upenn.edu

Robert Shumsky
W. E. Simon Graduate School of
Business Administration
University of Rochester

shumsky@simon.rochester.edu

1. Introduction

A variety of concepts and analytical tools fall under the label *yield management*. The term is used in many service industries to describe techniques to allocate limited resources, such as airplane seats or hotel rooms, among a variety of customers, such as business or leisure travelers. By adjusting this allocation a firm can optimize the total revenue or "yield" on the investment in capacity. Since these techniques are used by firms with extremely perishable goods, or by firms with services that cannot be stored at all, these concepts and tools are often called perishable asset *revenue management*.

The techniques of yield management are relatively new - the first research to deal directly with these issues appeared less than 20 years ago. These days, yield management has been an enormously important innovation in the service industries. American Airlines credits yield management techniques for a revenue increase of \$500 million/year and Delta Airlines uses similar systems to generate additional revenues of \$300 million per year² While the airlines are the oldest and most sophisticated users of yield management, these practices also appear in

other service industries. For example, Marriott Hotels credits its yield management system for additional revenues of \$100 million per year, with relatively small increases in capacity and costs. Broadcasting companies use yield management to determine how much inventory (advertising slots) to sell now to the "upfront market" and how much to reserve and perhaps sell later at a higher price to the "scatter market." Even manufacturers are using yield management techniques to increase profits. After all, manufacturing capacity is as perishable as an airline seat or an advertising slot - if it is not used when it is available, that opportunity to use the capacity is gone forever.

The purpose of this note is to introduce the fundamental concepts and trade-offs of yield management and to describe the parallels between yield management and the newsvendor framework that is an important model for inventory management. In particular, this note focuses on how a manager might allocate perishable inventory among a variety of customer segments. To acquire some intuition about the problem we will spend most of our time with an application that involves two types of customers. Consider an establishment we will call the Eastman Towers Hotel of Rochester, NY. The hotel has 210 King/Queen rooms used by both business and leisure travelers. The hotel must decide whether to sell rooms well in advance at a relatively low price (i.e., to leisure travelers), or to 'hold out' and wait for a sale at a higher price to late-booking business travelers.

Our solution to this problem will be simple - we will find a single 'cap' or 'booking limit' for the number of rooms to sell in advance to leisure travelers. There are also variations of the problem and the solution that are more complicated, and these variations can improve the performance of the yield management system. For example, we might change the booking limit up or down as time passes, or we might assume that some rooms may be used only for some customers while other rooms are more flexible (think

¹We would like to thank anonymous referees, Krishnan Anand, Gerard Cachon, Yu-Sheng Zheng, MBA students in the OPIM632 class at the Wharton School and the OMG402 class at the Simon School for comments and suggestions.

²Andrew Boyd, "Airline Alliance Revenue Management". *OR/MS Today*, Vol. 25, October 1998.

of 'economy' and 'deluxe' rooms). We will discuss these extensions in the last section of this note.

In general, the yield management/booking limit decision is just one aspect of a more general business question: How should a firm market and distribute goods to multiple customer segments? To answer this question, a firm must use tools for pricing and forecasting as well as for inventory management. All of these tools are often grouped under the term *revenue management*, although the boundaries between yield management and revenue management are often ambiguous. In this note we will specifically focus on just a single aspect of revenue management - yield management.

We next identify the environments in which yield management techniques are most successful. Sections 3-5 examine a specific solution to the basic Eastman Towers problem, while Section 6 will discuss a similar important problem, 'overbooking.' In Section 7 we describe complications and extensions of the basic problem. Appendix A is a list of suggested additional readings on yield management, Appendix B contains exercises that test our understanding of the yield management problem³.

2. Where and Why Firms Practice Yield Management

Business environments with the following five characteristics are appropriate for the practice of yield management (in parentheses we apply each characteristic to the Eastman Towers Hotel):

1. It is expensive or impossible to store excess resource (we cannot store tonight's room for use by tomorrow night's customer).
2. Commitments need to be made when future demand is uncertain (we must set aside rooms for business customers - "protect" them from low-priced leisure travelers - before we know how many business customers will arrive).
3. The firm can differentiate among customer segments, and each segment has a different demand curve (purchase restrictions and refundability re-

quirements help to segment the market between leisure and business customers. The latter are more indifferent to the price.).

4. The same unit of capacity can be used to deliver many different products or services (rooms are essentially the same, whether used by business or leisure travelers).
5. Producers are profit-oriented and have broad freedom of action (in the hotel industry, withholding rooms from current customers for future profit is not illegal or morally irresponsible. On the other hand, such practices are controversial in emergency wards or when allocating organs for transplantation).

Given these characteristics, how does yield management work? Suppose that our hotel has established two fare classes or buckets: *full price* and *discount price*. The hotel has 210 rooms available for March 29 (let us assume that this March 29 is a Monday night). It is now the end of February, and the hotel is beginning to take reservations for that night. The hotel could sell out all 210 rooms to leisure travelers at the discount price, but it also knows that an increasing number of business customers will request rooms as March 29 approaches and that these business customers are willing to pay full price. To simplify our problem, let us assume that leisure demand occurs first and then business demand occurs. Hence we must decide how many rooms we are willing to sell at the leisure fare or in other words, how many rooms shall we protect (i.e., reserve) for the full price payers. If too many rooms are protected, then there may be empty rooms when March 29 arrives. If too few are protected, then the hotel forgoes the extra revenue it may have received from business customers.

Notice we assume that the hotel can charge two different prices for the same product (a room). To separate the two customer segments with two different demand curves, the hotel must practice market segmentation. In the hotel example, we assume that the firm can charge different prices to business and lei-

³Solutions are available for interested faculty upon request

sure customers⁴. In order to differentiate between these two groups, a firm often introduces booking rules that create barriers or "fences" between market segments. For example, a Saturday-night stay may be required to receive a discounted room on Monday, because most business travelers prefer to go home on the weekend while leisure customer are more likely to accept, or may even prefer, the weekend stay. Again, leisure customers are more price-sensitive so the hotel wishes to sell as many rooms to business customers at a higher price as possible while keeping room utilization high.

booking limit = 210 - the protection level

Therefore, the hotel's task is to determine either a booking limit or a protection level, since knowing one allows you to calculate the other. We shall evaluate the protection level. Suppose the hotel considers protection level ' Q ' instead of current protection level $Q+1$ (Q might be anything from 0 to 209). Further, suppose that $210-Q-1$ rooms have already been sold (see Figure 1). Now a prospective customer calls and wishes to reserve the first 'protected' room at the discount price.




Figure 1: Protection Level and Booking Limit in the Hotel

Booking Limits and Protection Levels

Before we look at the mathematics that will help us to make this decision, some vocabulary would be helpful. We define a *booking limit* to be the maximum number of rooms that may be sold at the discount price. As we noted previously, we assume that leisure customers arrive *before* business customers so the booking limit constraints the number of rooms that these customers get: once the booking limit is reached, all future customers will be offered the full price. The *protection level* is the number of rooms we will *not* sell to leisure customers because of the possibility that business customers might book later in time. Since there are 210 rooms available in the hotel, and just two fare classes, in our example,

Should the hotel lower the protection level from $Q+1$ to Q and therefore allow the booking of the $(Q+1)$ th room at the discount price? Or should it refuse the booking to gamble that it will be able to sell the very same room to a full price customer in the future? The answer, of course, depends on (i) the relative sizes of the full and discount prices and (ii) the anticipated demand for full price rooms. The decision is illustrated as a decision tree in Figure 2.

⁴As we mentioned in the introduction, finding the appropriate price for each customer segment is an important part of revenue management. However, the pricing problem is not within the scope of this note. Here we assume that prices are established in advance.

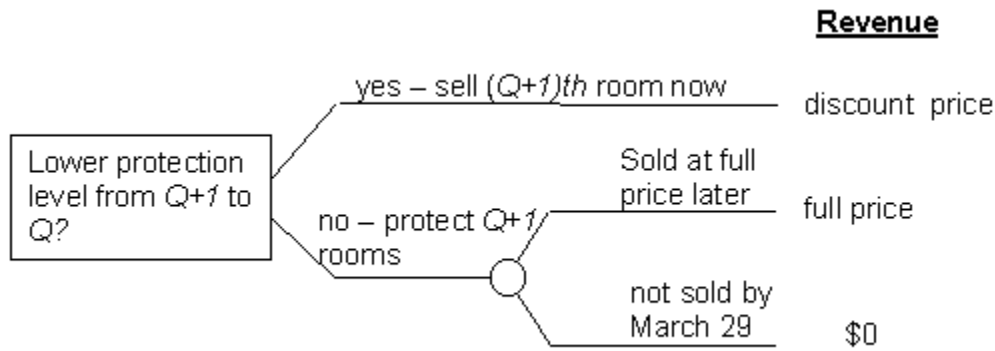


Figure 2 : Deciding the protection level

In the next section we associate numbers with this decision and find the optimal protection level, Q^* .

Solving the Problem

To determine the value of each branch of the decision tree in Figure 2, we need to know the probability for each 'chance' branch and the values at the end of the branches. Suppose that the discount price is \$105 per night while the full price is \$159 per night. To find the probability on each branch, define random variable D to represent the anticipated demand for rooms at the full price. The hotel may estimate the distribution of D from historical demand, as well as from forecasts based on the day of the week, whether there is a holiday, and other predictable events. Here we will assume that the distribution is derived directly from 123 days of historical demand, as shown in Table 1 below.

In the table, the 'Cumulative Probability' is the fraction of days with demand at or below the number of rooms in the first column (Q).

Table 1. Historical demand for rooms at the full fare.

Demand for rooms at full fare (Q)	# Days with Demand	Cumulative Probability $F(Q) = Prob\{D \leq Q\}$
0 - 70	12	0.098
71	3	0.122
72	3	0.146
73	2	0.163
74	0	0.163
75	4	0.195
76	4	0.228
77	5	0.268
78	2	0.285
79	7	0.341
80	4	0.374
81	10	0.455
82	13	0.561
83	12	0.659
84	4	0.691
85	9	0.764
86	10	0.846
above 86	19	1.000
Total	123	1.000

Now consider the decision displayed in Figure 2. If we decide to protect the $(Q+1)^{th}$ room from sale, then that room may, or may not, be sold later. It will be sold only if demand D at full fare is greater than or equal to $Q+1$, and this event has probability

$1-F(Q)$. Likewise, the protected room will not be sold if demand is less than or equal to Q , with probability $F(Q)$. Figure 3 shows our decision with these values included.

Now, $F(Q)$ is the third column in Table 1. We simply scan from the top of the table towards the bottom until we find the smallest Q with a cumulative value greater than or equal to 0.339. The answer here is that the optimal protection level is $Q^*=79$ with a cumulative value of 0.341. We can now evaluate our booking limit: $210 - 79 = 131$. If we choose a larger Q^* , then we would be protecting too many rooms thereby leaving too many rooms unsold on average.




Figure 3 : Booking Limit Decision with Data

Given Figure 3, we can calculate the value of lowering the protection level from $Q+1$ to Q . Lowering the protection level results in selling the $(Q+1)^{th}$ room at a discount which guarantees revenue of \$105. Protecting $Q+1$ rooms has an expected value equal to:

$$(1 - F(Q))(\$159) + F(Q)(\$0) = (1 - F(Q)) (\$159)$$

Therefore, we should lower the protection level to Q as long as:

$$(1 - F(Q))(\$159) \leq \$105$$

or

$$F(Q) (\$159 - \$105) / \$159 = 0.339.$$

If we set Q^* at a smaller value, we are likely to sell too many rooms at a discount thereby turning away too many business customers on average.

3. A General Formula

The solution described above is an example of a standard technique that was developed for the airline industry. The technique was named "Expected Marginal Seat Revenue" (EMSR) analysis by Peter Belobaba at MIT⁵. In our example we had two fare classes with prices r_L and r_H (r_H is the higher price, \$159, while r_L is the lower price, \$105). We also had a random variable, D , to represent the distribution of demand at the high fare. Since there are just two fare classes, the optimal booking limit for low fare class is equal to the total capacity minus Q^* .

The EMSR analysis can actually be described as a newsvendor problem. We make a fixed decision, Q , the protection level, then a random event occurs, the number of people requesting a room at the full fare. If we protect too many rooms, i.e., $D < Q$, then we

⁵Peter P. Belobaba "Application of a Probabilistic Decision Model to Airline Seat Inventory Control". *Operations Research* Vol. 37, No. 2, 1989.

end up earning nothing on $Q-D$ rooms. So the overage penalty C is r_L per unsold room. If we protect too few rooms, i.e., $D < Q$, then we forgo $r_H - r_L$ in potential revenue on each of the $D-Q$ rooms that could have been sold at the higher fare so the underage penalty B is $r_H - r_L$. The critical ratio is

$$\frac{B}{B+C} = \frac{r_H - r_L}{r_L + r_H - r_L} = \frac{r_H - r_L}{r_H}$$

From the newsvendor analysis, the optimal protection level is the smallest value Q^* such that

$$F(Q^*) \geq \frac{B}{B+C} = \frac{r_H - r_L}{r_H}$$

4. Overbooking

Another important component in the yield management toolbox is the use of overbooking when there is a chance that a customer may not appear. For example, it is possible for a customer to book a ticket on an airline flight and not show up for the departure. If that is the case, the airline may end up flying an empty seat resulting in lost revenue for the company. In order to account for such no-shows, airlines routinely overbook their flights: based on the historical rate of no-shows the firm books more customers than available seats. If, by chance, an unusually large proportion of the customers show up, then the firm will be forced to 'bump' some customers to another flight. Hotels, rental car agencies, some restaurants, and even certain non-emergency health-care providers also overbook. When determining the optimal level of overbooking, the calculation is similar to the calculation used for yield management. The optimal overbooking level balances (i) lost revenue due to empty seats and (ii) penalties (financial compensation to bumped customers) and loss of

customer goodwill when the firm is faced with more demand than available capacity.

Let X be the number of no-shows and suppose we forecast the distribution of no-shows with distribution function $F(x)$. Let Y be the number of seats we will overbook, i.e., if the airplane has S seats then we will sell up to $S+Y$ tickets⁶. As in the newsvendor model, define the underage penalty by B and the overage penalty by C . In this case C represents ill-will and the net penalties that are associated with refusing a seat to a passenger holding a confirmed reservation (the 'net' refers to the fact that the airline may collect and hold onto some revenue from this passenger as well). Here, B represents the opportunity cost of flying an empty seat. To explain further, if $X > Y$ then we could have sold $X-Y$ more seats and those passengers would have seats on the plane. So B equals the price of a ticket. If $X < Y$ then we need to bump $Y-X$ customers and each has a net cost of C . Thus, the formula for an optimal number overbooked seats takes following familiar form: is smallest value Y^* such that

$$F(Y^*) \geq \frac{B}{B+C}$$

As an example, suppose the Eastman Towers Hotel described above estimates that the number of customers who book a room but fail to show up on the night in question is Normally distributed with mean 20 and standard deviation 10. Moreover, Eastman Towers estimates that it costs \$300 to "bump" a customer (the hotel receives no revenue from this customer, and the \$300 is the cost of alternative accommodation plus a possible gift certificate for a one-night stay at Eastman in the future). On the other hand, if a room is not sold then the hotel loses revenue equal to one night sold at a discount since at

⁶Of course, the number of no-shows depends upon the total number of tickets that we sell. In the following calculations we assume that the distribution of X is the distribution of no-shows given that we sell exactly S tickets, and that the Y additional customers are guaranteed to show up. A more accurate model would account for no-shows among the extra Y customers, but for most applications this more complicated model would produce a

the very least this room could have been sold to leisure customers⁷. How many bookings should Eastman allow? The critical ratio is

$$\frac{B}{B+C} = \frac{105}{300+105} = 0.2592.$$

From the Normal distribution table we get $\Phi(-0.65) = 0.2578$ and $\Phi(-0.64) = 0.2611$, so the optimal z^* is approximately -0.645 and the optimal number of rooms to overbook is $Y^* = 20 - 0.645 * 10 = 13.5$ (one can also use Excel's Norminv function for this calculation: "`=Norminv(0.2592,20,10)`" gives us an answer of 13.5). If we round up to 14, this means that up to $210 + 14 = 224$ bookings should be allowed.

5. Complications and Extensions

There are a wide variety of complications we face when implementing a yield management system. Here we discuss a few of the more significant challenges.

Demand Forecasting

In the examples above, we used historical demand to predict future demand. In an actual application, we may use more elaborate models to generate demand forecasts that take into account a variety of predictable events, such as the day of the week, seasonality, and special events such as holidays. In some industries greater weight is given to the most recent demand patterns since customer preferences change rapidly. Another natural problem that arises during demand forecasting is censored data, i.e., company often does not record demand from customers who were denied a reservation.

In our example in Sections 3-5 we used a discrete, empirical distribution to determine the protection level. A statistical forecasting model would generate a continuous distribution, such as a Normal or t distribution. Given a theoretical distribution and its parameters, such as the mean and variance, we would again place the protection level where the distribu-

tion has a cumulative probability equal to the critical fractile.

Dynamic Booking Limits

By observing the pattern of customer arrivals, firms can update their demand forecasts, and this may lead to changes in the optimal booking limit. In fact, many airline yield management systems change booking limits over time in response to the latest demand information. For example, it is possible that an airline may raise a booking limit as it becomes clear that demand for business-class seats will be lower than was originally expected. Therefore, during one week a leisure customer may be told that economy-class seats are sold out, but that same customer may call back the next week to find that economy seats are available.

Variation and Mobility of Capacity

Up to now, we have assumed that all units of capacity are the same; in our Eastman example we assumed that all 210 hotel rooms were identical. However, rooms often vary in size and amenities. Airlines usually offer coach and first-classes. Car rental firms offer subcompact, compact, and luxury cars. In addition, car rental firms have the opportunity to move capacity among locations to accommodate surges in demand, particularly when a central office manages the regional allocation of cars. The EMSR framework described above can sometimes be adapted for these cases, but the calculations are much more complex. Solving such problems is an area of active research in the operations community.

Mobility of Customer Segments

In the example of Sections 3-5 we assumed that a leisure customer who is not able to book a room at the discount price (because of the booking limit) does not book any room from that hotel. In fact, some proportion of leisure customers who are shut out from discount rooms may then attempt book a room at the full price. The possibility that a customer may 'buy-up' complicates the model but, most modern yield management systems take such customer movements into account.

⁷Technically, this empty room could have been sold at either low or high price so it is possible that lost revenue is actually higher than we suggest. Also, often hotels charge a penalty for late cancellations.

Nonlinear Costs for Overbooking

In some cases the unit cost of overbooking increases as the number of 'extra' customers increases. This is particularly true if capacity is slightly flexible. For example, a rental car agency faced with an unexpected surge of customers may be able to obtain, for a nominal price, a few extra cars from a nearby sister facility. However, at some point the agency's extra supply will run out - and then customers will be lost, resulting in a much greater cost.

Customers in a Fare Class Are Not All Alike

While two leisure travelers may be willing to pay the same price for a particular night's stay, one may be staying for just one day while the other may occupy the room for a week. A business traveler on an airplane flight may book a ticket on just one leg or may be continuing on multiple legs. Not selling a ticket to the latter passenger means that revenue from all flight legs will be lost.

In each of these cases, the total revenue generated by the customer should be incorporated into the yield management calculation, not just the revenue generated by a single night's stay or a single flight leg.

There are additional complications when code-sharing partners (distinct airlines that offer connecting flights among one another) operate these flight legs. If code-sharing occurs, then each of the partners must have an incentive to take into consideration the other partners' revenue streams.

A similar complication occurs when there are group bookings. How should the hotel consider a group booking request for 210 rooms at a discount rate? Clearly, this is different from booking 210 individual rooms since denying a reservation to one out of 210 group customers may mean that all 210 will be lost to the hotel. Frequently, companies restrict group reservations during the peak seasons but up till now there are no general rules for handling group reservations.

Summary

Table 2 summarizes the impact of these issues on three industries: airlines, hotels and car rental firms⁸. Sophisticated yield management tools have been developed in all three industries, and these tools take industry-specific factors into account. However, all of these tools are based on the basic EMSR model described above.

Parameter	Airline	Hotel	Car rental
Unit of capacity	Seat	Room	Car
Number of resource types	2-3 (e.g., 1 st -class and coach seats)	2-10+	5-20+
"Capacity" at a location fixed or variable	Fixed	Fixed	Variable
Mobility of capacity	Small	None	Considerable
Number of possible prices per unit	Many (3-7+)	Few (2-3+)	Many(4-20+)
Duration of use	Fixed	Variable	Variable
Corporate discounts	Occasional	Yes	Yes
Capacity managed locally or centrally	Central	Central/local	Central/regional/local

Table 1: Comparison of yield management applications

⁸Adopted from Carrol W. J. and R. C. Grimes. "Evolutionary change in product management: experiences the car rental industry." *Interfaces*, 1995, vol.25, no.5, 84-104.

Appendix A. Suggested further readings and cases.

Handbook of Airline Economics. 1995. Aviation Week Group, a division of McGraw-Hill Companies, Chapters 47 - 50.

Handbook of Airline Marketing. 1998. Aviation Week Group, a division of McGraw-Hill Companies, Chapters 23 - 30.

R. G. Cross. 1996. *Revenue Management: Hard-Core Tactics for Market Domination*. Broadway Books.

Dhebar, A. and A. Brandenburger. 1993. American Airlines, Inc.: Revenue Management. Harvard Business School Case.

M. K. Geraghty and E. Johnson. 1997. "Revenue Management Saves National Car Rental". *Interfaces*. Vol.27, pp. 107-127.

J. I. McGill and G. J. Van Ryzin. 1999. "Revenue Management: Research Overview and Prospects." *Transportation Science*. Vol.33, pp. 233-256.

R. D. Metters. Yield Management at Pinko Air. Southern Methodist University Case.

B. C. Smith, J. F. Leimkuhler, and R. M. Darrow. 1992. "Yield Management at American Airlines". *Interfaces*. Vol.22, pp. 8-31.

K. T. Talluri and G. J. van Ryzin. 2002. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, Dordrecht, The Netherlands. To be published.

G. J. van Ryzin. 1998. Transportation National Group. Columbia University Case.

Appendix B. Suggested exercises.*Problem 1*

Recall the yield management problem faced by the Eastman Towers Hotel, as described in the note "Introduction to the Theory and Practice of Yield Management." A competing hotel has 150 rooms with standard Queen-size beds and two rates: a full price of \$200 and a discount price of \$120. To receive the discount price, a customer must purchase the room at least two weeks in advance (this helps to distinguish between leisure travelers, who tend to book early, and business travelers who value the flexibility of booking late). You may assume that if a leisure traveler is not able to get the discount rate, she will choose to book at another hotel.

For a particular Tuesday night, the hotel estimates that the average demand by business travelers has a mean of 70 rooms and a standard deviation of 29 rooms. Assume that demand follows a Normal distribution around the forecast.

- a) Find the optimal protection level for full price rooms (the number of rooms to be protected from sale at a discount price).
- b) Find the booking limit for discount rooms.
- c) Suppose that for a short time, the hotel's forecast of business customer demand is biased upward: the forecast of 70 rooms is too high and fewer business customers appear, on average. Qualitatively describe the economic consequences of using the protection level and booking limit derived in (a) and (b).
- d) Suppose that for a short time, the hotel's forecast of business customer demand is biased downward: the forecast of 70 rooms is too low and more business customers appear, on average. Qualitatively describe the economic consequences of using the protection level and booking limit derived in (a) and (b).

Problem 2

An airline offers two fare classes for coach seats on a particular flight: full-fare class at \$440/ticket and economy class at \$218/ticket. There are 230 coach seats on the aircraft. Demand for full-fare seats has a mean of 43, a standard deviation of 8, and the following empirical distribution:

Economy-class customers must buy their tickets three weeks in advance, and these tickets are expected to sell out.

- a) Find the (i) protection level and (ii) booking limit for low-fare seats.
- b) Suppose that unsold seats may sometimes be sold at the last minute at a very reduced rate (similar to USAirways' "esavers" for last-minute travel). What effect will this have on the protection level calculated in (a)?

The protection level (Q^*) will be...

Circle one: Higher Lower The same

Justification:

Full-fare demand (Q)	Probability Prob{d=Q}	Cumulative probability Prob{d≤Q}
40	0.02	0.25
41	0.06	0.31
42	0.04	0.35
43	0.01	0.36
44	0.06	0.42
45	0.07	0.49
46	0.02	0.51
47	0.03	0.54
48	0.03	0.57
49	0.05	0.62
50	0.03	0.65
51	0.05	0.70
52	0.04	0.74
53	0.06	0.80
54	0.09	0.89
55	0.11	1.00

Problem 3.

Sunshine Airlines flies a direct flight from Baltimore to Orlando. The airline offers two fares: a 14-day advance purchase economy fare of \$96 and a full fare of \$146 with no advance purchase requirement. In its yield management system, Sunshine has two fare classes and, based on historical demand patterns, has established a booking limit of 90 seats for the economy fare class.

Because KidsWorld is located near Orlando, Sunshine strikes a deal with KidsWorld Inc. The deal allows Sunshine to offer a 5-day KidsWorld pass to anyone who purchases an economy-class Sunshine ticket (the customer pays \$96 for the Sunshine ticket and an additional \$180 for the KidsWorld pass, a discount from the regular price of \$200 for a 5-day pass purchased directly from KidsWorld). Sunshine passes along the \$180 pass revenue to KidsWorld, but as part of the agreement, Sunshine receives a \$50 incentive or 'kick-back' from KidsWorld for every pass sold through Sunshine.

Given that Sunshine has accepted the deal with KidsWorld, should Sunshine change the booking limits and protection levels for the Orlando flight? Should it make other changes to its fare class structure? If not, why not? If yes, how? Your answer should be qualitative, but you should be as specific as possible.

Problem 4.

WZMU is a television station that has 25 thirty-second advertising slots during each evening. The station is now selling advertising for the first few days in November. They could sell all the slots now for \$4,000 each, but because November 7 will be an election day, the station may be able to sell slots to political candidates at the last minute for a price of \$10,000 each. For now, assume that a slot not sold in advance and not sold at the last minute is worthless to WZMU.

To help make this decision, the sales force has created the following probability distribution for last-minute sales:

- a) How many slots should WZMU sell in advance?
- b) Now suppose that if a slot is not sold in advance and is not sold at the last minute, it may be used for a promotional message worth \$2500. Now how many slots should WZMU sell in advance?

Demand	Probability	Cumulative
d		Prob(D ≤ d)
8	0	0
9	0.05	0.05
10	0.1	0.15
11	0.15	0.3
12	0.2	0.5
13	0.1	0.6
14	0.1	0.7
15	0.1	0.8
16	0.1	0.9
17	0.05	0.95
18	0.05	1
19	0	1

Problem 5

An aircraft has 100 seats, and there are two types of fares: full (\$499) and discount (\$99).

- a) While there is unlimited demand for discount fares, demand for full fares is estimated to be Poisson with mean $\lambda=20$ (the table below gives the distribution function). How many seats should be protected for full-fare passengers?
- b) An airline has found that the number of people who purchased tickets and did not show up for a flight is normally distributed with mean of 20 and standard deviation of 10. The airline estimates that the ill will and penalty costs associated with not being able to board a passenger holding confirmed reservation are estimated to be \$600. Assume that opportunity cost of flying an empty seat is \$99 (price that discount passenger would pay). How much should airline overbook the flight?

x	Prob(D ≤ x)
20	0.559
21	0.644
22	0.721
23	0.787
24	0.843
25	0.888
26	0.922
27	0.948
28	0.966
29	0.978
30	0.987