COLUMBIA BUSINESS SCHOOL

Lecture 10

- Introduction to Yield Management
- Ski Jacket Case
  - Profit calculation
  - Spreadsheet simulation
  - Analysis of results
- Summary and Preparation for next class

Introduction to Yield Management

Yield management is the process of allocating different types of capacity to different customers at different prices in order to maximize revenue.

Examples include

- Hotel industry
  - How many rooms to allocate for each market segment (over time)?
- Airline industry
  - How many seats to allocate for each fare class?
  - How to set fare classes and restrictions?
- Other industries
  - Cruise lines
  - Railroads
  - Car and truck rental
  - Theater and concert ticketing
  - Advertising networks

In-class yield management examples:

- Retail Pricing: Quantity fixed, price over time to be decided (lecture 8)
- Ski Jacket: Price fixed, quantity to be decided
- AA Case

Ski Jacket Production

- **Problem**: Consider a retailer who must determine how many ski jackets to produce (and stock) at the beginning of the selling season given an uncertain level of demand? Because of production batch size requirements, the quantity produced must be a multiple of 2,000.

- Quantity to produce, *Q* — Random demand, *D* — Profit (*Q*,*D*)

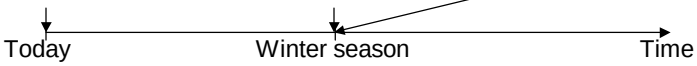

Variable production cost per unit (*C*):	$80
Selling price per unit (*S*):	$100
Salvage value per unit (*V*):	$30
Fixed production cost (*F*):	$100,000

S is the amount the manufacturer receives from the retailer. All full-price jackets unsold at the end of the season can be salvaged for *V* per unit. *F* is the fixed cost of plant and equipment. Let *Q* denote the quantity of ski jackets to produce (decision variable).

Exact demand for ski jackets next winter is unknown. What can we say about the random demand *D*?

Estimating Demand

- How do we estimate demand? One method (sometimes called the "Delphi" method) is to have a number of managers each estimate the product's demand based on their own judgment. Suppose twelve managers give the following demand estimates:

14,000	16,000
13,000	8,000
14,000	5,000
14,000	11,000
15,500	8,000
10,500	15,000

These forecasts have mean m = 12,000 and a standard deviation s = 3,500 (3,497 to be precise). The Delphi method uses numbers as the mean and standard deviation.

- What distribution for demand should we use in the simulation?

Without further information about how well historical forecasts have done, we'll assume that demand is normally distributed. That is, we'll assume that

$D \sim N(m = 12{,}000,\ s = 3{,}500)$. (~ = "is distributed as")

Profit Calculation

- Suppose you decide to produce Q = 12,000. Demand is random, but suppose it turns out to be D = 13,000. What is the profit for the season?

  Profit = Revenue - Variable cost - Fixed cost
  = 100(12,000) - 80(12,000) - 100,000
  = $140,000

  Note: Even though D = 13,000, only 12,000 jackets were sold.

- Suppose Q = 12,000 as before. Now suppose demand turns out to be D = 11,000.

  What is the profit for the season?

  Profit = Rev + Salvage val - Var cost - Fixed cost
  = 100(11,000) + 30(1,000) - 80(12,000) - 100,000
  = $70,000

- What is a general formula?

Profit Formula

- The general formula is

  Profit = Revenue + Salvage value - Var cost - Fixed cost.

- To compute revenue, there are two cases to consider: $D < Q$ and $D \geq Q$. If $D < Q$ the revenue is $S * D$; if $D \geq Q$ the revenue is $S * Q$. This can be incorporated in a spreadsheet with the single formula

  Revenue = S * *IF*($D < Q$, D, Q) .

  If $D < Q$ the salvage value is $V * (Q - D)$; if $D \geq Q$ the salvage value is 0. This can be done in a spreadsheet with the formula

  Salvage value = V * *IF*($D < Q$, $Q - D$, 0) .

  The variable cost is $C * Q$ and the fixed cost is F.

- The ski jacket production problem can be formulated as an optimization model:

  $$\max_{Q} \text{E[Profit}(Q,D)] ,$$

  i.e., find the Q among 6000, 8000, ..., 14000, etc., which maximizes expected profit.

Ski Jacket Simulation Spreadsheet

	A	B	C	D	E	F	G
1	SKI.XLS			**Ski Jacket Simulation**			
2							
3	Variable Cost (C)		80		*Mean demand*		12,000
4	Selling Price (S)		100		*Standard deviation*		3,500
5	Salvage value (V)		30		*Demand (D)*		**12,000**
6	Fixed Cost (F)		100,000				
7							
8			*Salvage*	*Variable*			
9	***Quantity (Q)***	***Revenue***	***Value***	***Cost***	***Profit***		
10	6,000	600,000	0	480,000	**20,000**		
11	8,000	800,000	0	640,000	**60,000**		
12	10,000	1,000,000	0	800,000	**100,000**		
13	12,000	1,200,000	0	960,000	**140,000**		
14	14,000	1,200,000	60,000	1,120,000	**40,000**		
15							

Assumption cell

Forecast cells

- Cell B10: =$C$4*IF($G$5 < A10, $G$5, A10)
  Cell C10: =$C$5*IF($G$5 < A10, A10 - $G$5, 0)
  Cell D10: =$C$3*A10
  Cell E10: =B10 + C10 - D10 - $C$6
- To compare the five production quantities with a Crystal Ball simulation:
  (1) Define assumption cell (cell G5 is normal with mean =G3 and standard deviation =G4)
  (2) Define forecast cells (E10 through E14)*
  (3) Set simulation run preferences (500 trials and seed 123)
  (4) Run the simulation

  *See next slide for details of step 2.

Defining Multiple Forecasts in Crystal Ball

- In step 2 (define forecasts) we want to define cells E10:E14 as forecast cells. A quick way to do this is:
  (2a) Define forecast cell E10 (move cursor to cell E10 and then click on the Crystal Ball "Define Forecast" icon)
  (2b) With the cursor on cell E10, click on the Crystal Ball "Copy Data" icon (not the regular Excel copy)
  (2c) Highlight cells E11:E14. Then click on the Crystal Ball "Paste Data" icon

  This procedure defines cells E10:E14 as forecast cells.

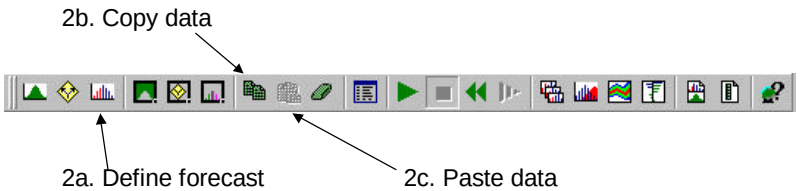

- After setting the run preferences, click on the "Start Simulation" icon.

Ski Jacket Simulation Results - 500 Trials, Seed 123

- To compare the results of the five forecasts simultaneously, use the "Extract Data" icon from the Crystal Ball toolbar (choose "Statistics" as the Type of Data) to get summary results:

Quantity (Q)	Average Profit	Std. Error
6,000	15,945	1,202
8,000	46,228	2,351
10,000	58,936	4,105
12,000	45,729	6,260
14,000	3,645	8,321

- According to the simulation, the optimal quantity to produce is $Q^* = 10{,}000$.

Confidence Interval Computation

- Give a 95% confidence interval for the average profit corresponding to Q = 10,000.The mean profit is 58,936 with a standard error of 4,105.
- General formula for a (1- a) confidence interval:

$$\overline{X} \pm z_{a/2}\, SE$$

where the standard error = $SE = s/\sqrt{n}$.

- For a 95% confidence interval: a = 0.05, and $z_{a/2}$ = 1.96. In general, $z_{a/2}$ can be computed in a spreadsheet using `NORMSINV(1- a/2)`, e.g., `NORMSINV(0.975) = 1.959961.`
- For our example, the 95% confidence interval based on n = 500 trials is:

$$58{,}936 \pm 1.96(4{,}105),$$

i.e., the 95% confidence interval is

$$[50890,\ 66982] .$$

- To obtain a narrower confidence interval, run more simulation trials.

Results From 10,000 Simulation Trials

Quantity (Q)	500 trials Avg. Profit	10,000 trials Avg. Profit	Std Dev
6,000	15,945	15,744	27,606
8,000	46,228	44,789	54,799
10,000	58,936	57,056	94,940
12,000	45,729	42,743	142,503
14,000	3,645	-2,684	186,741

- With many more simulation trials, the estimates of the mean profit change slightly, but the optimal quantity to produce still appears to be

  $Q^* = 10{,}000$.

  With 10,000 trials the 95% confidence interval for the average profit at $Q = 10{,}000$ is:

  $57{,}056 \pm 1.96(949)$,

  i.e., the 95% confidence interval based on n = 10,000 trials is

  [55196, 58916] .

Are we sure that Q = 10,000 is better than Q = 12,000?

- A more direct and accurate way to compare two strategies, e.g., compare the profit for $Q = 10{,}000$ and $Q = 12{,}000$, is as follows:
  - Define a new forecast cell in the spreadsheet which is the difference in the profits under the two strategies.
  - Run the simulation.
  - Check that the average difference in profit is positive and compute a confidence interval to show that the difference is statistically significant.

How do you know when a difference is statistically significant?

Consider this new forecast cell (with the difference in profits in it). It has a mean and a standard error associated with it (after the simulation is run). If a 95% confidence interval around the mean does not contain 0, then the difference is statistically significant at the 5% significance level. If 0 is contained in the confidence interval then the difference is not significant at the 5% level.

What about risk?

- Maximizing average profit does not consider the *distribution* of profits that will actually occur. For example, the table of average profit vs. quantity shows that Q = 12,000 and Q = 8,000 have the same average profit of about $44,000.
- Would management be indifferent between these two production quantities?
- Probably not.

Average Profit, Std Dev vs. Quantity *Q*

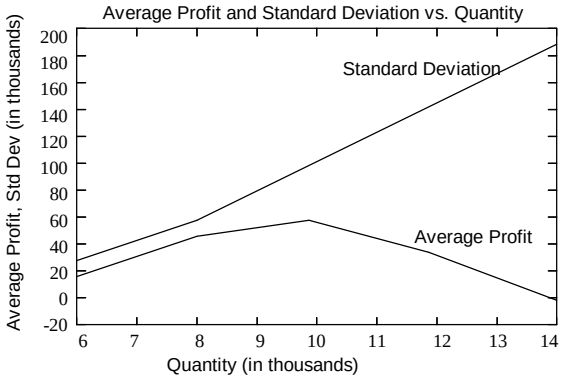

- The standard deviation of profit for Q = 8,000 is approximately $55,000. The standard deviation of profit for Q = 12,000 is approximately $143,000. Even though both quantities have the same expected profit, the risk is nearly *3 times greater* for Q = 12,000 compared to Q = 8,000.

Efficient Frontier

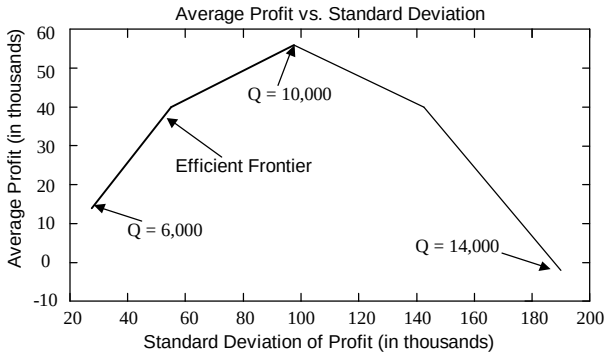

- Production quantities $Q \le 10{,}000$ are efficient; production quantities $Q > 10{,}000$ are inefficient. Management might prefer $Q < 10{,}000$ because of the reduced risk.
- Average profit and standard deviation do not tell the whole story though. A frequency histogram shows the entire distribution of profit for a fixed Q. Since demand is normally distributed, is profit also normally distributed?

Distribution of Profit ($Q = m = 12{,}000$)

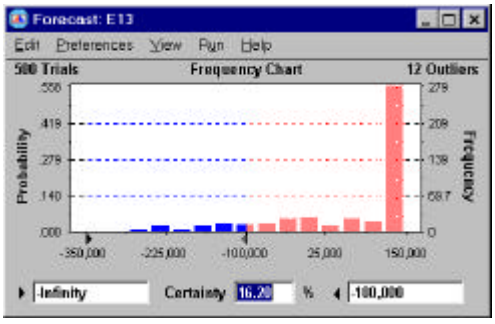

- The histogram shows that the distribution of profit is highly skewed to the left. Losses greater than $100,000 happen about 16% of the time. (In the cell E13 frequency chart, drag the right arrow to -100,000 and read the probability in the window titled "Certainty.")
- Even though demand is normally distributed, profit is *not* normally distributed (as one can plainly see in the histogram above). With $Q = m = 12{,}000$, the most likely outcome is a profit of 140,000. This occurs whenever $D \ge m$, which happens 50% of the time.

Distribution of Profit ($Q^* = 10{,}000$)

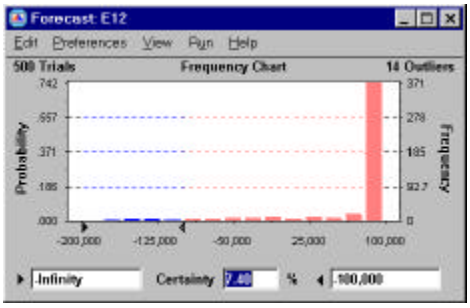

- The histogram shows that the distribution of profit is also highly skewed to the left. Losses greater than $100,000 happen about 7% of the time. (In the E12 frequency chart, drag the right arrow to -100,000 and read the probability in the window titled "Certainty.")
- Even though demand is normally distributed, profit is not normally distributed (as one can plainly see in the histogram above). With $Q^* = 10{,}000$, the most likely outcome is a profit of $100,000. This occurs whenever $D \geq 10{,}000$, which happens about 75% of the time.

Modifying Histograms in Crystal Ball

- The histograms on the previous two slides were automatically generated when the simulation was run. However, some adjustments were made to make them more presentable. Specifically in the E12 and E13 forecast windows:
  - *To change the number of bins:* Choose "Preferences" and then choose "Chart." Change the number of groups (bins) to 15.
  - *To change the x-axis range:* Choose "Preferences" and then choose "Display Range." In the section "Using Fixed End-Points" change the "Min" to -350,000 and the "Max" to 150,000.
  - *To print the histogram:* Choose "Edit" and then "Copy." Return to the Excel spreadsheet (or word processor) and select "Paste."

Remember: It is a good idea to have only one simulation spreadsheet open at a time when a Crystal Ball simulation is run. This is because Crystal Ball will simultaneously simulate all open spreadsheets. This will slow down the computer and also means that the random number seed will not work as expected.

The Intuition Behind Our Result

- We observe in this example that $Q^* < m$, i.e., the optimal order quantity is less than the average demand. Why is this? Why might we have expected this in our answer?
- Furthermore, this result does *not* appear to be due to simulation error.

- Intuition:

  Compare the marginal profit of stocking one more ski jacket with the marginal cost. If the extra ski jacket is sold, the added profit is $S - C = \$20$. If the extra jacket remains unsold, the added profit is $V - C = -\$50$, i.e., a \$50 cost. The *cost* of not selling one more jacket is *much greater* than the *benefit* of selling one more jacket. Hence the optimal Q will be less than m.

An Analytical Solution

- The ski jacket optimization model is:

  max E[Profit(Q, D)]

  s.t. $Q \geq 0$

  A formula which finds Q^* for all costs and demand distributions:

$$Q^* = F^{-1}\left(\frac{C_u}{C_u + C_0}\right)$$

where $C_u = S - C$, $C_o = C - V$, and where $F^{-1}(x)$ is the inverse of $F(x) = P(D \leq x)$.

Analytical Solution

- We have C_u = 100 - 80 = 20, C_o = 80 - 30 = 50, and $C_u/(C_u + C_o)$ = 20/(20+50) = 2/7. Since demand is distributed N(m = 12,000, s = 3,500), $F^{-1}(2/7)$ = `NORMINV(2/7,12000,3500)` = 10019.18, i.e., $P(D \leq 10019.18) = 2/7$.
- Optimal average profit: The formula for Q^* gives the optimal order quantity, but not the optimal average profit.
- The optimal (continuous) quantity is Q^* = 10019.18. Why bother simulating?
- The formula for Q* gives the optimal value if Q is allowed to be continuous, but does not handle batch size requirements as easily. (In this case, the optimal Q turned out to be very close to an allowable batch size of 10,000.)

Other Versions of the "Newsvendor" Model

- The main characteristics of the ski jacket case are:
  - A decision needs to be made with imperfect information
  - The item is "perishable" in that it cannot be restocked after better information about demand becomes known
  - There are costs for overstocking (or overestimating demand) and costs for under stocking (or underestimating demand)

Other applications:

- Retailing
- PC manufacturing
- Number of H&H bagels to stock in the Uris deli each morning

For next class

- Read “Identifying, Measuring, and Hedging Currency Risk at Merck,” “Merck’s 1995 Annual Report,” and “Managing Risk” in the readings book.
- At this point we have covered enough material on simulation for you to begin the “Yield Management at American Airlines” case. The case is due March 1st (at the start of the final exam).