

Advanced Information Systems and Business Analytics for Air Transportation, June 2015

Group Case Study: SkyJet¹

Basic instructions

This case consists of two “independent” cases, which are accompanied by the corresponding excel files. These cases require some forecasting analysis, understanding the basics of revenue management, basic simulation skills (in Excel), and use of Excel’s solver.

Due date: end of the week.

Please read the cases completely, and if you have questions/concerns please contact the instructors soon enough so that we can respond accordingly.

You are required to support your work with the relevant (Excel) files. Make sure one can follow your work and name your files (and tabs) properly. For example, if you use a separate file for each part, then you might have “SkyJet Case1A PartI MyName.xls”.

Individual submissions.

Case 1: Single Leg Revenue Management at SkyJet

Case 1A

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

SkyJet has constructed a two-tier fare structure: advance purchase tickets (nonrefundable tickets purchased at least 14 days in advance) cost \$114 one-way. Full-fare refundable tickets purchased at any time cost \$174. Assume that there is heavy demand for advance purchase tickets, so SkyJet could sell out the aircraft with these discount passengers. Therefore, SkyJet’s revenue management system may need to protect a certain number of seats for full-fare tickets.

On March 5, 2007, SkyJet is setting its booking limits for the week of March 19, 2007 – March 25, 2007. To make these decisions, SkyJet’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 the Excel spreadsheet, SkyJet (A) demand.xls.

Find the optimal (profit-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 7 protection levels and 7 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 Section 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.

¹ Adopted from Robert A. Shumsky. 2008. Introductory, Integrative Cases on Airline Revenue Management.

² Netessine, S. and R. Shumsky. 2002. Introduction to the Theory and Practice of Yield Management, INFORMS Transactions on Education, Vol. 2, No. 1.

Case 1B

SkyJet operates Flight 97, a nonstop flight from JFK to Salt Lake City that departs at 9:30 pm.³ For this route they fly an Airbus A320 that can carry 146 passengers. On the airplane, all seats are economy-class seats (there is no business or 1st-class), and the marginal cost of each additional passenger in these seats is negligible. SkyJet has constructed a two-tier fare structure: advance purchase tickets (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 portion of the 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, and 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 (profit-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 Section 5 of Netessine and Shumsky, 2002). While this formula can be useful, in practice it may need to be modified because the formula is based upon 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

Construct a simulation model using Excel (or any other software of your choice)⁴. For any simulation model you want to construct many repetitions for each instance. I would recommend using about 1000-5000 replications (but you can have even more)—the more replications you have in your simulations, the smaller is the “noise” in your decisions.

Find the optimal protection level using the simulation model you have constructed. 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.⁵ If you are using Excel, refine your steps to make sure you are selecting the proper protection level.

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 simulation model and the general formula agree? Why or why not?

If you are using Excel, you may find the following command useful:

RAND() generates random numbers between 0 and 1

INT(x) rounds the number x down to the closest integer.

³ The information about Flight 97 in this paragraph is identical to the information in “Revenue Management at SkyJet (A).” The demand information given in the next paragraph, however, differs from the demand information in the (A) case.

⁴ If you are interested in installing Crystal Ball please refer to http://www.decisioneering.com/crystal_ball/index.html.

⁵ Another alternative is to use a tool that automates this sensitivity analysis, such as the Sensitivity Toolkit’s Crystal Ball Sensitivity.

ROUND(x,y) will round the number x to y digits.

Thus, INT(RAND()*6)+1 in excel will simulate dice shuffling.

NORM.DIST(x,mean,standard deviation, cumulative) returns the normal distribution for the specified mean and standard deviation.

NORM.DIST(probability,mean,standard deviation,) returns Returns the inverse of the normal cumulative distribution for the specified mean and standard deviation

MIN() and MAX() return the smallest and the largest values, respectively, from a set values.

Keep in mind that when you simulate arrival from a distribution, you need to truncate negative values.

Part 2: Buy-Up Behavior

The protection level formula was derived under the assumption that the market can be perfectly segmented, i.e., 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 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. 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 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 the Excel function BINOM.INV will help you find the (random) number of passengers from the 36 that are willing to buy up.⁶ Note that BINOM.INV will ask you to input the number of trials, probability, and the criterion value (which is a number between 0 and 1).

Create a spreadsheet that takes 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 re-running the simulation.

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

⁶ 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.

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, SkyJet overbooks the airplane. Within SkyJet's information system, this is represented by a virtual capacity that is larger than the actual capacity of 146 seats. For example, SkyJet 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 \$180.⁷ 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 this example.

Example: Suppose that SkyJet 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 decide to buy up (recall that the actual number of passengers 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 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 re-booked. Therefore, the total revenue from the flight is $104 * \$114 + 46 * \$174 - 4 * \$180 = \$19,140$.

Reconstruct the Excel file 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 manually.⁸

Comment on the optimal results you obtain.

⁷ The \$180 includes the expected the 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.

⁸ You can also use the Sensitivity Toolkit to conduct a "Two-Way" sensitivity analysis, or use other software as instructed in previous parts.

Case 1C

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 SkyJet 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 – e.g., there is no systematic variation for holidays, day-of-week or time-of-year.

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

Now SkyJet is faced with a low-cost competitor in this market that offers \$108 seats with no Saturday-night stay requirement. To remain competitive, SkyJet 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.

Simulate this new environment in which some business customers buy-down (the logic is similar to your solution to the buy-up problem in SkyJet (1B) part 2, but is a bit more complicated).

Suppose that SkyJet 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 the number of trials to run to 92.

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

During the quarter SkyJet observes the number of high-fare ticket sales, and uses these observations to update its own 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 (i) the mean number of high-fare tickets sold, (ii) the new optimal protection level.

Now, SkyJet 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 the following table. Note that the first column of the table can be completed using the numbers you calculated, above.

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

What do you observe? Why?

Finally, suppose that SkyJet uses the simulation itself to find an optimal protection level (remember to change the number of trials to run to something reasonably high).

Question 5. Using the simulation, find the (i) optimal protection level, (ii) the mean daily revenue.

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

Case 2: Network Revenue Management at SkyJet

Case 2A

SkyJet flies three airplanes between Houston and three cities, Chicago, Miami and Phoenix. These three cities are the spokes connected by the Houston hub. A few times each day the three airplanes fly from the spoke cities to Houston. They arrive simultaneously at Houston, connecting passengers change aircraft during a 45-minute layover, and the three airplanes depart for the spokes. One set of six flights (3 inbound to Houston and 3 outbound) is called a bank. Each bank can serve passengers flying on 12 different routes: 3 inbound direct routes (Chicago-Houston or C-H, M-H, and P-H), 3 outbound direct routes (H-C, H-M, and H-P), and 6 routes requiring 2 flights each (C-M, C-P, M-C, M-P, P-C, and P-M).

SkyJet charges a single fee for a one-way coach-class ticket on each passenger route. The table below shows the prices charged by SkyJet. The marginal cost of flying a passenger on each route is virtually zero.

		Destination			
		Houston	Chicago	Miami	Phoenix
Origin	Houston	-	\$197	\$110	\$125
	Chicago	\$190	-	\$282	\$195
	Miami	\$108	\$292	-	\$238
	Phoenix	\$110	\$192	\$230	-

Each of the three airplanes currently has 240 coach seats. The next table shows demand for the routes in a bank, and we can see that passenger demand exceeds airplane capacity on every flight. For example, on the flight from Chicago to the Houston hub (C-H), the total demand is the sum of demands for three passenger routes, C-H, C-M, and C-P, totaling $130 + 98 + 88 = 316$ passengers (this is the sum of the second row of the table). Because only 240 passengers can travel on the C-H flight, at least 76 passengers must be turned away.

		Destination				Total Demand for flight to hub
		Houston	Chicago	Miami	Phoenix	
Origin	Houston	0	123	80	110	-
	Chicago	130	0	98	88	316 (C-H)
	Miami	72	105	0	68	245 (M-H)
	Phoenix	115	90	66	0	271 (P-H)
Total demand for flight to spoke		-	318 (H-C)	244 (H-M)	266 (H-P)	

When the total demand for a particular flight may be larger than the available capacity, an airline can decide whether to accept or reject an offer to buy a ticket for a particular route. Controlling sales in this way to maximize revenue is called revenue management. For example, SkyJet may decide that it is optimal to sell large numbers of tickets for the C-H and C-M routes, but might severely restrict the number of C-P tickets sold. Given the data above, SkyJet might sell tickets for 130 C-H routes, 98 C-M routes, and only 12 C-P routes, thus filling all 240 seats on the C-H flight.

Build a model to determine the number of tickets to sell for each route. The objective is to maximize revenues over a single bank. Note that the tables above are available in the spreadsheet SkyJetNRM_data.xls.

Case 2B

Frequently other airlines offer to pay SkyJet to fly passengers on its planes (this is one purpose of an alliance among airlines). For example, Air France often flies passengers from Paris to Miami. A few of Air France's passengers then go on from Miami to Houston, while other passengers are headed to Chicago.

- 1) Air France asks SkyJet to reserve 5 seats on SkyJet's flights from Miami to Houston, to be used for connecting Air France passengers. Air France offers to pay \$104 per seat. Should SkyJet accept the offer?
- 2) In addition to the offer in (1), Air France offers SkyJet \$285 per seat to reserve 10 seats for passengers traveling from Miami to Chicago. Should SkyJet accept this offer?

Case 2C

SkyJet is replacing the three aircraft that fly in and out of its Houston hub, and it plans to purchase the three new aircraft from the Airbus A330/A340 family. SkyJet has already decided to configure the three aircraft with only coach seats and no first-class cabin, but the airline has not yet decided on the size of the aircraft. The A330/A340 family comes in a wide range of sizes, from 240 to 380 coach seats.

To decide on an aircraft size, SkyJet must consider both the cost and revenue implications. On the cost side, a larger aircraft is more expensive to purchase and more costly to operate. The purchasing terms and performance data show that the total cost of one flight from, say, Houston to Miami, includes a fixed cost of \$12,000 and an additional cost of \$40 per seat. These numbers are calculated from all costs associated with the flight, including fuel, labor, and maintenance. The cost parameters for the other five flights in each bank in and out of Houston are nearly identical because all six flights are approximately the same length.

- 1) Assume that SkyJet purchases three identical aircraft. How many coach seats should SkyJet order for the three new aircraft?
- 2) Now suppose that the three aircraft can be different sizes, between 240 and 380 coach seats.
 - a. How do you think the three aircraft should be allocated among the six routes? In other words, should the same aircraft always fly the same routes? Why or why not? (Hint: you do not need solver to answer this question).
 - b. How many coach seats should SkyJet order for each of the three new aircraft?
- 3) Airbus is offering SkyJet a one-time, \$5 million discount if it will order three identical aircraft, because it is cheaper to manufacture three identical planes. Should SkyJet take the discount? In deciding this, you may assume that SkyJet operates 3 banks per weekday through Houston, and that the revenues and demands for every bank on every weekday are equal to the demands in the two tables provided earlier in Case 2A.