B6015 – Decision Models
Review Session 5

The primary aims of this review session are:

- To practice with constructing more complicated simulation models.
- To become more familiar with the features of Risk Solver Platform in the context of these more advanced examples.
- To review the construction of confidence intervals around the results of a simulation.

■ □ ■

Question 1 (GAP's Return Policy)

GAP has (or had, when I wrote this question!) a rather innovative return policy. If you buy a product full-price from GAP and it is subsequently put on sale, you can turn up with the receipt and get the difference refunded. This policy seems counterproductive – how could it possibly increase GAP's overall income? In this question, we will construct a very simple model to explore the benefits of this policy.

Assume that GAP only sells two products – call them A and B. GAP knows the demand for each product is normally distributed, but does not know what the *actual* demand will eventually be (or, indeed, whether one product will be more popular than the other). Table 1 summarizes the data about these products. Assume that supply of each product is basically unlimited, so that all demand is satisfied.

	Mean	Standard deviation
Company A	2,000	300
Company B	5,000	400

Table 1: Demand for products A and B is normally distributed. This table summarizes the mean and standard deviation of demand at a particular store.

GAP follows the following pricing strategy

- Initially, during a first phase which we will call Phase 1, it sells both product for $50
- After Phase 1, it looks at quantities sold and reevaluates its pricing policy. Any product which has sold fewer than 500

units has its price reduced to \$45, as part of a sale. Any product which has sold more than 500 units remains at \$50.

- These prices stay in place during Phase 2. If the return policy is in action, it is assumed that *every* customer who bought the product in phase 1 will ask for the refund in phase 2.

Assume that GAP's customers behave as follows

- If the return policy described above is in action, then 90% of customers will buy products in Phase 1 (knowing they might get a refund) and 10% will buy products in Phase 2. So for example, if total demand was 10 000 units, 9 000 would be bought in Phase 1 and 1 000 in Phase 2.
- If the return policy described above is not in action, 40% of customers will buy products in Phase 1 while 60% will buy products in Phase 2 (in other words, they choose to wait for sales to begin before making their purchases).

For the sake of this question, assume that fractional numbers of customers are possible.

Part A

Use Monte Carlo simulation to determine GAP's expected revenue in both scenarios (using the refund policy and not using the refund policy), and to estimate the expected difference in revenue. Calculate 95% confidence intervals on these figures.

Can you say with 95% confidence that one policy is better than the other? If not, estimate how many trials you might need to be able to tell, run this number of trials and come up with an answer.

Solution

As usual, let us begin by identifying the sources of randomness in our problem (ie: our *assumption* variables) and their distributions. Here, there are two basic underlying sources of uncertainty

Demand for product A : This is normally distributed, with mean 2,000 and standard deviation 300. We will denote this by x_A. Thus,

$$x_A \sim N(2000, 300^2)$$

Demand for product B : This is also normally distributed, with mean 2,000 and standard deviation 400. We will denote this by x_B, and

$$x_B \sim N(5000, 400^2)$$

Having established what our assumption variables are, let's now move on to our *result* variables. What do we actually care about in this case? Answer: we care about the *revenues* made under each scheme, and about the *difference* in revenues under the two schemes. Let's find out how we can work out each one of those using our assumption variables

Revenue without refund scheme : We'll call this revenue p_1. Clearly, since there are no refunds, it'll simply be equal to

$$p_1 = \text{Phase 1 revenue} + \text{Phase 2 revenue}$$

Phase 1 is pretty easy to figure out. According to the question, demand is always met, and 40% of customers will buy in phase 1 if the refund scheme is not available. Thus, phase 1 sales are simply equal to $0.4x_A$ and $0.4x_B$ for products A and B respectively. The phase 1 price is simply equal to $50. Thus,

$$\text{Phase 1 revenue} = 0.4x_A \times 50 + 0.4x_B \times 50 = 0.4 \times 50(x_A + x_B)$$

Phase 2 is more complicated, because now, the price of each product depends on whether more or fewer than 500 product were sold in Phase 1. If over 500 items of a given product were sold, it is not discounted. If fewer than 500 items were sold, it is discounted. Thus, for example

$$\text{Price of A in phase 2} = \text{IF}(0.4x_A \leq 500, 45, 50)$$

A similar formula applies for product B. Furthermore, we know that 60% of demand is realized in phase 2. Thus

$$\begin{aligned} \text{Phase 2 revenue} &= 0.6x_A \times \text{IF}(0.4x_A \leq 500, 45, 50) \\ &\quad + 0.6x_B \times \text{IF}(0.4x_B \leq 500, 45, 50) \end{aligned}$$

As such, the total profit in this situation is

$$\begin{aligned} p_1 &= 0.4 \times 50(x_A + x_B) \\ &\quad + 0.6x_A \times \text{IF}(0.4x_A \leq 500, 45, 50) \\ &\quad + 0.6x_B \times \text{IF}(0.4x_B \leq 500, 45, 50) \end{aligned}$$

Revenue with refund scheme : We'll call this revenue p_2. There are now refunds, so it'll be equal to

$$p_2 = \text{Phase 1 revenue} + \text{Phase 2 revenue} - \text{Refunds}$$

Once again, phase 1 is simple. This time, 90% of demand happens in phase 1, and the price is still $50. Thus

$$\text{Phase 1 revenue} = 0.9x_A \times 50 + 0.9x_B \times 50 = 0.9 \times 50(x_A + x_B)$$

Let's first deal with phase 2 revenues, before refunds. The situation will be similar to that in phase 1 – 10% of demand will happen in phase 2, and the price will depend on whether a discount is applied. Thus, for example

$$\begin{aligned}\text{Phase 2 revenue} &= 0.1x_A \times \text{IF}(0.9x_A \leq 500, 45, 50) \\ &\quad +0.1x_B \times \text{IF}(0.9x_B \leq 500, 45, 50)\end{aligned}$$

Finally, we need to deal with the refunds. These are actually easier than you'd expect. All you need to do is multiply the number of refunds (which in this case, is simply equal to the number of sales in phase 1) by the price difference in price in phase 1 and phase 2. In other words

$$\begin{aligned}\text{Refunds} &= 0.9x_A \{50 - \text{IF}(0.9x_A \leq 500, 45, 50)\} \\ &\quad +0.9x_B \{50 - \text{IF}(0.9x_B \leq 500, 45, 50)\}\end{aligned}$$

And finally

$$\begin{aligned}p_2 &= 0.9 \times 50(x_A + x_B) \\ &\quad +0.1x_A \times \text{IF}(0.9x_A \leq 500, 45, 50) \\ &\quad +0.1x_B \times \text{IF}(0.9x_B \leq 500, 45, 50) \\ &\quad -0.9x_A \{50 - \text{IF}(0.9x_A \leq 500, 45, 50)\} \\ &\quad -0.9x_B \{50 - \text{IF}(0.9x_B \leq 500, 45, 50)\}\end{aligned}$$

Difference in revenues : This is trivial – the difference in values is simply given by

$$D = p_2 - p_1$$

Running this simulation with 500 trials and a seed of 123, we find the mean of our result variables are

- $\mathbb{E}(p_1)$ =\$350, 015.69.
- $\mathbb{E}(p_2)$ =\$350,036.59.
- $\mathbb{E}(D)$ =\$20.90.

It seems, therefore, that the refund scheme *does* lead to higher profits, but it's very close – the refund scheme only outperforms the no-refund scheme by \$20.90. How can we be sure, therefore, that this difference isn't just the result of statistical variation? We need to calculate a 95% confidence interval for this difference.

Consider that the standard deviation of D, as given by Crystal Ball, is 269.62. Thus, since we have 500 simulations as discussed in review session last week, the 95% confidence interval around the mean difference is

$$20.90 \pm 1.96 \times \frac{269.62}{\sqrt{500}}$$

This gives a confidence interval of

$$-2.73 \text{ to } 44.53$$

Notice that 0 is unfortunately inside the confidence interval. This means that we can't be sure the difference is actually positive – it is still possible (with greater than 5% possibility) that the no-refund scheme actually *outperforms* the refund scheme!

We must therefore increase our number of simulations to get a more accurate result. In particular, we want the half-width of our confidence interval to be *no greater* than $20.90 (that way, the lower end of our confidence interval doesn't cross 0). We can find the number, n, of simulations that are required to do this by solving the following equation

$$1.96 \times \frac{269.62}{\sqrt{n}} = 20.90$$

Solving, we find that this gives

$$n = 639.3$$

To be on the safe side, let's run 1,000 simulations. We then find that the confidence interval becomes

$$4.17 \text{ to } 37.47$$

0 is no longer in the confidence interval, so we are now *sure* the difference is positive. In other words, the refund policy outperforms the no-refund policy.

Part B

GAP is also interested in the probability it will end up having to discount *either* product. Find this probability in each scenario, and use this to explain why the refund scheme might actually increase GAP's overall profits.

Solution

We are now interested in the *probability* that a discount is applied to either product. A general strategy for finding probabilities is to define a new result variable (call it P) given by

$$P = \text{IF}(\text{OR}(\text{Phase 1 A price} \neq \text{Phase 2 A price}, \text{Phase 1 B price} \neq \text{Phase 2 B price}), 1, 0)$$

The value of the variable P will end up being 1 if either product is ever discounted and 0 otherwise.

Finding the *expected value* of this variable P is then equivalent of finding the probability that it is equal to 1.

Running this model is 1,000 simulations, we find that the probability of discount is 0.006 without the refund policy and 0 with the refund policy. This helps us understand the reason the refund policy might actually be *helpful*. Because it means that a greater numbers of customers will buy the product early, it allows GAP to preemptively know that the product really doesn't need to be discounted at all. Thus, even though the refund could potential cause GAP to lose money, it happens less often.

■ □ ■

Question 2 (Renovations to a Holiday Home)

Your British friend, John Doe, is planning extensive renovations to his holiday home in Texas, three years from now. He anticipates the renovation will cost $7,500. John has gone through his budget and finds that he can invest $200 per month for the next three years. John has opened accounts at two mutual funds - the first follows an investment strategy designed to match the return of the FTSE 100. The second fund invests in short-term Treasury Bills. Both funds have very low fees. John is worried about not having enough to cover renovation, and is therefore primarily interested in minimizing his risk of falling short of the required $7,500.

John has decided to follow a strategy in which, each month, he contributes the same fixed fraction of the $200 to each fund. He gets advice from two sources

- John's financial advisor suggested that he invest 20% in the FTSE 100 fund at 80% in T-Bills, because the latter are backed by the United States government and are very safe. If you follow this allocation, he said, your average return will be lower, but at least you'll have enough to reach your target in three years.
- John's sister, an experienced investor herself, said just the opposite - she advised John to invest 80% of the $200 in the FTSE 100 fund each month, and the other 20% in the T-Bill fund. Even though stock returns are risky investments in the short run, the risk would be fairly minimal over the longer three-year period.

Not knowing who to believe, John has for to you for help. After some investigation, you discover that the historical average monthly return on the FTSE 100 fund is 0.95% (with standard deviation 4.2%), and you expect the returns to be roughly similar in the years ahead. Moreover, these monthly returns appear approximately to have a normal distribution. T-Bill rates also vary from month to month, but they are much more stable. You estimate the monthly T-Bill returns over the next three years will be in the range of 0.4% to 0.7%, with all values in this range being equally likely.

Note that when money is invested in a fund, returns in a given month are *reinvested* in the *same fund* the next month. Thus, John is only able to control the percentage of the \$200 he adds in each month – he can't reshuffle moneys between funds.

Based on these assumptions, develop a spreadsheet model to simulate the two suggested investment strategies over the three-year period (36 months). Based on your simulation results, which of the two strategies would you recommend. Why?

Solution

As usual, our first step must be to identify our uncertain assumption variables. In this case, what is our source of uncertainty? It is the return on the FTSE 100 fund and the T-Bill fund. However, the return is uncertain *each month*. As such, our problem will have $36 \times 2 = 72$ assumption variables. We will call then $f_1, \ldots, f_36$ (for the FTSE returns) and $t_1, \cdots, t_36$ (for the T-Bill returns). The distributions of these variables will be the same each month – namely

$$f_i \sim N(0.95, 4.2^2)$$

$$t_i \sim \text{Uniform}[0.4, 0.7]$$

Next, we must consider our *result* variables. In this case, what we are interested in is the final amount of money that can be obtained under each investment strategy. We'll work the details for the first investment strategy, and the steps are similar for the second one.

Let's first consider the first month. Under the first strategy, we invest 20% of the \$200 (ie: \$40) in the FTSE 100, and \$160 in the T-Bills. Let's denote the value of our FTSE 100 and T-Bill holdings be at the end of the month by F_1 and T_1. It is clear that

$$F_1 = 40 \times (1 + f_1)$$

$$T_1 = 160 \times (1 + t_1)$$

What about month 2? We now *leave* the amounts F_1 and T_1 in our respective funds, and we *add* an extra \$40 and \$160 in the FTSE 100 and T-Bill accounts respectively. Thus, at the end of month 2, our holdings are

$$F_2 = (F_1 + 40)(1 + f_2)$$

$$T_2 = (T_1 + 160)(1 + t_2)$$

Similarly, for month 3

$$F_3 = (F_2 + 40)(1 + f_3)$$

$$T_3 = (T_2 + 160)(1 + t_3)$$

And so on. Our results cell is then the sum of our holdings at the end of month 32 – namely, $F_32 + T_32$.

Running this simulation in Risk Solver Platform, we find that the expected final value of each of the strategies is

- $8,109 for strategy 1.
- $8,492 for strategy 2.

It is clear, as expected, that strategy 2 results in higher returns than strategy 1. However, let's now consider the probability of an outcome that is less than $7,500

- In strategy 1, the chance this will happen is 0.2.
- In strategy 2, the chance this will happen is 19.1.

It is therefore clear that in this case, since John Doe's objective is to make the $7,500 mark, strategy 1 should be chosen.

■ □ ■