[SQUEAKING]

[RUSTLING]

[CLICKING]

RICHARD DE NEUFVILLE: So what do we do in practice? The simple approach is to calculate the upper bound. This is known as the expected value of perfect information. So this is where I'm going to be leading you in this exercise.

So the expected value of perfect information is totally hypothetical. It does not exist. But it simplifies the analysis and gives you an upper bound, which is useful because we can definitely exclude some tests if they don't meet-- if they exceed the upper bound of the possible value.

So the concept is that there is a perfect test, which exactly which event is going to occur. I will call this-- I'm calling this a Cassandra machine. The Greek mythology, not that anybody should necessarily know it, is that this lady was Cassandra who had the gift, so to speak, of always seeing what the future would be and with the curse that nobody would ever believe her. It was all about the battle for Troy, and she could predict that somebody would die, but nobody would believe her, and so on. So that was the thing, that she could predict absolutely.

So this is the Cassandra machine. It predicts absolutely. Now, of course, that does not exist, but what is the benefit of it? So the Cassandra machine is this black box which predicts exactly what test result will occur. It produces the best possible information.

Therefore, I mean, it's absolutely true. So if you had the best possible information-- that is, if you ended up with no uncertainty-- you could make the best decisions, and you would have the maximum gain over the decision you'd make without the information.

Therefore, it gives you the upper limit on the value of the test. So what we can do if we can get the real value, we can at least set some limits on it. And this is what the Cassandra machine, this hypothetical thing, does.

The further beauty of it is that the calculation for this expected value of perfect information is almost trivial. The first time you try to do it, you'll flounder around a bit because you're not used to it. But basically, it's a very simple device and is pretty transparent.

So let me give you an example. So remember, this is the raincoat problem, that you're getting up in the morning, and should you decide to take the raincoat or not. There may be a chance of rain, rain or no rain, as in the diagram.

There's a prior assessment that the probability is 0.4 or 0.6 of rain or no rain. And there is these values of the situation in the end of 5, minus 2, minus 10, and 4, values that I've invented just to give us numbers to play with. But that's the original problem.

Now, just to set ourselves back into this context, the better choice at this point is what? So how do we remind you how it is? If we take the raincoat, which is a top decision, we have a 0.4 chance of having 5, which averages to 2. We have a 0.6 chance of getting a minus 2, which is a minus 1.2. 2 minus 1.2 is 0.8.

Conversely, if I don't take the raincoat, I have a probability of 4 of the minus 10. That's minus 4. A probability of 0.6 of 4, 2.4. So it's a negative number. So my best choice is to take the raincoat, right? So I hope that you're all now back into this particular groove.

So the setup for the probabilities for the EVPI. It's either going to say it's going to rain or not rain. That is the perfect machine. You press the button, say truth please, and it'll say Cassandra machine says rain. Or it will say the contrary because I've only given it two choices.

Now, what are the probabilities it's gonna say that? Well, your best estimate of these probabilities is what you have started with before. You thought there was a 40% chance of getting rain, so 40% of the time, it's going to tell you it's going to rain. And conversely, so that you-- a priori, you have your best estimates for the value of the outcome is what your priors were.

Now, in general, once you know what the resolution of the uncertainty is, the best decision is pretty obvious. If you knew it would rain-- going back here, if you knew it was going to rain, you're comparing 5 versus 10. Well, you'd take the raincoat. If you knew for sure it wasn't going to rain, you compare minus 2 and 4, and you don't take the raincoat.

So you don't have to do any complicated calculations of the possibilities. You would have a probability of 0.4 that would take the raincoat to get it to 5, and the probability of 0.6 of not taking the raincoat and getting a result of 4. So you can write down your best decisions pretty automatically.

Now, the expected value after this case, as I'm running this calculation, doing the 0.4 times the result of 5 and the 0.6 of 4, your perfect value is 4.4. Now, the expected value of perfect information is the value of the increment, not the value at the end of it. It's the value of the increment. What did you get from it? So the expected value after the test was 4.4. The expected value before the test was 0.8. The expected value of perfect information in this case is 3.6.

All right, so the nice thing about it is the prior-- the expected value of simple-- perfect information is simple to calculate. The prior probability must equal the probability of the perfect test result. Once you've run the perfect test, the posterior probabilities are either 0 or 1. Either this happens, or it doesn't happen. There's no doubt about that.

Once you know that there is a whole lot of 0 probabilities and there's a one probability, you know what you're going to choose, and pretty much you can write the expected value of perfect information formula directly. There's no need to use Bayes' Theorem.

So Robert, since you asked the question before-- I'm not trying to put you on the spot. But since you've been doing it, and I'm showing you this. Why did you not use this approach, if you knew about it? Or maybe you didn't know about it. So what is the-- can you have a discussion about this, please, or your response to this, if I'm not embarrassing you or putting you on the spot, which I don't mean.

AUDIENCE: No, that's fine. I actually haven't heard this perfect information idea. Yeah, so that's probably the main reason I didn't use this. But the other one was at Chevron, we kind of have a defined procedure in terms of how we justify value of information. So that procedure was utilizing the conditional probabilities.

But myself, I was personally a bit uncomfortable because like you were saying, it's hard to figure out what the percentages would be for either a good, medium, or bad outcome. So it was pretty challenging to actually apply, and I think at the end, people weren't super confident in it, like were saying.

RICHARD DE NEUFVILLE: Well, I think that that's right. So about 30 or 40 years ago-- I mean, I'm old enough to remember that-- there was a phase where decision analysis was really taught a lot in business schools, which has since gone by.

And companies picked up on that because it was certainly better than not looking at uncertainty. And it was nice to have this notion of a Bayes' rule, a Bayes' Theorem. And people thought of themselves as being Bayesian or not.

And so there's a whole history about why this was the correct improvement in approach, not thinking about the fact that these other issues, the probabilities-- the associated probabilities and the probability of occurrence and so forth, which are the part of that special factor for the operation of the Bayes' Theorem are just basically hard to get and hard to document and hard to prove.

So I think that you are reflecting, which may exist elsewhere-- I'm not pointing a finger at Chevron-- but that they have a decision analysis manual. They have an approach, a procedure, and so forth. But it doesn't really fly. And as you suggest, that a lot of people don't believe in it.

And I would be one of those, if I were in Chevron, because it's such a spongy-- the calculation is correct, but the data are spongy. And it's the GIGO, garbage in, garbage out kind of situation, I think.

So I understand that you're encouraged to do that, and this EVPI is trivial mathematically. I mean, you don't have to be smart or think through the complexities, but it gives you an upper bound. And that is the useful aspect of it. And I thank you for noting that it is not the usual way of approaching things.

But I think that it's a reasonable way to do it, and productive. And I would encourage you to take the message back. Now, I can see that they're not going to change the procedures for you. And so what people do in the case if the procedure is wrong, they sort of try to ignore it, which is what you reflected.