



INFORMS Transactions on Education

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Case Series—BlueSky Airlines: Single-Leg Revenue Management

Robert A. Shumsky,

To cite this article:

Robert A. Shumsky, (2009) Case Series—BlueSky Airlines: Single-Leg Revenue Management. INFORMS Transactions on Education 9(3):140-144. <https://doi.org/10.1287/ited.1090.0033cs1>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.


The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2009, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>



Case Series

BlueSky Airlines: Single-Leg Revenue Management

Robert A. Shumsky

Tuck School of Business, Dartmouth College, Hanover, New Hampshire 03755,
robert.shumsky@dartmouth.edu

Case A

BlueSky operates Flight 97, a nonstop flight from JFK to Salt Lake City that departs at 9:30 P.M. For this route they fly an Airbus A320 that can carry 146 passengers. On the airplane, all seats are economy-class (there is no business or 1st-class), and the marginal cost of each additional passenger in these seats is negligible.

All passengers on Flight 97 purchase one-way tickets on that flight alone, and therefore are not connecting from other flights operated by BlueSky. The airline has constructed a two-tier fare structure for this flight: Advance purchase tickets (i.e., nonrefundable tickets purchased at least 14 days in advance) cost \$114 one-way. Full-fare refundable tickets purchased at any time cost \$174. There is heavy demand for advance purchase tickets, so BlueSky could sell out the aircraft with these discount passengers. Therefore, BlueSky's revenue management system may need to protect a certain number of seats for full-fare tickets.

On March 5, 2007, BlueSky is setting its booking limits for the week of March 19, 2007 through March 25, 2007. To make these decisions, BlueSky's revenue managers have collected the daily demand for full-fare tickets over the previous 12 months. Note that these are estimates of actual demand, not the number of tickets sold. These data are available in *BlueSky Single-Leg (A) demand.xls*.¹

Find the optimal (revenue-maximizing) protection level for full-fare seats and the optimal booking limit for economy-class seats for each day of the week beginning on March 19. You should find seven protection levels and seven booking limits, the best possible pair for each day of that week.

You may want to use a simple formula to calculate the protection level (see, for example, the end of §5 of Netessine and Shumsky 2002²) Think carefully, however, about how you use the historical data to estimate $F(Q)$, the distribution of demand for full-fare tickets.

¹This Excel spreadsheet file can be found and downloaded from <http://ite.pubs.informs.org>

²Netessine, S., R. Shumsky. 2002. Introduction to the theory and practice of yield management. *INFORMS Trans. Ed.* 2(1) 34–44. Available online at <http://archive.ite.journal.informs.org/Vol3No1/NetessineShumsky>.

Case B

BlueSky operates Flight 97, a nonstop flight from JFK to Salt Lake City that departs at 9:30 P.M.¹ For this route they fly an Airbus A320 that can carry 146 passengers. On the airplane, all seats are economy-class (there are no business or 1st-class seats), and the marginal cost of each additional passenger in these seats is negligible. BlueSky has constructed a two-tier fare structure: Advance purchase tickets (i.e., nonrefundable tickets purchased at least 14 days in advance) cost \$114 one-way. Full-fare refundable tickets purchased at any time cost \$174.

Throughout this case we will focus on a single Flight 97 on a single day. Assume that the demand distributions for this flight have been carefully estimated. In particular, demand for full-fare tickets is normally distributed with a mean of 92 and a standard deviation of 30. Demand for advance purchase (*low-fare*) tickets is normally distributed with a mean of 80 and a standard deviation of 25. Note that a model with negative demand would not make any sense; therefore, we will censor both distributions at zero. That is, if the demand falls below zero, we assume that the actual demand is equal to zero. The censoring, however, makes little difference to us because the probability that demand is less than zero is small for both high and low-fare demand distributions.

Our goal is to find the optimal (revenue-maximizing) protection level for full-fare seats and the optimal booking limit for economy-class seats for this flight. One method is to use the standard protection limit formula (see, for example, the general formula at the end of §5 of Netessine and Shumsky 2002²). While this formula can be useful, in practice it may need to be modified because the formula is based on assumptions that ignore many real-world phenomena. In this case we will consider two such phenomena: buy-up behavior, in which some supposed low-fare customers are willing to purchase a full-fare ticket if no low-fare ticket is available, and no-show behavior, in which some passengers buy a ticket but do not show up.

Part 1: The General Formula vs. Simulation

The file *BlueSky Single-Leg (B1).xls*³ contains a Crystal Ball (CB) model of the basic revenue management problem. Audit the model to make sure that

you understand its logic; then find the optimal protection level. This may be done by choosing a protection level, running the simulation, recording the expected revenue, and repeating this process until the expected revenue is maximized.⁴

Also use the general formula to find the optimal protection level for full-fare seats (note that you only need the distribution for full-fare tickets to use the formula). Do the answers from the CB model and the general formula agree? Why or why not?

Part 2: Buy-Up Behavior

The protection level formula was derived under the assumption that customers can be categorized into just two groups, those who only buy advance purchase nonrefundable tickets and those who only buy full-fare refundable tickets. It is reasonable to assume that there is a third group: those who prefer a low-fare ticket but are willing to *buy up* to a full-fare ticket if a low-fare ticket is not available (think of a price-sensitive small business owner who must fly to a meeting). This implies that the protection limit itself affects the number of full-fare customers: Whenever we stop selling low-fare tickets, some of the remaining low-fare passengers buy up to full-fare.

Specifically, assume that each low-fare customer has a 30% chance of buying a full-fare ticket if no low-fare ticket is available. To see how a model with buy-up might work, consider the following example.

EXAMPLE. Suppose that we set the protection level to 50 seats, so that the booking limit is $146 - 50 = 96$ seats. Now suppose that 132 passengers arrive whose first choice is to buy a low-fare ticket and that 40 passengers arrive later to buy full-fare tickets (these 40 will only buy full-fare refundable tickets). Given the 132 low-fare passengers, the booking limit of 96 is reached, 96 low-fare tickets are sold, and $132 - 96 = 36$ low-fare passengers who are turned away. Of these, suppose that 12 passengers (30%) are willing to buy up to a full-fare ticket, so that the total demand for full-fare tickets is $40 + 12 = 52$. We cannot sell all 52 full-fare tickets, however, because we have already sold 96 low-fare tickets and $96 + 52 = 148$, which is greater than the 146 seats on the airplane. Therefore, we sell 50 full-fare tickets, and the total revenue is $96 * \$114 + 50 * \$174 = \$19,644$.

Of course, this example fixed many of the quantities that will be random. Recall that we assumed, above, that the number of initial low-fare and full-fare passengers are both normally distributed. In addition, the number of passengers willing to buy up is not

¹The information about Flight 97 in this paragraph is identical to the information in “BlueSky Airlines: Single-Leg Revenue Management (A).” The demand information given in the next paragraph, however, differs from the demand information in the (A) Case.

²Netessine, S., R. Shumsky. 2002. Introduction to the theory and practice of yield management. *INFORMS Trans. Ed.* 2(1) 34–44. Available online at <http://archive.itejournal.informs.org/Vol3No1/NetessineShumsky>.

³This Excel spreadsheet file can be found and downloaded from <http://ite.pubs.informs.org>.

⁴A faster alternative is to use a tool that automates this sensitivity analysis, such as the Sensitivity Toolkit’s CB Sensitivity (Tuck School of Business at Dartmouth. 2009. Sensitivity toolkit. <http://mba.tuck.dartmouth.edu/toolkit/download.html>. Last accessed March 26, 2009.)

always 30% of the number of low-fare passengers shut out of low-fare tickets. In fact, the number of buy-up passengers is distributed according to the binomial distribution: If 36 passengers are turned away, then in CB the function `CB.Binomial(0.3,36)` gives the (random) number of passengers from the 36 that are willing to buy up.⁵ Note that the second value given to `CB.Binomial` must be “1” or larger—`CB.Binomial(0.3,0)` will produce an error—you should use the `if()` function to check.

Copy the file *BlueSky Single-Leg (B1).xls*, name it “*BlueSky (B2).xls*,” and open that model (recall that to avoid confusion it is best to open only one CB model at a time). Then alter “*BlueSky (B2).xls*” to take this buy-up behavior into account. Once the model is complete and correct, find the optimal protection level. Again, this may be done by repeatedly changing the protection level and rerunning the simulation, or by using the CB Sensitivity Toolkit.

How has the optimal protection level changed from Part 1? Does the change make sense?

Part 3: No-Shows

In Part 2 we assumed that all passengers who bought a ticket will show up. In practice, passengers miss flights, and full-fare (refundable) ticket-holders are particularly likely to miss a flight.

Specifically, assume that each passenger who purchases a full-fare ticket has a 92% chance of showing up, and assume that all passengers with low-fare tickets always show up. Each full-fare no-show receives a full refund (with no penalty). Therefore, even if the airline sells tickets for all 146 seats, there is a significant risk that the airplane will fly with empty seats that do not generate revenue.

Therefore, BlueSky overbooks the airplane. Within BlueSky’s information system, this is represented by a virtual capacity that is larger than the actual capacity of 146 seats. For example, BlueSky may find it optimal to pretend that the airplane has 154 seats and therefore will sell more tickets than the actual capacity if there is sufficient demand.

Of course, there is a potential cost: If too many customers show up, the airlines must pay to divert

customers from the flight, e.g., by finding volunteers to take a later flight in exchange for cash or a voucher (these passengers are sometimes called *bumped* passengers). Suppose that the penalty cost of each bumped passenger is \$180.⁶ The airline, then, must find the right balance between too little overbooking and leaving seats empty, and too much overbooking and paying too many of these penalties. To see how a model with buy-up and no-shows might work, consider the following example.

EXAMPLE. Suppose that BlueSky sets the protection level to 50 and that the virtual capacity is 154 seats (8 higher than the actual capacity). Therefore, the booking limit is $154 - 50 = 104$. As in the example in Part 2, suppose that low-fare demand is 132 and (pure) full fare demand is 40. Therefore, 104 low-fare tickets are sold and 28 low-fare passengers are shut out. Suppose that of these 28, 12 passengers decide to buy up (recall that the actual number that buy up is a random number with a binomial distribution). Therefore, the total full-fare demand is $40 + 12 = 52$. Given that the virtual capacity is 154 and that we have already sold 104 tickets, we can only sell 50 full-fare tickets. Therefore, we sell 104 low-fare and 50 full-fare tickets, and the virtual airplane is full.

Now, suppose that of the 50 full-fare tickets sold, 46 passengers show up and pay full fare (although the actual number that shows up is also a random number with a binomial distribution). Therefore, the total number of actual passengers is $104 + 46 = 150$, and we have 4 extra passengers that will not fit into the actual airplane and must be rebooked. Therefore, the total revenue from the flight is $104 * \$114 + 46 * \$174 - 4 * \$180 = \$19,140$.

Copy your completed buy-up model “*BlueSky (B2).xls*,” name it “*BlueSky (B3).xls*” and open that model. Then alter “*BlueSky (B3).xls*” to add no-show behavior to the model. Note that there are now two decision variables: the protection level and the virtual capacity. Use the model to find the optimal combination of virtual capacity and protection level. Here you must adjust both decision variables by hand or use the CB Sensitivity Tool to conduct a two-way sensitivity analysis.

⁵ The binomial distribution is the model used for the classic probability question: How many heads do you see when you flip a coin N times? In the buy-up model here, the coin is not fair (it has a 30% chance of landing on heads and producing buy-up behavior), and “ N ” itself is not known in advance, but is equal to the number of low-fare passengers turned away.

⁶ The \$180 includes the expected cost of giving a seat to the bumped passenger on a later flight. In practice this displacement cost is sometimes close to zero because the airline provides a seat that would have flown empty, anyway.

Case C

This case examines what has happened to many traditional airline revenue management systems as low-cost competition has prompted the removal of fare fences and changed customer behavior.

We begin in a traditional yield management environment in which customers can be reliably segmented into leisure and business customers. The primary fences used to separate the two types are a Saturday-night stay requirement and advance-purchase requirements. Here, the low-fare revenue is only available with a Saturday-night stay requirement.

Consider the following data for one BlueSky flight:

Capacity of airplane	250
Low-fare revenue	\$108
High-fare revenue	\$230
Mean low-fare demand	52
Standard dev. of low-fare demand	18
Mean high-fare demand	120
Standard dev. of high-fare demand	43

You may assume that these demand distributions are good for every day, occasions, e.g., there is no systematic variation for holidays, day of the week, or time of year.

Question 1: Based on these data, find the optimal protection level:

Critical fractile _____

Optimal protection level _____

Now BlueSky is faced with a low-cost competitor in this market that offers 108 seats with no Saturday-night stay requirement. To remain competitive, BlueSky removes its own Saturday-night stay requirement (although it keeps the advance-purchase requirement). As a result, a large proportion (60%) of the business customers who are willing to buy the high-fare ticket would prefer to buy the low-fare ticket.

The spreadsheet *BlueSky Single-Leg (C).xls*¹ allows you to simulate this new environment in which some business customers buy down (the logic is similar to your solution to the buy-up problem in BlueSky Single-Leg (B) Part 2, but is a bit more complicated).

Suppose that BlueSky continues to use the protection level you calculated in Question 1 over the next quarter of the year (92 days). Simulate this by entering the appropriate booking limit and changing “Run ► Run Preferences... ► Number of trials to run” to 92.

Question 2: What is the average daily revenue over the quarter?

Mean daily revenue _____

During the quarter BlueSky observes the number of high-fare ticket sales, and uses these observations to update its forecast of the distribution of high-fare ticket sales for the next quarter. Then, this forecast is used to find a new optimal protection level and booking limit.

Question 3: Given the results of the simulation run for Question 2, find

Mean number of high-fare tickets sold _____

New optimal protection level _____

Now, BlueSky operates for another quarter with this new protection level. At the end of the second quarter, it records its average revenue, observes the new high-fare demand distribution, and calculates another optimal protection level. This is followed by another quarter, and another.

Question 4: Simulate this process and complete Table 1. Note that the first column of the table can be completed using the numbers you calculated, above.

What do you observe? Why?

Finally, suppose that BlueSky uses the simulation itself to find an optimal protection level (remember to reset “Number of trials to run” to something reasonably high).

Question 5: Using the simulation, find the

Optimal protection level _____

Mean daily revenue _____

How do these values compare to the results from Question 4? Why?

¹This Excel spreadsheet file can be found and downloaded from <http://ite.pubs.informs.org>.

Table 1

	Quarter				
	1	2	3	4	5
Protection level during quarter					
Mean daily revenue					
Mean number of high-fare tickets sold					
New optimal protection level (for next quarter)					