# Lecture 16: Risk Analysis II

In the last lecture, we discussed how to deal with the uncertainty that must always be present in predicting the future. Here, we take into account the fact that we may be able to estimate the relative probabilities of different futures, and show how these probability data may be incorporated in the planning process.

There's no point in doing this unless we expect to be able to get good probability data to work on. How is this going to happen?

There are two ways of estimating a probability. The *frequentist* method is to repeat a given situtation a large number of times, and keep count of how often each possible outcome occurs. The probability of each outcome is then the number of times it has occurred divided by the total number of trials. This method gives an objective estimate, and can be refined to any degree of precision by increasing the number of trials. However, cases where we can repeat a situation identically many times over are rather rare. One such example might be an oil company sinking exploratory holes; the company has a large database of different geologies, and can use this to come up with a reasonably good estimate of the likelihood of success with a new hole in known terrain.

A second example might be a communications conglomerate looking at the chances that a given launcher will take their satellite into orbit, rather than blowing it up on the launch-pad or losing it in space. Statistical data are available on the success rates of the different space agencies, and might be used to calculate the advantages of spending more money to get a launch on the reliable Ariane, rather than a cheaper Chinese rocket.

The other definition of probability is the *intuitionist* definition: ``The probability of an event is a numerical measure of the expectation a rational person would have that that event will occur." This definition is less objective. It leaves unanswered the question of where we are going to find a rational person, and how that person generates their estimate. On the other hand, it can be applied to a much broader range of situations.

For example, if you are asked to estimate the rate of inflation in Canada over the next ten years, you cannot give a precise answer based on historical data, because the future may not resemble the past, and we don't know all the root causes of inflation. On the other hand, you can say with some confidence that the inflation rate will probably be positive; that it is most likely to be between 2% and 10%; that it is quite unlikely to be over 20%; and that it is almost certain it will not be over 100%.

We conclude that there are at least some situtations in which we can generate plausible estimates of probabilities. We must next review some simple concepts from statistics and probability.

If there are only a few possible outcomes for an event -- for example, a satellite launch either succeeds or fails -- the situation can be described by giving the probability of each outcome. If, on the other hand, the outcomes form a continuous range -- for example, an oil well can produce any yield from zero to a thousand barrels a day -- the situation must be described by the *probability density function* for the variable in question. This function is defined by the rule

**P(x in x+ delta x) = f(x) * delta x**

A probability distribution can be [imperfectly] characterised by its *mean* and its *variance*; the mean is the average value, and the variance is a measure of the spread about that value. Many different probability distributions may have the same values for mean and variance.

## The Basic Procedure

Our approach will be to use available information about the probabilities of different parameter values to calculate the probability distribution for present value of each strategy. This may be more or less difficult, depending on how many different possibilities we are considering. In the simplest case, we have a small number of possible outcomes, and the probability of each outcome and its associated present value may be calculated explicitly.

It may also happen that there are a large number of possible outcomes -- for example, inflation rate may take any value between 0 and 50% -- making it impractical to evaluate each separately. In this case, we may have to settle for calculating the expected value and the variance of present worth. The expected value can be calculated using the formula

**E(PW) = Sum_over_j (expected value of cash_flows for year j)$(1+i)^{-j}$**

where a second calculation may be necessary to determine the expected value of cash flows in a given year from more basic probability information.

Calculating the variance is somewhat more complicated, as it depends on the degree to which the variations in individual parameters are correlated. For example, the variation in costs due to inflation is probably well-correlated with the variation in revenues due to inflation, and neither of these is correlated with the incidence of work stoppages due to labour unrest. For the case of perfect correlation between cash flows, the associated variance in present worth is given by

**Var(PW) = Sum_over_j (Variance in cash flow in year $j)^{0.5}(1+i)^{-j}$**

In the case of zero correlation between cash flows, the variance in present worth is

**Var(PW) = Sum_over_j (Variance in cash flow in year j)$(1+i)^{-2j}$**

In practice, the cash flows may interact in complex ways. One strategy for dealing with this is the *Monte Carlo method*. The basis of this method is to set up a computer model of the present-worth calculation, and run it many times, selecting the values of the input parameters at random, but ensuring that the random distributions from which they are selected match their real distributions. (For example, if I know that there is a 75% chance that expenditures will be $3.00/item, and a 25% chance that they'll be $4.00/item, I set the value of expenditure by generating a random number between 0 and 1, and setting **expenditure=$3** if the number is between 0 and 0.75.)

After many runs of this model, I will have generated a distribution of present worths which (I hope) will approximate the true distribution.

## Making a Decision

Having generated a distribution, how do I reach a conclusion? If one strategy leads to a higher expected present worth, and a lower variance, than any other, that's clearly the way to go. But what if the strategy with the highest present worth also has a high variance, so that it may possibly lead to large losses? Should I still choose it, or would it be better to play safe?

There is no canonical answer to this question. It depends on what importance you attach to avoiding risk. Among the possible meta-strategies you could use are:

- Ignore variance, go for the highest expected value
- If the different possible outcomes can be described as a small number of possible scenarios, choose the strategy that gives the best result in the most probable scenario (this works best when one scenario is much more probable than the others).
- Decide on your goal -- for example, ``to meet the MARR'' or ``not to go out of business'', and choose the strategy with the highest probability of meeting that goal.

# Two Examples

- A communications company is looking for a launcher for a new communications satellite. They have collected data on the success rates of five national space programs -- the Russian, US, European, Japanese and Chinese -- and have quotes for launch costs from each agency. If the satellite is destroyed or lost on launch, insurance will cover the cost, but there will be a twelve-month delay before a replacement satellite can be ready. Some of the national space programs also have a habit of delaying commercial launches when a military payload needs to be put up.

Once the satellite is launched, it will start generating revenue at a rate of $100,000/month. The communications company has a pre-tax MARR of 2% per month, compounded monthly. Which agency should it choose?

| | | | Probability of Delay | | | | |
|---|---|---|---|---|---|---|---|
| | | Launch Cost ($000,000) | **On Time** | **3 months** | **6 months** | **9 months** | **12 months** |
| Agency | **US** | 1.2 | 0.5 | 0.2 | 0.1 | 0.05 | 0.15 |
| | **Russia** | 0.9 | 0.5 | 0.1 | 0 | 0 | 0.4 |
| | **PRC** | 0.75 | 0.4 | 0.15 | 0.10 | 0.05 | 0.3 |
| | **Europe** | 1.0 | 0.7 | 0.0 | 0.0 | 0.0 | 0.3 |
| | **Japan** | 1.5 | 0.9 | 0.05 | 0 | 0 | 0.05 |

*The first thing to do is to calculate the value of a launch now, three months from now, and so on.*

*The value of a launch now is an infinite series of payments of $100,000/month, where the interest rate is 2%. The capitalised cost of this series is $100,000/0.02 = $5,000,000. The capitalised cost of the same series, starting 3 months later, is $5,000,000(P/F,0.02,3) = 5,000,000 * 0.9423. A six-month delay gives a capitalised cost of 5,000,000 * 0.888; nine months, 5,000,000 * 0.8368; a year, 5,000,000 * 0.7885.*

*(Alternatively, we could observe that all options lead to the same outcome at the end of one year; at that time, the satellite is up and earning. So it would be equally reasonable to take a one-year study period.)*

*So the expected value of a launch with the US rocket is*

*-1,2000,000+5,000,000(0.5 + 0.2 * 0.9423 + 0.1 * 0.888 + 0.05 * 0.8368 + 0.15 * 0.7885)*

*= -1,200,000 + 5,000,000 * 0.937 = 3.487 million.*

*Variance = $$867,000*

*Using the Russian rocket gives us*

*-9000,000+5,000,000(0.5 + 0.1 * 0.9423 + 0.4 * 0.7885)*

*= -900,000 + 5,000,000 * 0.9096= 3.648 million.*

*Variance = $$1,009,000*

*Using the Chinese rocket gives us*

*-750,000+5,000,000(0.4 + 0.15 * 0.9423 +0.1 * 0.888 + .05 * 0.8368 + 0.3 * 0.7885)*

*= -750,000 + 5,000,000 * 0.9085= 3.792 million.*

*Variance = $$1,084,000*

*Using the European rocket gives us*

*-1,000,000+5,000,000(0.7 + 0.3 * 0.7885)*

*= -1,000,000 + 5,000,000 * 0.9365= 3.682 million.*

*Variance = $$960,000*

*Using the Japanese rocket gives us*

*-1,500,000+5,000,000(0.9 + 0.05 * 0.9423 + 0.05 * 0.7885)*

*= -1,500,000 + 5,000,000 * 0.9865= 3.492 million.*

*Variance = $$735,000*

*So the Chinese launcher looks like the best bet from the viewpoint of expected value.*

*Should we also take variance into account?*

*One way to think about this question is to consider that, if we change our choice in order to* *get a lower variance, we are paying a premium of several hundred thousand dollars to reduce our uncertainty.*

*Why might we want to do this? Well, if the company is short of cash, and a twelve-month delay might cause the company to fail, we might reasonably attach a higher importance to staying in business than to maximizing expected profit. Secondly, in making plans for the year ahead, it would be convenient to be fairly sure what our cash flows are going to be. Are either of these considerations worth losing $300,000 in expected profit? There's no algorithm that will tell us.*

*A comparison may be helpful. You are offered a chance to play one of two games. It costs ten cents to play either of them. Game 1 involves tossing a coin once; you get $1 if it comes up heads. Game 2 involves tossing a coin twice: you get $4 if there are 2 heads, otherwise you get nothing. Which game would you rather play?*

*Now consider two different games. It costs all you possess to play either of them. Game 3 involves tossing a coin once; you get ten times what you possess if it comes up heads. Game 4 involves tossing a coin twice: you get forty times what you possess if there are 2 heads, otherwise you get nothing. Which game would you rather play?*

*If our company's resources are large compared with the launch costs, and if we launch satellites on a regular basis, we will probably consider that variance in individual cases will average out in the long term, and hence make a decision based solely on expected value.*

- Your company produces a high-energy-density battery. Unfortunately, due to imperfections in the manufacturing process, an average of one battery in every ten thousand is defective. A defective battery doesn't just fail to work, it disintegrates messily, oozing a corrosive fluid.

A second company uses your batteries to power portable communications devices. They buy your batteries in lots of 2,000/year, placing a repeat order every year. If they ever get a defective battery, they will never buy from you again. (Though they won't sue you to recover damages.)

You could ensure they don't get defective batteries by putting each battery you produce through a complex test, but this will be expensive. If you make a profit of $5.00 on every battery sold, and if your pre-tax MARR is 25%, what is the most it is worth spending to test a single battery? (This question is asking you to calculate the ``expected value of perfect information'', a topic which we didn't cover in the lecture. Nevertheless, you know enough to be able to solve it.)

(Hint: if your answer is not between 0 and $5.00, you've made a mistake.)

*To answer this question, we need to choose a study period. How long do we expect the second company to be making this model of communications device? Well, mobile communications is a fast-changing field, so it may not be that long. Certainly not as long as ten years. On the other hand, there's also the question of our reputation with the other company. If we maintain our reputation, they may use our batteries in the successor to the current device. Knowledge of the market would be useful here -- are there many other customers for our batteries, or does this customer dominate the market?*

*It might be reasonable to get a high and a low estimate; for the low estimate, assume that the other company is only one of a number of possible customers, and that they're only going to* *be making this model for another three years. For the high estimate, assume that we really need to maintain a good reputation with this company, and that fulfilling this order successfully may lead to repeat business for many years to come -- ten years, for example.*

***Three-Year Study***

*Suppose the cost of inspecting a battery is **c**. Then our expected profit on inspected batteries is*

$2000 * (5-c) * (P/A,0.25,3) = 2000 * (5-c) * 1.952 = 3,904*(5-c)$

*If on the other hand we don't do the inspection, we definitely get the first year's profit, 2,000 * 5 * (P/F,0.25,1). The chance that the first lot of batteries contains no defectives is* $0.9999^{2000} = 0.819$*. So we have a 0.819 chance of getting an additional 2,000 * 5 * (P/F,0.25,2). And we have a 0.819 * 0.819 chance of getting an additional 2,000 * 5 * (P/F,0.25,3) So our expected profit is*

$2000 * 5 * ((P/F,0.25,1) + 0.819(P/F,0.25,2) + 0.671(P/F,0.25,3))$

$= 2000 * 5 * (0.8 + 0.819 * 0.64 + 0.671 * 0.512)$

$= 10,000 * 1.67 = 16,700$

*This gives the equation*

$3,904 * (5-c) = 16,700$

*which can be solved to give **c = 0.73**. So it's worth spending about 73 cents per battery on inspection.*

***Ten-Year Study***

*If we do a ten-year study, the expected profit on inspected batteries is*

$2,000 * (5-c) * (P/A,0.25,10) = 2000 * (5-c) * 3.571 = 7,142(5-c)$

*With no inspection, our expected profit is*

$2000 * 5 * 0.75 * (1 + \text{sum\_from\_i=1\_to\_9}\ (0.819 * 0.75)^i)$

$= 7,500 * (1 + (1.614^9-1)/(0.614*1.614^9))$

$=7,500 * 2.607 = 19,550$

*Solving for **c**, we get \$2.26.*

*So our overall conclusion is that it's worth spending at least 75 cents per battery; if we are* *concerned about establishing a long-term relationship as a supplier to the other company, it may be worth spending up to $2.25 per battery.*

---

---

*John Jones*
*Jan 2 10:38:23 PDT 2008*