[SQUEAKING]

[RUSTLING]

[CLICKING]

RICHARD DE NEUFVILLE: Now, given that you have a valuation model, given that you have uncertainties, you can see what the outcome, the performance of your system is when there's no flexibility-- simply saying, OK, if we build this plant and then see what the revenues are, we can think about, well, our production may go up and down because of mechanical problems or whatever.

The demand may go up and down, depending upon competition or whatever. And all together, we may have the following results in terms of benefits and costs. To do this, you generate a dynamic scenario and apply it to the model period by period. That is, the demand starts at one level, goes up, goes up, goes down, up, down, whatever. And you do that for all the different demand possibilities that you simulate, and then you can think about-- that would also happen in terms of production possibilities.

And for each of those scenarios, combination nation of the variables, you'll get a performance measure-- for example, the net present value for that. You go through the simulation-- we'll go through this process. 2,000 examples is a standard number a lot of people use. There's a way of saying-- of determining whether you think 2,000 is better than 1,000 or better than 10,000, what's the right balance.

We won't get into that for the moment, but what I'm saying is that doing 2,000 is a reasonable first order approximation at least so that you have 2,000 measures of performance, such as the present value, and you can calculate the averages-- in particular, the ENPV, the expected-- the average. And you can then think about what the extremes are and so on. You can plot the distribution. So that's the basic result of a simulation.

Now, it's important to recognize the nature of the product from a simulation analysis. A traditional discounted cash flow spreadsheet analysis to use for net present value is numbers in, numbers out. You use your best estimate for investment costs, production, sales, et cetera. You put the numbers in, you run the machine, you get a number. The net present value is such. The return on investment is such.

In a simulation, in uncertainty analysis, don't put a number in. You put in a shape, a distribution in and you get a distribution out. And the target curve is a generic name to describe what that distribution is. Typically, it's either of two shapes. One is the one I personally prefer-- but it's a matter of choice-- is a cumulative distribution.

It shows from-- there may be some chances things are really bad, more chances that things are better, and et cetera, It goes on up, a cumulative distribution to there's nothing-- it's not likely to be any better than a particular maximum. But you could also show it as a distribution. Think of it as a normal shape or something-- it could be or it could be bimodal, but a shape of a frequency or the cumulative.

How do you create the target curve? Or how does the process create the target curve? So the simulation produces many files. Now, the probability of each scenario-- we've done the simulation correctly, the way it should be done-- is that each one is equally probable. Why? Because you sample from the underlying distribution according to probability.

If the high value is 4 times more probable than the low value, you sample the high value 4 times and the low value 1 time so that you sampled it from the-- according to the distribution. So the consequence is that each of your 2,000 or however many simulations has the same probability. Given that, given your 2,000 results, the way that it proceeded mechanically is to bend the results.

But as you look at outcome ranges-- the first 1%, 2%, 3%-- maybe 100 times, 100 bins, and that-- so the number of samples in each bin is that frequency, and that creates a histogram, if you want. If you have enough 1's, it is no longer a bar chart, but it's a distribution. And that can easily be calculated into a cumulative distribution. So this entire process is easily automated.

And the garage case template, which is a very trivial spreadsheet-- but it does that for you. So if you wanted to do it yourself, you can do it, but it's not particularly necessary. It's all done for you for your projects, if you're using it, or for after class or later time.

So here is an example of it, that there's all these individual bins have been put together, and here is the probability of each, and here is the cumulative distribution, the target curve-- a target curve. By the way, just to make it clear here, so this is the-- going from 0% to 100%. So there's nothing worse than this particular number in your distribution. That's the lowest one. This is the highest one, 100%. And this one, for example, here is that, at the 0 level, for some reason it's 50%. But it doesn't have to be that way, but it is in this particular case. So this is an example of a cumulative distribution.

Now, the interesting aspect is, how do you look at flexibility? So the first question is, of course, well, what flexibility? What are we talking about? And the answer to that in general is the flexibility that is-- has the most impact regarding the uncertainty Is it high if you had a really bad case, what would you do about it? And what kind of flexibility would help you in the-- would be most impactful?

In the case of the garage case, the uncertainty is in the demand, so what kind of flexibility is the capacity meet the demand, so it's a number of levels in the garage. In a particular case, you'd have to think about that for your problem and what kind of flexibility might be most impactful. And for those of you who are doing projects, we'll look at that in detail. That is one of the questions that you will want to wrestle with.

So then the thing is, OK, I have this flexibility, for example, of building an extra floor on my garage. When should I exercise it? Should I do it right away? Should I do it in anticipation of future demand? Should I wait until demand is there? In simulation, we can't do this. We don't have a way of calculating when's the optimal time.

There are some methods for dealing with that, such as dynamic programming, but it-- dynamic programming implies there's nothing you can do about it to affect future aspects. This is the so-called path independence function. And the whole point of having flexibility is precisely that you can do things so that that powerful method does not work for us.

And there's so many paths. If we think about 2,000 samples of what may happen, that's only 2,000 samples of what may be a million possibility of paths as a combination of the different probabilities at different times of prices, and quantities, and so forth. It's much too large to be searched. So the procedure is to define what we know as decision rules.