US Fish and Wildlife Service | Making Decisions under Uncertainty_ Monitoring and Adaptive Management _Part 1_

[MUSIC PLAYING]

So we've been talking about uncertainty and what we do in the face of uncertainty. In some cases, we can make decisions anyway in the face of that uncertainty. And that's what we talked about in the risk module. About when we have to live with uncertainty, how do we make decisions in the face of that, and guard against risk, and fold in our risk attitudes appropriately?

But now we want to turn to another set of questions. What if we can conduct research to reduce that uncertainty and then make a decision later? How should we proceed there? So I'm going to talk about the value of information and the roles of monitoring in that case.

And then I'll talk about the idea of, what if we can make decisions at the same time that we're reducing uncertainty that will do these things simultaneously? And that's the idea behind adaptive management. And we'll talk a little bit about adaptive management as well.

OK. Reducing uncertainty through monitoring and the value of information. Let's talk about this topic. We haven't really talked about monitoring yet. Jeez, we've gotten to the 10th module of this course, 10th and last module, and we haven't even talked about monitoring. Not really.

As scientists-- I'm one--we have a strong tendency to ask for more information. We love to ask for more information. And this leads to a tendency to postpone a decision in order to gather more information.

Why do we do that? Why would we postpone a decision waiting for more information? What do we think we're going to do with that information? Presumably, we think that information is going to enhance our ability to make the decision and that the costs of delay are more than offset by the gain in performance that we get with that new information.

That's a powerful concept, but how often do we really trace out that decision? Notice that's a decision in and of itself. The idea to postpone a decision until we have more information is a decision of its own.

And there must be some trade-offs. There must be some analysis that's associated with that. How do we do that analysis? How do we evaluate the importance of information?

So we do. We seek information to reduce uncertainty. And we can ask, will that information change our decision, enhance our performance? Or the other possibility is that the information is really not relevant to the decision, and I'm afraid that a lot of the time, when we delay decisions in order to gather information, that actually if we had stopped and done the analysis, we would have found that that information wasn't going to be relevant to the decision. That in fact, we're just using that as a stalling tactic.

So I think we've got to be concerned that when we go out and gather information, and otherwise delay decisions to gather that information, that we're doing it for the right reasons. There are a set of techniques that go under the title of the Expected Value of Perfect Information, or the Expected Value of Partial Information, or the Expected Value Sample Information.

These are analytical techniques that come to us from the decision analysis literature. These techniques allow us to assess how much our management might improve if we can resolve uncertainty. And they can help us decide if it's worth the cost of gathering that information.

Is the cost is gathering that information-- the direct cost, the monetary and staff time cost of gathering it, but also the opportunity cost of delaying your decision--are those costs worth it? Is that information going to change how we make our decisions in the future so that our expected performance is going to be higher?

Now I might note that in a decision analysis setting, the expected value of information is exactly how you should do sensitivity analysis. Because what it's asking you is, how much would your decision change if you could reduce uncertainty? And how much would the expected performance in the long run change if you could reduce uncertainty? And that's a critical question.

So how does it work? Let me demonstrate for you an example of the expected value of information. And then I'll give you a chance to try it yourself.

So suppose that you're managing a forest, and you're managing a forest primarily for timber production. OK, there's a forestry management setting. And you're trying to minimize the lost revenue from a gypsy moth infestation.

So gypsy moths have come into this forest, or they could come into this forest. And if they do, they're going to reduce the yield from the timber operations. And you want to minimize the lost revenue due to that infestation from gypsy moths.

And let's suppose you have three possible actions. You can do nothing. You could reduce colonization, the colonization rate of the gypsy moth. So you could reduce the rate at which they're coming into the forest. Or you could eradicate large patches of gypsy moth infestation.

And let's suppose those are three quite different strategies, and they have different costs associated with them, et cetera. Let's suppose also that you have two different models of how gypsy moths, and hence your timber sales, will respond to these actions.

So you've got some uncertainty. You don't know which model is the case. Model one says that if you do nothing, your lost revenue will be $299,000. But if you reduce colonization, the loss will only be $202,000. And if you eradicate large patches, the loss will only be $140,000.

Model two is different. It says if you do nothing, your loss is going to be $493,000. If you reduce colonization, you'll reduce that loss to $256,000. But if you eradicate large patches, the loss will be reduced to $273,000.

So these models have different predictions about what's going to happen under the three actions. And let's suppose based on prior research, you think the likelihood of model one is 30% and the likelihood of model two is 70%.

You could make the decision in the face of uncertainty. You could treat this uncertainty, the model one versus model two, as branches on the decision tree, and do an expected value calculation, just as we did in the last module. And then, perhaps you could even invoke the concepts of utility if you wanted.

But let's suppose there's also a possibility of finding out, before you have to decide on a course of action in your forest, suppose you could find out which model is in fact the case? How much is that worth to you? How much is it worth it to acquire that information, to reduce that uncertainty?

Let's look at how you would do that calculation. So the first thing we can do is calculate the expected value under uncertainty. So the idea here is, suppose we had to make the decision right now, and we didn't have the opportunity to reduce uncertainty? So we're making the decision in the face of that uncertainty.

What would we do? Well, let's calculate the expected value of these three actions. There's two possibilities. If model one is the case, then our lost revenue is going to be $299,000. If model two is the case, our lost revenue was $493,000.

So if we take the weighted average of those two things, weighted by our belief in model one, 0.3, our belief in model two, 0.7. So weighted average of the two losses. We get an expected loss of $434,800. So that's our expected loss if we do nothing in the face of uncertainty about which model is the case.

We do similar calculation, weighting by model one and model two, for the action of reducing colonization. And we get an expected loss of $239,800 for that action. And if we eradicate large patches, our expected loss in the face of uncertainty averaging over the two models is $233,100.

And so since we're wanting to minimize the expected loss-- let's suppose we're risk neutral here. If we want to minimize the expected loss, then we choose the action that's got the lowest expected lost. That's $233,100. And that's the "eradicate large patches."

So what we would do if we had to make the decision in the face of uncertainty is choose to eradicate large patches. Implement that. Either our loss will be $140,000, or it'll be $273,000. We believe in advance that the way to weight those is 30% and 70%. And so our expected loss is $233,000.

That's the smart thing to do in the face of uncertainty if we have to live with that uncertainty. But if we could reduce the uncertainty first, what would happen? Well, the expected value under certainty--if we knew model one is the case, then the smartest thing to do would be to eradicate large patches. Because under model one, that has the lowest loss. Right? Action three, eradicate large patches.

So if we knew it was model one, we're going to eradicate large patches, and our loss would be $140,000. If we knew it was model two, the best thing to do is reduce colonization, because that has a loss of $256,000 instead of $273,000, which is the loss under "eradicate large patches" for that model.

Notice that the point here is, if model one is the case, then the best thing to do is eradicate large patches. If model two is the case, the best thing to do is reduce colonization. So this is uncertainty that matters. Depending on which model is the case, we would do something different.

Uncertainty matters to the decision. It doesn't just matter to our expected loss. It matters to the decision that we would make. So it's important. It's important to the decision.

So if we take the expected value now, we don't know in advance what we're going to find out if we reduce this uncertainty. There's a 30% chance it would be model one. If it is, then we would eradicate large patches and have a loss of $140,000.

There's a 70% chance it's going to be a model two. If it is, we would reduce colonization and have a loss of $256,000. 30% times $140,000 plus 70% times $256,000, that's an expected loss of $221,000.

What this means is, if we have to make the decision in the face of uncertainty, the smart thing to do is eradicate large patches, and our expected loss is $233,000. If we can reduce the uncertainty first, well then, we won't say in advance what we're going to do. It depends on what we find out. But prior to finding that out, our expected loss is $221,000.

Now that difference, the difference between $233,000 and $221,000 that difference of $11,900, that is the reduction in the expected loss that we get by reducing uncertainty. That is the Expected Value of Perfect Information.

If we could perfectly resolve our uncertainty about whether model one or model two is the case, that reduces our expected loss by $11,900. That's the value of that information.

How much would you pay for that Information? If you had to pay to do the research study that was going to resolve that uncertainty, how much is it worth to you? Well, it's worth no more than $11,900.

If that study is going to cost you more than $11,900, then you're better off just making the decision in the face of uncertainty right now, eradicating the large patches and moving on. If that study is going to cost you less than $11,900, then it's worth it to do the study first and then make your management decision based on the outcome of that study.

So that's how the Expected Value of information works. Powerful, powerful tool in terms of telling us whether information is going to affect our decision and how much. And how much it's worth it to us to acquire that information.

OK. Now I have an exercise for you to try. This is an example that comes from Carl Walter's book on adaptive management from 1986. And he paints the following hypothetical picture.

Suppose you're managing a salmon fishery, a sockeye salmon fishery. And your objective is to maximize the net economic value of that fishery. But you've got a decision here. You've got a decision of whether it would be advantageous to build artificial spawning channels.

And let's say you're making this decision in the face of uncertainty. OK, and let's suppose you got two options. You don't build the channels, or you build the channels. Let's suppose that you know that currently, the net economic value of this fishery is $240 million.

And so if you don't do anything, if you don't build the channels, then you expect the net economic value is going to be $240 million in the future. But if you build the channels it's going to cost you. It's going to cost you, let's say, $105 million to build these artificial spawning channels. And the question is whether there's going to be any increase to the productivity of that fishery as a result of these channels.

And let's suppose there's uncertainty. Let's suppose there's two models. One that says, no, these things don't work. They're not going to do anything for you. And so, if that's the case under model one, the no response model, if you build the channels, you're out the $105 million, and there's no gain. And so the net economic value is reduced to $135 million.

But let's suppose these things work. If they work, then you expect really great increased productivity, and the economic value of the fishery will increase from $240 million to $564 million, even accounting for the cost of building those artificial spawning channels.

Which do you do? Well, what I want you to do is to figure out what is the-- and let's further suppose that we don't know which of these models is the case. And with the current information, we'll put a 50/50 probability on these two models.

I want you to figure out, what is the expected value of perfect information in this case? How would you figure out how much it's worth to conduct a study that could perfectly resolve uncertainty and tell you in advance of building these channels, or in advance of choosing to build these channels, whether model one, the no response model, or model two, the good response model, is in fact the case.

OK. So take some time. Pause the video. Either do this individually or work with the group that you're taking the course with and calculate the expected value of perfect information for this problem.

OK. So you've had a chance to work on this exercise. Let's walk through it. I'll show you how I do this calculation, and you can see if it matched up with how you did this exercise.

So the first thing I would do is I would say, well, what I do in the absence of new information? What would I do if I had to make the decision today? Would I build the channels or not build the channels in the face of uncertainty?

So let's calculate expected value. So the value of not building-- well, under both models, the net economic value is expected to be $240 million. So if I don't build the channels-- actually no, my net economic value is going to be $240 million.

I'm weighting the $240 million from the no response model with the $240 million from the good response model, and I'm weighting those 50-50. But because the outcome is the same, it comes out to be $240 milion.

In the case of the "build the channel" option, in the face of uncertainty, there's a 50% chance that I'm going to get a return of $135 million and a 50% chance that I'm going to get a return of $564 million. That averages out to be $349.5 million.

So what would I do? In the face of uncertainty, actually, the expected value, the expected net economic value of the sockeye fishery, is higher under the build option versus the not build option.

Now there's some risk there. If I have to make the decision in the face of uncertainty, the higher expected value is associated with the "build the channel" option. And so that's what I would choose if I was a risk neutral decision maker and I had to make this decision in the face of uncertainty.

Now if I could resolve uncertainty before I made my decision, if I found out the no response model was the case, what would I do? Well, I would not build, and my return would be $240 million. Right? If I knew it was a no response model, there's no way I'd build the channels.

Likewise, if I knew it was the good response model, I would build the channels, and my return would be $564 milion. So the two possible responses I get here, if I could reduce uncertainty, are on the diagonal. Right? If it's the no response model, I get $240 million. If it's the good response model, I get the $564 million.

Before I've done the study to reduce the uncertainty, I don't know what I'm going to find out. So I still have to do this 50-50 weighting. So it's a 50% weight on $240 million and 50% weight on $564 million. Those are the diagonal elements of that table. And that averages $402 million.

So if I can wait to make my decision about building the channels until I've resolved uncertainty, then my expected outcome is $402 million. Now the difference between the $349.5 million and the $402 million, that difference is $52.5 million. That's how much my expected performance is increasing by reducing the uncertainty before I have to commit to the decision to build the channels or not.

So that's the expected value of perfect information. $52.5 million. How much would I be willing to pay for that study that could resolve that uncertainty before I built the channels? I'd be willing to pay up to $52.5 million. If somebody offered to do it for $25 million, I'd take the deal.

So hopefully that's another example that gives you a sense of the value of information. And again, I'll say, this is really a powerful technique that we ought to be using a lot more often when we're considering the issue of, do we delay a decision in order to do some research, or do some monitoring, to reduce uncertainty, and then make a decision with a better set of information?

If it's going to change our decision, and if the value of changing that decision has a high expected value of information, sure then that makes sense. But I think a lot of times, as I said earlier, we don't stop to do this calculation or even think through it properly. And so, there's plenty of cases, I think, where we're making decisions to gather information which are really largely unwarranted in the applied context of making decisions.