US Fish and Wildlife Service | Consequences _part 1_

OK, welcome back. This is Module E, Consequences. And before we start, I'd like to acknowledge a couple of other people who have contributed to developing this module-- Dave Smith, Jennifer Szymanski, and Jim Lyons.

In Module C, Mike talked to you about objectives, where is it that we want to go. And in the last module, Module D, I talked to you about alternatives, what are our options for getting there. Now we're going to talk about the thing that links them together, predicting the consequences of different actions in terms of our objectives.

So again, we can picture this as this sort of three-legged stool. The analysis, or the logic, which we're going to talk about a bit in the next module, is what holds the stool together. That's the top. The three different legs are the objectives, where do we want to go, the alternatives, how can we get there, and the consequences. And really the consequences, that part of the process captures what we know and what predictions we can make. It's sort of the science step of the process.

It's now about understanding the consequences of different actions in terms of our objectives. So the key thing to remember about the consequences step in PrOACT, which we'll often call prediction or modeling, links objectives and actions together. And we make statements like, if I take action A, the outcome in terms of the objectives that I have-- the things that I care about-- is X. Those are the kinds of statements that we need to make in order to understand which of our actions might be the best action to take.

So, the models are really the tools that we use to predict consequences. And that's what modeling is about in structured decision making. And really, that's what science is about in structured decision making. It's about allowing us to predict the consequences of taking a particular action.

But it's important to say, oftentimes when people hear the word "model," they feel intimidated and they think that it's going to be something very complex. But models don't necessarily have to be complex. It depends, really, on the decision that you have to make and the knowledge that you have. And so on.

So, for example, if I have an 8:30 meeting at the office, I need to make a prediction about what time to leave the house. So one of my alternatives is to leave at 7:45. And I would then need to make a prediction. If I leave at 7:45, will I make it to my 8:30 meeting?

Now we all make predictions like this every day. We all have sort of mental models in our head that are the product of our experiences over time. Or lots of different ways that we have of bringing information together to bear on decisions. One thing that we do do is-- in structured decision making-- we try to take that in-your-head kind of stuff and put it down and make it explicit so that everyone can see it. So that we can get better at using the knowledge that we have. And that we can be more complete in bringing knowledge into the decision-making process.

So let's, again, start with an example to motivate these ideas. So I have a decision to make, and in this case, I want to make a decision about whether to get a new pet. And my question is, what kind of pet should I get?

So I start by thinking about-- now that I know my decision problem-- what matters to me. What do I care about? So maybe I start by thinking of these things in terms of my concerns, as we did in Module C.

So my concerns are that, well, I travel a lot. So that seems to be important to keep in mind when I'm choosing a pet. I'd like to avoid the high vet bills that might come with certain kinds of pets. I'd like an animal that's good with children. I'd like to keep the cost of feeding the pet low. I certainly don't want any kind of a pet that is going to bite visitors who come to my house. I also recognize that, yeah, I don't have a lot of time to devote to care of a pet, partly because I travel frequently. And when I am home, I'm pretty busy.

These are my concerns. Now I need to translate these into objectives. It might look like this.

Well, the concern about avoiding high vet bills and keeping food costs low. I can more simply state that as I just need to minimize, or I want to minimize, the total cost of owning a pet over time. I might measure that in terms of dollars of expected costs per year.

A couple of other issues that I had that were pretty closely related. I wanted the pet that I chose to be good with children and not to bite visitors. So I want to maximize the friendliness of the pet that I choose. And so maybe I would measure that on a scale from one to five, where one is low friendliness and five is high friendliness. So that's a kind of a constructed scale.

I had the concern about traveling frequently, so maybe in this case, what I realized that means to me is that I want to maximize the number of willing pet sitters amongst my friends. And so I measure that by asking my friends and counting up the number who would be willing to pet sit for a pet of a particular kind.

And then last, I recognized that I didn't have a lot of time to devote to care of the pet. So I translate that into an objective, which is I want to minimize the required care. And I might measure that in hours per week. So I've translated my issues or concerns into objectives. And I've devised a way to measure them.

OK. Now I need to think about what my alternatives are. Well, I might do a lot of different things. Perhaps I just go down to the local pet store, or I open up the newspaper, or I go to the Humane Society. And I come up with a list. So, for example, I could choose a cat. I could choose a hamster. Or I could choose a goldfish. Or I could choose a very cute poodle, for example.

Then I need to understand the consequences of different kinds of alternatives in terms of the things I care about. So this is what is called a consequence table. It's currently blank, but we're going to fill it in in the next couple of minutes.

You're going to see consequence tables a number of times throughout the rest of the course. And consequence tables can be a really useful way, when you have multiple objective problems, of visualizing the consequences of different alternatives. They can also be useful in the long term in actually analyzing the decision. And we're going to come back to that in the multiple objectives module a little bit later on.

But let's look through this consequence table. So, I've listed my objectives down the left side of the table-- minimize cost, maximize friendliness, maximize pet sitters, and minimize the hours of required care. I've listed my measurement criteria. And I've also listed the desired direction. I want to minimize total cost, so I'm interested in low numbers there. I want to maximize pet sitters, so I'm interested in high numbers there. Now, along the top, I have the rest of all of my alternatives listed-- poodle, cat, fish, and hamster.

Now the cells, what I can do is I can use those cells to fill in my predictions of the consequences of choosing each one of these alternative pets on each of the objectives that I care about. Perhaps I ask experts at the pet store, what do you think the total cost of caring for one of these kinds of pets would be? Or maybe I talk to friends who own those different kinds of pets. In some form, I make some predictions about the cost of owning each of these different kinds of pets.

And then I need to make predictions about friendliness next. I have my constructed scale. And again, maybe I survey my friends. Or I survey people at the pet store. In some way, I develop predictions that the poodle will be the friendliest and the fish will be the least friendly.

Then I need to do the same thing with pet sitters. Ask my friends whether they would be willing to. And then I also need to do the same thing with required care. Again, maybe my pet store would be a useful place to get information like this.

So now, what I have is I have a prediction of the consequences of taking any particular action-- that is, buying or adopting any particular pet-- on each of the four different objectives that I care about. Again, this consequence table is a tool that we're going to see a number of different times. And we're going to focus quite a lot of attention on this tool during the multiple objectives module.

The question is now, which pet should I choose? Well, it might not be entirely apparent. It depends on how much I care about cost versus friendliness versus pet sitters, and so on. So there will be more analysis that's needed.

But even at this point, there might be some things that are really apparent to me. For example, if I really think I care a lot about friendliness, and maybe the difference between $20 a year and $300 a year doesn't seem overwhelming to me, maybe I eliminate the fish right off. Maybe I don't. Maybe I keep it in the analysis. But even with just doing these steps, sometimes the decision that I'm going to make becomes a little bit more clear.

We're going to talk, as I said, in the multiple objectives module a lot more about how to analyze a consequence table like this. But it is a very useful tool for visualization. There's a lot of information in the consequence table. It's very concise. It's orderly. It's easy to see. And it's an initial framework for addressing the trade-offs.

Now it's also important to say that I made predictions, or I used models, to fill in that consequence table. And that required me to make projections into the future. I had to predict how much that poodle or hamster would cost me into the future. Or how many of my friends would be willing to pet sit next year when I actually had that pet at home and that was a reality. So we're making projections into the future when we make predictions about consequences.

It was important that I had a common scale within each objective. And Mike talked to you quite a bit about that in the objectives module-- developing that way of measuring each objective. It was on a common scale. And my models, it's important to realize, might have included a little bit of hard data and some more subjective assessment. Like friendliness is hard to get from real data. Whereas cost, we might actually have data out there in some way to make estimates of total cost.

So the models that we use are often some combination of data, subjective assessment, and so on. And we also use expert opinion and things like that. So we try to make the most of available information to predict consequences. Oftentimes we have to go to unexpected places to get the kind of information that we have. But we shouldn't feel overly restricted by the availability of hard data.

And we also want to report the appropriate level of precision. Maybe it's more sensible to say it's $300 give or take $50, rather than it's exactly $329 a year. So we don't want to be overly precise or underly precise in making predictions. We might want to make the precision reflect what we really do know.

And last, it is clear that we didn't incorporate uncertainty in any formal way into our consequence table. But in the long term, we might recognize that we do have some uncertainty. We'll talk a little bit in the multiple objectives module about how we can use sensitivity analysis to deal with the kind of uncertainty that we might have. For instance, we don't know the exact cost. We only have a prediction of the cost. And it might be $50 one way or the other. It might also be a reasonable estimate.

OK. Let's look at one more consequence table. And then we'll talk a little bit more about modeling. So again, this consequence table, which comes out of the Cultus Lake sockeye example from the Gregory and Long 2009 paper, which will be included in the list of useful references for this course. This is an example where we've got a number of alternatives along the top, and a number of objectives along the side. And what the color coding does here is it, again, emphasizes how consequence tables can just be useful for understanding very quickly the kind of information and the kind of alternatives we have.

So what they've done in this consequence table is they've coded each of the cells depending on its relationship to the status quo. So the first alternative listed in blue is the status quo alternative. The cells that are red are cells that perform more poorly than the status quo on a given objective. The cells that appear in green do better than the status quo on a given objective. And the ones in yellow are ones where the difference is so small, it can basically be ignored. So, basically, it performs very similar to the status quo.

So this is a situation where we can look pretty quickly and we can see that for some of the objectives, we've developed some alternatives that do very well. And for some of the objectives, we haven't really developed alternatives that do very well. That might inspire us to, say, develop some more objectives that are a bit more in between. So some that, say, do very well on the objectives that currently we don't seem to be able to do very well on, and vice versa.

Again, it's just a quick visualization of the decision problem. It includes a lot of information pretty concisely. And, as I said, we'll go back to these kinds of tables at quite a bit of length in the multiple objectives module to talk about how to analyze them.

Another example that might be useful is the Florida whooping crane release example. Now this was a problem that I and a colleague were involved in, where we were making a decision about how many whooping cranes should be released into a non-migratory population in Florida. And, again, here we have our alternatives along the top. So the alternatives, in this case, were a combination of the number of years to continue releasing cranes, the number of cohorts of cranes to release each year, and whether or not there should be a delay of about five years in order to get another reintroduction project off the ground.

And there were six objectives that were important to the decision makers. And, again, we have estimates of things like cost. We have estimates of some of the other objectives, like the probability of success in the population, which is the next-to-last objective. The public relations outcomes of different alternatives, and so on. So again, this was a problem that we analyzed using multiple objective techniques that we'll talk about in the multiple objectives module a little bit later.

But let's now talk about where does the information come from to populate these kinds of consequence tables? How do we build models, and how do we use models in structured decision making?

So it's important to recognize that the models, again, don't have to be incredibly complex. They can range from very simple to very complex, depending on the needs of the decision framework. They may or may not be highly quantitative. They may or may not involve a great deal of uncertainty. It really depends on the alternatives, on the objectives, and what is known.

The characteristics of a model are case-dependent. So it's important to say that when you're using models to help you make decisions, oftentimes it's not possible to simply pull a model off the shelf. The model has to really be structured in such a way that it really does address your objectives and your alternatives. So, the role that the model's going to be used and the other components of the decision framework really need to be understood before the model is developed.

It's important to recognize, again, that models link our alternatives and our objectives. So, inputs into our models will be based on the actions that we could take. And outputs from our models will be linked to the objectives that we have. So we put into our models, let's take this action. And what we learn is how taking that action will affect the things that we care about, our objectives.

Through models, and using models in decision analysis, we do three different things. And let's talk through those in greater detail.

First of all, they help us to structure the analytical problem itself. They just help to clarify our thinking. They lend transparency to the analysis. So that makes it important that our models not be black-box models. But that they actually help us think about the problem, be transparent about the problem, communicate the problem to others, and so on. And last and most obvious, again, they help us develop the predictions of consequences so that we can evaluate the alternatives that we're considering.

So let's talk about a little bit more in-depth about these three benefits. So first, as I said, models help us to structure the analysis. So, oftentimes in determining the consequences of alternative actions that we're considering, that process involves multiple analytical steps. So models can be useful for thinking through what these analytical steps have to be.

They can help us by initially displaying the problem graphically. And we'll talk quite a bit about that in this module. What are the key elements? What are the relationships between the key elements in the system that we're trying to make predictions about? And, by decomposing the problem in this way, seeing the steps, we can understand the steps that we then need to take in the analysis in order to get the point where we're making predictions that we feel confident about.

Second, models allow us to lend transparency to the analysis. Remember, one of the characteristics of a good decision is it's a decision using a process that will be compelling not just to you, but to other people. And that's especially relevant when we're talking about managing publicly owned resources. We need to be able to communicate about why we're taking the decision we're taking. And we also need to be able to communicate that within the team we work in.

If we understand why we're making particular predictions about how our system works, we also understand what we know, what we don't know. And it allows us to get better over time. If I just make a guess, sort of a black-box model in my head about the impact of taking a particular action, no one can challenge that prediction with data. And so I don't get better over time at making that prediction. Being explicit about the models that we build to make predictions helps us get better over time. And, again, it helps us to communicate to other people why we're taking a particular decision.

So every decision maker uses some kind of a model to predict consequences. But that doesn't mean that some models aren't more transparent than others. We want to make our models quite explicit and available to everyone involved. So during the model development process, we need to illustrate the key elements of the problem, key elements of the model, the relationships among them. And capture all the kind of complex information in a way that can be communicated to various different kinds of people, including everyone on the team.

And last, and really, probably most relevant and important and obvious, is that these models allow us to evaluate the different alternatives that we have. To compare and contrast alternatives, we have to predict the future outcomes of taking different kinds of actions. And, specifically, we make those predictions in terms that are relevant to our objectives, so in the measurable criteria that we've developed to correspond with each of the objectives that we have.

It's useful to look at a couple of different quotes that have been attributed to various people about what models do and don't do. And perhaps the one that people are most familiar with is the top on this list. All models are wrong. Some models are useful. And I think that's particularly apt in the case of using models in decision analysis.

We recognize that our models can't possibly capture all of the complexity of the system that we're making predictions about. The question really is, does it capture enough? Does it capture the information in a useful way? Does the model we build to help us make decisions help us be more effective in making decisions? Help us be more transparent and more likely to get to where we want to go?

So let's take a few minutes and do an exercise. Imagine for a second that you're a manager of a Fish and Wildlife Service refuge. And you have to make a decision. You have three adjacent areas, and you are interested in acquiring one of those. Each of those adjacent areas contain some either current or potential habitat for an endangered butterfly that you care about. As a manager, you want to minimize the probability that the butterfly will go extinct while also minimizing the cost to you of acquiring that land.

The decision is which of the pieces of land to acquire. You want to minimize butterfly extinction. And you want to minimize cost. You need to then build one or more models to help you make that decision.

Take a few minutes and answer these questions. One, what will the models need to be able to predict? And two, what are the benefits of building some explicit predictive models in this case?

OK, so take a minute. Turn off the video and jot down your answers to those questions. And then we'll come back, and we'll talk about them.

So we're back, and hopefully you had a chance to make some notes about the questions of first, what is it that the models you need to build have to be able to predict? Well, they need to make predictions about the objectives that we have. They need to make predictions about the probability of extinction of the butterfly in the cases of buying each of those different parcels of land. And they need to make predictions about the cost to the refuge of buying each of those different parcels of land. So, at the most basic level, the models need to be able to make predictions of the objectives, choosing each of the different alternatives.

And then also, another question that I asked you to think about is, what are the benefits of building explicit predictive models in this case? Well, there's lots of different benefits. For one thing, you might realize simply by being explicit about building a model that there's quite a bit of complication in predicting, say, the probability of extinction of this butterfly. You might realize that you don't know everything that you need to know and perhaps you need to bring some scientists in who know about the ecology of the butterfly and can help you build those models.

You might realize that, by being explicit, say, your staff biologists know a bit more about this butterfly than you do. And you can get them involved in building the model. And they can help you to build the best model possible to be most accurate about the predictions that you might make. So there are lots of benefits of building explicit models in this case.

You might also have to justify your decision to your supervisor or others up the chain if a lot of money is going to be spent. Say, on buying a particular parcel. So you might find that being quite explicit in building our models help you to explain your decision to other people that you have to convince.

But again, it's important to say here that, at the most basic, what we need to be able to do is make predictions, in terms of our objectives, of taking any particular action. And that's what our models help us to do.