US Fish and Wildlife Service | Trade Offs_ Overview of Optimization

[MUSIC PLAYING]

Welcome back. This is Module G, Trade-offs. This is an overview of optimization and analytical methods, and this module was developed by Mike Runge. So if you think back to our three-legged stool, in this module we're going to talk about the seat of the stool, the thing that holds everything together. And ultimately optimization analytical methods are the way that we actually develop a solution to our decision problem.

So an outline of this module, we're going to talk about the rule of optimization, both a general overview, and then we're going to talk a little bit about classes of solution methods, so depending on whether you have single objective, multiple objective problems, whether you have uncertainty in your decision problem or not, and so on. And then we're going to talk a little bit about some single objective tools that apply both in a deterministic setting, when we don't have uncertainty, and in a stochastic setting, when we do have uncertainty. And then we're going to talk also a little bit about multiple objective tools, although we'll cover those in much greater detail in the next module. And then we'll finish up by talking a little bit about linked decisions, and linked decisions we'll also cover in a lot more detail-- Mike Runge will cover in a lot more detail in the last module in the course.

So the role of analytical methods, as I said, this is the seat of the three-legged stool. This is the bit that holds it all together. We integrate all of the components that we've developed so far.

Our objectives, where is it that we want to go? Our alternative actions, what are the different ways we have for getting there? And then the predictions that we make that link those two things together, the objectives and the actions, and make predictions about the impacts of taking different actions.

We need to integrate all those components now to find and identify the truly smartest solution that we can identify amongst the set that we're considering. Identifying the smartest solution, the best or optimal solution, is easiest with a single objective problem. But we often are actually dealing with multiple objective problems. And we have a lot of tools available to deal with those multiple objective problems. We're going to introduce those here and then cover them in greater detail in the following module.

It's also certainly easiest to identify optimal solutions when we don't have uncertainty, but we can deal with situations where we have either natural uncertainty, what we would call process uncertainty, so

things like-- it's very hard to predict exactly what the weather is going to be like today-- or where we have knowledge uncertainty, epistemic uncertainty, where we just really don't understand exactly how our system functions.

It's important to recognize that in all of this, the solution method depends on the particular problem that we're dealing with. So depending on whether we have a multiple objective or a single objective problem, depending on how important uncertainty is and exactly how large it is, depending on whether or not we have a linked decision, we're going to choose different solution methods. And the solution method is developed after all the other parts are understood, after the objectives are understood, and the alternative actions, and the models are understood, because the solution method has to accommodate all of those other pieces.

We can look at this graphic that's useful. And what we are looking at, basically, is a table that considers whether or not we have a single objective or multiple objectives, and considers whether we have uncertainty or whether we can ignore any uncertainty that we have, and so we basically can treat it as a situation where there's no uncertainty. But if we have a single objective and we have no uncertainty, or at least we can ignore any uncertainty that we do have, we have a fairly wide array of solution options available to us, graphical solutions, numerical solutions.

We can use calculus. We can use linear or non-linear programming. These are very traditional solution methods that come out of management science. And we're going to introduce you to just one or two of those here in this module.

When we're in a single objective setting where we do have some important uncertainty that we can't ignore, the classic approach to dealing with that situation is a decision tree. And I'm going to introduce you to a decision tree today. There are also related methods like Bayesian nets, Bayesian belief networks, which we talked a little bit about in the consequences module. We can also use things like simulation to deal with uncertainty.

Now, when we have multiple objectives settings, it becomes a bit more complicated. When we don't have uncertainty, there are some things that we can do. We can use even swaps to make trade-offs between objectives by putting them on the same scale. And I'm going to introduce that concept here today and also in the next module in more detail.

We're also going to talk about a method in the next module called SMART, which is Simple Multi-Attribute Rating Technique, which is used for making trade-offs amongst multiple objectives. And there are a couple of other methods that are commonly used, analytic hierarchy process, goal programming, that are used to deal with trade-offs amongst multiple objectives.

Now, when we have multiple objectives but we also have uncertainty, what we need to do is actually modify techniques that apply to multiple objective problems with no uncertainty. Or perhaps what we can do is take a multiple objective problem and transform it in some way to make it into a single objective problem by combining multiple objectives into a single measure. If we can do that, combine multiple objectives into a single measure, we can use things like decision trees with what we would call multi-attribute utility, which is basically a combination of multiple objectives into a single measure. We could also use the methods that are available for multiple objective problems like SMART analytic hierarchy process with probabilities.

When we get into the realm of linked decisions, repeated decisions over time, our options are less. The idea of linked decisions is that we're going to make a single decision, but we're going to make it multiple times. So an example would be every year I have to set duck harvest regulations. I'm making the same kind of decision, but I'm making it every year. So that's an example of a linked decision.

When we have linked decisions, if we don't have uncertainty and we have just a single objective, we can use dynamic programming. When we do have uncertainty, we can use tools like stochastic dynamic programming, or also adaptive stochastic dynamic programming.

Now, when we have linked decisions with multiple objectives, there are very few tools available to us. And so in most cases, what we would try to do is combine multiple objectives into a single objective in some way. And we're going to talk at length about how to do that in the next module on multiple objective problems.

So with that overview, let's talk a little bit about single objective tools. In a deterministic setting, so in a setting either where we don't have uncertainty, or where we have uncertainty that we think is low enough that we can afford to ignore it, we have a number of alternatives available for finding optimal solutions. Imagine a situation where we have a single, continuous decision variable like a harvest rate, the amount of an herbicide to apply, the size of a biocontrol release, something like that, some kind of a continuous decision variable.

We can predict the outcomes, the objective, as a function of the decision variables. So the idea here would be that we would predict, say, the amount of harvest of a species based on some harvest rate, that harvest rate being the decision variable. And there are a number of methods for finding the solutions once we have those predictions.

We can do brute force, where we just plug in a bunch of different options until we find one that seems to be the best. We can use closed form solutions. We can use calculus for finding those solutions. We can use numerical solutions, so mathematical search methods. The computer can search efficiently using particular algorithms for finding the optimal solution. We can also use some constrained optimization methods where we're finding some mathematical solution to a problem where we have, say, that single continuous decision variable, but also some kind of a constraint, like budget.

So let's look at an example. This is the example of harvest rate. We're interested in choosing some harvest rate. And so our different harvest rates become our alternatives.

The objective, maybe, is to maximize the sustainable harvest of the population. So in that case, our objective is sustainable harvest, and we're interested in how different harvest rates impact sustainable harvest. And so we might find a situation where we have a parabola like this. And we can then use a variety of different methods to find the maximum sustainable harvest.

We can look at it graphically. And we could then see that the optimal harvest will be at the top of that parabola. That will be the maximum sustainable harvest. We can also use calculus to find the solution in a case like this if we have an equation that describes that parabola.

So that's a situation where we're in a deterministic setting. We're ignoring any uncertainty that we have. Or maybe we don't have any uncertainty.

More commonly, we're in a stochastic setting. We have some uncertainty about the way our system functions, or we have just variation in the functioning of some system, like weather varying over time. So in stochastic settings, we're actually going to be dealing with probability. So we need to predict outcomes and their probabilities both in order to deal with the uncertainty.

So a really classic tool that's used in these kinds of settings are decision trees. So what we're doing here with decision trees is we're basically saying, well, there's some probability that we'll get this kind of

outcome. And there's some other probability that we'll get some other outcome. And what we're going to do is we're going to sum up the expected value of each of those possible outcomes. And then we can identify, using that kind of information, the optimal solution. So let's go through an example.

Here's our example. So in this situation, we have to decide whether to take some protective measures against the arrival of some invasive species. So we don't know, when we make a decision, whether that invasive species is really going to show up in our habitat. Now, certainly if it shows up, that's going to result in an outcome that we don't like. It's not going to be as good. But if we take some protective measures against the invasive species showing up, things might not be quite as bad as if we had done nothing.

So what we can do is we can diagram this as a decision tree. So what we're going to do is we're going to link the decision action that we take with the objective outcome. So we have our decision, or action node, indicated with a square, and this is whether or not we take protective measures against the invasive species.

The outcome nodes are represented by hexagons, as we did in the influence diagrams in the consequences modules. And then we have chance nodes, which are basically, in this case, whether or not the invasive species shows up. So let's work through some of these different outcomes.

So we could decide to take protective measures against the invasive species. So then we would be in the top of the decision tree. So we would follow the arrow coming out of our decision node that says yes.

So once we've done that, we then go into a chance node. And then the question in the chance node is will, in fact, that invasive species show up? There are two possibilities, yes or no.

Now, if the answer is no, we would, again, be at the very top of the decision tree. If we take protective measures and, in fact, the invasive species doesn't show up, maybe what that means-- the situation is that we've sacrificed some habitat by taking those protective measures. And so the acres of habitat that we end up with are 70 acres. So that's our outcome.

If we take protective measures and the species does, in fact, show up, then we still end up with some habitat, but much less, 20 acres in this case. Now, if we don't take protective measures and the species doesn't show up-- so now we're at the third outcome from the top-- that's the best situation, because we

haven't sacrificed any habitat to developing these protective measures. And so we end up with 100 acres of native habitat.

But if we don't take any action and the species does show up, that's the worst outcome of all. And we end up with zero acres of native habitat. So there's some trade-offs in here, and there's some uncertainty.

So what we then do is we consider the probability of each of the outcomes and the value of each of the outcomes. So the probability of getting the outcome of 70, at the very top, is 0.65. The probability of getting the outcome of 20, which is the second one down, is 0.35. And of course those add to 1. The probability of no and the probability of yes have to add to 1.

Then what we do is we take the product of the probability and the outcome, so 0.65, and multiply it by 70. And add to that the product of the probability and the outcome for the other direction out of the chance node, so in that case, 0.35 multiplied by 0.2. So if you get the idea of this, why don't you pause the tape for just a second and do the calculations and see what you end up with for the expected value of each of the different decisions that you can take, either to take the protective measures or to not take the protective measures.

See if you can calculate and identify which is the smarter thing to do. So pause the tape now. Do that. And we'll come back.

Hopefully you were able to do the calculations and see that the expected value of taking an action, of doing something to take protective measures against the invasive species, was that you would end up with 50-- your expected number of acres that you end up with this 52.5. That is, the product of the probability that the species doesn't show up and the value that you get if it doesn't show up, given that you've taken protective measures, plus the probability that the species does show up and the value that you get if it does show up, given that you've taken protective measures-- if you add those two together, you get the expected value of that action.

You do the same thing for the lower half of the decision tree, and you get the expected value of that action. Now it's clear from this that you should take protective measures. It seems smarter to do the thing that is going to get you to the point where you're expected to have a better outcome. Now, because of chance, it might be that you take protective measures and, in fact, the species doesn't show

up, and you're worse off than you would have been if you hadn't taken protective measures.

But what you've done by using a decision tree and integrating the uncertainty is you've done the thing that is most likely to lead you to a good solution. And so that's the best we can do when we can't know the outcome of some chance node when we make our decision.

We talked a little bit about the idea of single objective problems. So it might be useful now to stop the tape and spend a few minutes thinking about your work setting and identify some single objective problems in your work setting. So pause the tape here and write down a few ideas of single objective problems that you have encountered in your work setting, or that you can imagine encountering in your work setting.

If you're taking the course with a colleague, you might want to stop and work in pairs at this point. And then we'll come back and we'll talk a little bit more about multiple objective problems.

Maybe you were able to think of some single objective problems. You might have found it challenging. Oftentimes, we find that most of our problems are actually multiple objective problems.

But you might have been able to think of a single objective problem with some constraints, say, how do we optimally allocate some budget to different endangered species projects. That might be a single objective problem with a budget constraint. It's more often the case that we're dealing with multiple objective problems, especially in natural resource management, public natural resource management. We're often balancing the opinions and the needs of many different stakeholders.

So oftentimes, we're in the realm of using and working with multiple objective tools. So let's talk a little bit about some multiple objective tools. And then, as I said, we're going to talk about these in much greater detail in the next module, Multiple Objective Optimization Techniques.

There a couple of different strategies that we use when we are dealing with multiple objective problems. But in all cases, the first thing we do is we simplify the problem as much as possible, and we'll talk about some ways to simplify multiple objective problems using very simple techniques for doing so. Once we've simplified a multiple objective problem as much as we can, then we have to confront the trade-offs in some way.

We can do this with a couple of different methods. We can reduce multiple objective problems to a

single objective problem by combining multiple objectives into single objectives. And we'll talk more about this in the next module, as well. We can treat multiple objective problems as explicit trade-off problems. And we're going to talk about this at great length in the next module.

Or in some cases, what we do is we negotiate a solution. So we just graphically display the different options. And we allow people to compromise from there.

We start by simplifying the problem. And there's a couple of different, very easy methods to apply for doing so. What we're doing is we're looking for dominated alternatives, irrelevant objectives, and even swaps.

Once we've done that, we've simplified the problem as much as we can. Then we have to manage the trade-offs. And as I said, we can do that in a couple of ways.

We can combine multiple objectives into single objectives. We can do quantitative trade-offs. Or we can just display the trade-offs graphically, in some way, to allow for people to deal with those trade-offs through compromise. And we're going to talk about all of those approaches in the next module.

So this would be a good time, again, to stop the video and think about some multiple objective problems in your work setting. Now, you almost certainly have multiple objective problems in your work setting. So take a few extra minutes and, for a couple of those problems, try to write down what the objectives are. And then we'll come back and we'll talk about linked decisions.

Well, let's now talk about the concept of linked decisions. Now, we're going to start by introducing the concept with special case of linked decisions, where we just have two linked decisions, say, over time. And they're different kinds of decisions.

Now, imagine a situation where you're deciding whether to establish a captive facility for some endangered species. If you do that, establish the captive facility, you'll have to bring animals in from the wild. When you do that, when you establish the captive facility, bring the animals in from the wild, you can't know for sure, at that point, whether captive propagation will be successful.

If it's not, then you've put resources, and taken animals in from the wild, and the population has not successfully propagated in captivity. But if it does successfully propagate in captivity, at that point, then, sometime, at least, in the future, you have to decide whether or not you would reintroduce those

animals into the wild. Maybe you would consider different numbers of cohorts to reintroduce.

So the key here is that to understand the smartest thing to do at the leftmost decision node, whether or not to establish a captive facility, you have to understand, already, the smartest thing to do at the rightmost decision node, whether to reintroduce one or two cohorts. So what we do in the case of linked decisions is we start from the end, and we work backwards.

So in this case, we would actually start by solving the decision tree that begins at reintroduce one or two cohorts. We would understand the smartest thing to do there. And then we would use that to work backwards through time. So that's a special case of a linked decision.

More commonly with linked decisions, we think of making the same decision over and over and over again through time. And we think in particular of dynamic decisions. Imagine the situation where we're deciding on setting a duck harvest rate every year.

By taking some action, setting some particular harvest rate, we actually have an impact on that population. It changes the state of the duck population. We also get some return. We get some number of ducks harvested.

That system state, then, has some follow-up impact on the system state in the next year, where we take another action and get another return, and so on through time. So this is an example of a dynamic decision, where what I do now has an impact on the system state, which then influences the system state in the next year, where I'm going to have to, again, decide what to do, and so on through time. And again, what we try to do in these settings is we use-- we think about starting at the end and working backward through time to figure out what's the smartest thing to do given a particular, say, system state.

So we'll use, oftentimes, a technique called dynamic programming. And so dynamic programming is a method for solving these kinds of dynamic decision problems where the temporal dynamics are important. So what we mean by that is what I do now affects the system state. That system state then has an impact on the system state next year. So when I'm talking about duck population, say, the system state is how many ducks do I have in my population, for example.

And what we do is we recognize that each action that we take at each decision point accrues a certain return, so the number of ducks that I harvest this year. But it also changes the state of the system,

which can ultimately affect future returns. And so what I want to do is I want to look at the long-term effect of taking any decision now or at any time in the future.

Oftentimes, when we're in those kinds of dynamic decision settings, we also have some uncertainty. So we have stochastic dynamic programming to deal with those kinds of situations. So in that case, the impact of my action on my system state, or the impact of the system state at this time period on the system state at the next time period, is subject to some uncertainty. And so the optimal strategies that I would choose would anticipate the risk of doing something that leads to an outcome that I don't like because of my uncertainty. So we can use stochastic dynamic programming in a setting like that. And again, Mike will talk about this quite a bit more in the last module in the course.

Another example of this for the development of this idea of stochastic dynamic decisions are adaptive stochastic dynamic decisions. So we use adaptive stochastic dynamic programming. So here, the idea is that we have a dynamic decision. So what we do now influences the future state of the system. We have stochasticity. We have some uncertainty.

Now, the adaptive part is the idea that we have some uncertainty, but we can monitor the system. And we can learn about the system in a way that reduces our uncertainty over time. If we integrate into our analysis the potential to learn, what we can see is that what might be smartest in the short-term is to do something to reduce our uncertainty so that in the long term, we can manage with more certainty. So we can actually integrate the potential for learning directly into our optimization so that we can see if there are some particular decisions that will let us learn now and reap the rewards of knowing that information later on in our decision making process so that our optimal strategies, then, anticipate the benefits that we get from learning.

Now, we've talked about a number of methods that are used for dynamic decisions that are basically common in the rubric of dynamic programming. Now, dynamic programming is actually a method for finding exact optional solutions. It's truly the smartest thing you can do.

Now, sometimes these decision problems become so big that all we can hope for is finding approximate solutions. Some big problems are simply too big to be solved with any of the methods that we've talked about. So in those cases, we often are using things like genetic algorithms, simulated annealing, or other kinds of mathematical search methods where what we're doing is we're finding a very, very good solution, but we can't guarantee that, in fact, it's the optimal solution.

An example of this that people may be familiar with is reserve design. We have many, many different pieces of land that we could buy, and all different combinations of those lands, which we could think of portfolios of actions, are so large-- all those different possible actions are so large that we can't actually evaluate each one individually. And so we might use something like simulated annealing in a case like that to find the approximately optimal solution, so a very, very good solution, but not necessarily the absolute best solution.

So we've talked a little bit about a number of methods for solving a lot of different kinds of problems, single objective problems with and without uncertainty, multiple objective problems. We're going to talk in much greater detail about multiple objective problems in the next module. But the goal here is to solve the problem, to identify the best alternative action as measured by the impact of those actions on our objectives using the predictions from the model. So again, optimization methods are the bit that brings everything else together so that we can say, this seems to be the smartest thing we can do given everything that we now and given where we're trying to go to, given what we care about, what our objectives are.

There are a very wide variety of analytical tools that exist to find solutions. Now, some of them are fairly technical, and oftentimes would require collaboration with scientists or specialists who work with these kind of analytical tools. Now, the key is knowing where to look in the toolbox. And in particular, the key is knowing whether or not you have a single objective problem, or a multiple objective problem, or a linked decision problem, and whether or not you have an important source of uncertainty, or whether you have little uncertainty, or whether you have uncertainty that you can, in fact, ignore.

So the key is really knowing where to look. And then, once you know where to look, you can figure out, at least, the class of decision tools that might be available to you. So with that in mind, let's end with just a little bit of individual work. And let's work through this quiz.

Spend a little time individually thinking about this question. What type of decision problems are the following examples and what tools might be useful for solving them, whether to relocate or start a captive breeding program for the Devil's Hole pupfish, what kind of decision problem might that be and what tools might you use; management of early successional fields on refuges; stocking levels of bass in Adirondack lakes; annual harvest regulations for woodcock; or whether to invest in a study of the costs and benefits of outsourcing. So stop the video. Spend a little bit of time for each of those

examples writing down your thoughts about what kind of decision problem each of those is and what tools might be useful for solving them. We'll come back and talk a little bit about those.

So you've taken a few minutes to think about each of these different examples. And it might be that different people get different answers here. It really depends on how you conceive of the problem. There's not necessarily a single right answer.

But let's take the first example, whether to relocate or start captive breeding of Devil's Hole pupfish. Well, you might think that, maybe, this is a single objective problem, or, say, a single objective problem with a constraint, but it's almost certainly going to be a linked decision problem, because if you start captive breeding, presumably, in the long term, you would ultimately want to do releases of pupfish from the captive breeding population. And you almost certainly have some important uncertainty in that case, because you probably don't know a tremendous amount, or have quite a bit of uncertainty, about whether captive breeding for Devil's Hole pupfish would work. So that might be a case where you would use a decision tree with a linked decision.

The next example, management of early successional fields on refuges in the Northeast. Well, management of fields on refuges is almost certainly going to be a multiple objective problem. There's a lot of different benefits we want to get out of a habitat like an early successional field. We would want to probably maintain populations of a number of grassland species, maybe some other kinds of insect species. Perhaps we might want to have some kind of user benefits, like the ability to do bird watching on those early successional fields. So almost certainly that'll be a multiple objective problem. Now, depending on the kind of benefits you're trying to get, you may have some substantial uncertainty as well.

The next example, stocking levels of bass in Adirondack lakes-- you might think of this as a single objective problem, potentially, although it may not be, depending on the different user groups involved and so on. But you might think of this as a situation where you're trying to maximize the sustainable yield of bass in those lakes.

The annual harvest regulations for woodcock. Well, that's almost certain to be a dynamic decision, a linked decision that goes on over time. And you'll almost certainly have some uncertainty about woodcock populations and how they'll respond to harvest, or other things that are going on in the environment. So that's almost certainly going to be a problem that you would want to take on with

adaptive management, and as I said, Mike will talk more about those ideas in the last module in the course.

And then our last example was whether to invest in a study of the costs and benefits of outsourcing. Now, this is a special case of a linked decision. Our first decision is do we invest in the study of the costs and benefits of outsourcing. And then once we invest in that study, do we, in fact, go ahead and do the outsourcing?

This is a special case of a linked decision where what we're interested in is the value of the information that we get by investing in a research study. And Mike is going to talk about that kind of decision at some length in the last module in the course.

So hopefully, you've found some similar ideas to what we talked about. But you may have thought about each of these quiz examples a little bit differently. Again, it really depends on the particular situation that you're in. It depends on how many users are involved, how many objectives you have, and so on. And again, the big message is the decision tools that you use, the optimization tools that you use, have to be selected to fit with the other elements of the problem, the objectives, the models, the number of actions you have, and so on.

So what comes next? Well, we've talked about problem definition. We've talked about setting identifying-- setting objectives. We've talked about identifying a good set of actions or alternatives. We've talked about building models to consider the consequences of taking different actions. And now we've just started talking about trade-offs.

In this module, you've used optimization. We've talked about the use of optimization to identify the action, among a set of alternatives, that best achieves your objectives given what you understand about the system, which you have codified in some model. So that is the goal in optimization. And the goal then, at the end, is to be able to say, this is the smartest thing that we can do given where we want to go, given our options for getting there, and given what we do know and what we don't know about the system that we're trying to manage.

[MUSIC PLAYING]