US Fish and Wildlife Service | Making Decisions with Multiple Objectives _Part 2_

You've just seen how we use the smart method to deal with the trade-offs between multiple objectives by normalizing it and then weighting within objectives and then summing across alternatives to develop a single score for each alternative. You're going to get a chance to practice that in a few minutes.

But before we do that, let's talk a little bit more about the concept of weights and how do you assign weights. It's important to say that the weights are actually a part of the statement of the objectives. So in multiple objective problems, it's not just what is the list of the objectives but also how relatively important are they.

And it's key that we work with the decision maker to develop the weights. Just as Mike talked to you about in the objectives module, we work with the decision maker to understand that decision maker's values and therefore objectives. We need to work with the decision maker to understand how relatively important those different objectives are.

So weights have to be elicited from the decision maker. And there are a number of different methods that can be used. One could directly elicit weights from the decision maker. But there's a challenge with that.

Imagine a situation where you had two different cars that you were concerned by. Say you had two different objectives. You had an objective that had to do with costs and an objective that had to do with fuel efficiency. It would be difficult, without knowing the values for the cost and the fuel efficiency for the two different cars, to just abstractly decide how relatively important those two things were.

So for example, if the difference between the cost of the cars was $5,000 but the difference between the fuel efficiency was only two miles per gallon, you might weigh those objectives differently than if the difference in the cost of the cars was say $1,000 and the difference in the fuel efficiency was say 10 miles per gallon.

So it's important often to recognize the actual values before one develops weights. Direct elicitation sometimes works. But more often we use a method called swing weighting. And we're going to talk about how to do swing weighting.

The idea of swing weighting is that it recognizes that weights are context dependent. It depends on the actual range of scores for the given objective, how important that objective is to you.

Multiple objective methods require, of course, converting all these objectives to the same scale, these kind of quantitative trade-offs that we've talked about, and then putting in weights on those normalized scores. So we use swing weights to develop those weights.

So the idea here is that the relative weight or preference for an objective depends upon a particular value. So our preferences among objectives are context dependent. It's hard to think about weights in the abstract.

So swing weights use a swing or the range from worst to best consequence values across the actual alternatives to develop the context specific weights. So we go through a series of steps to elicit swing weights.

The first thing we do is we identify the worst case scenario or benchmark scenario for each objective. So this worst case scenario is an imaginary alternative that scores worse on each of the different objectives. So it's the most expensive. It has the least environmental benefit. It has the most disturbance, for example.

Then we develop a number of additional scenarios in which one objective is swung from that worst value to its best value. And so basically what we're thinking then about in asking the question is, if just one of the attributes could be moved from the worst value to the best value, which would it be?

That imaginary alternative becomes the top-ranked alternative. And the objective associated with that alternative becomes the top ranked objective. And we would do that for all of the different objectives, developing a scenario in which that objective was swung from the worst case to the best case.

Then once we've done that, we develop, rather than just weights, actual relative scores in terms of the attractiveness of those different scenarios. And then we normalize the scores. So let's go through an example.

Let's imagine a situation where we have decided to pick a puppy. We've decided to get a puppy for our pet. And now we actually have to decide which puppy that we are going to buy.

We have three different objectives. We like the personality of large dogs. So we think that a good attribute for that would be the pounds that we expect the dog to weigh as an adult. Because we think that would be a good measure of whether or not that dog will have a large dog personality, which we like.

We are not so much in favor of long-haired dogs. Because there's annoying hair all over the place. So we measure that attribute in terms of the inches of hair that that dog would be expected to have. And then we also care about the friendliness of the dog. And so we would maybe, in that case, use a constructed attribute.

In this table, what we see is that we must have a number of different alternatives that we're considering. We don't actually see those alternatives in that table. Instead, what we see is the range, the worst and the best in each case.

So the worst is a 20-pound dog with six-inch long hair and a friendliness score of only one. So you can think of that dog as the small, hairy, mean dog. Now, we also have a best case for each of the different objectives.

The best case is a 65-pound dog that only has one-inch long hair and scores a four on our constructed scale for friendliness. So that's the best case. That's the big, short-haired, friendly dog.

We start by developing a benchmark scenario. That benchmark scenario is the worst case. Then we develop three other scenarios. Because we have three objectives. And each of those scenarios is associated with a different objective where we swing the score for that single objective from the worst case to the best case.

Scenario one is still a dog that isn't very friendly. It scores a one. It's still a dog with long hair. It has a six in the length of the hair. But it now is a much bigger dog. And it scores well on the large dog personality, 65 pounds.

So now note that this is not an actual alternative. It's a pretend alternative to help us get at these swing weights. We develop these three different pretend alternatives where one objective is swung from its worst to its best.

Now we have four different alternatives, pretend alternatives. We have the benchmark alternatives, one, two, and three. And what we do is we use these pretend alternatives to have a conversation with our decision maker where we say, now, clearly the benchmark alternative is the least favorable. That's a four. We're going to rank that fourth.

But which of the other three is actually the best alternative? If you could choose one objective to swing from its worst case to its best case, which would you swing? And our decision maker then, thinking about that, could rank those different pretend alternatives or those different scenarios, one, two, and three.

And so maybe, in this case, our decision maker thinks he'll give a one to the not annoying hair objective. Because the swing from six to one really means a lot to that decision maker. Maybe our decision maker gives a two to the friendliness objective, the scenario associated with the friendliness objective where we now have, yeah, the dog is still small and hairy but at least it's friendly. And then last the large dog personality is ranked third.

Now we know the importance of those objectives in terms of the weights. But now we actually have to get scores. Yes, one is ranked better than two. But how much better is it?

So we would ask our decision maker, if you gave 100 points to scenario two, the dog that has the short hair, how much then would the scenario that you ranked second score? And in this case, our decision maker says 60. It would be a score of 60 compared to 100 for scenario two. And for scenario one, our decision maker thinks that that's a 40.

So then we take these scores of 100, 60, and 40. And we normalize them so that they actually sum to one total. And how we do that is we simply take say 100 and divide that by the sum of 100, 400, and 60. And that gives us a normalized score of 100.

So now our weights come to 100. Or if we divided by 100, they sum to one. And they recognize, not just in the abstract how important are these concepts, but given the swing across the alternatives that are in the mix, how important are these objectives.

So this is how we develop swing weights. And those swing weights then we can use in doing and use in our smart method to weight the different objectives once we've normalized them.

Let's talk a little bit about the concept of facilitating a negotiated solution. So the idea here is that we haven't been able to combine multiple objectives into a single objective. And we haven't used the quantitative trade-off techniques.

Sometimes what's actually most powerful is to simply graphically display the trade-off between the different objectives to facilitate compromise amongst different parties or to facilitate even a single decision maker figuring out what makes most sense. So the idea here is to facilitate looking at trade-offs directly.

We can see this graphically with a problem that I worked on, American Shad Management in Fish and Wildlife Service Region five. So what we have here is we have a number of different options for American Shad management.

Each of these alternatives is actually a bundle of actions. So we could think of these as different portfolios. It's a bundle of different actions that has some overall population benefit, which we'd like to maximize, and some overall cost, which we'd like to minimize.

So each of these points on this graphic represents a different alternative with some cost and some benefit. We can imagine then that there's some option that's the least costly. There's some option that has the highest environmental benefit.

Along this frontier, what we call this efficient frontier, basically, the options that are very close to this line, those are good options. And they're all options that would be reasonable to choose.

There are also a number of inefficient options. Why would we ever choose an option that has the same population benefit but actually costs a lot more money? So the options over to the right will be inefficient.

But once we can graphically display this efficiency frontier, we can allow a conversation about the negotiation of trade-offs. So a decision maker might think, well, I'm willing to spend a little bit more money to get a certain amount more population benefit.

And by being able to see that graphically, that's a powerful way to negotiate trade-offs. So this is another option for facilitating the direct trade-offs between different objectives and multiple objective problems.

So let's briefly talk about the concept of uncertainty. Now, multiple objective decision techniques are not extremely powerful at handling uncertainty. But what we can do is we can do sensitivity analysis.

So the idea here is that we recognize that we have say some uncertainty in the scores that exist in our consequence table. And we could repeat the analysis while varying those scores. And we can check to see whether our decision is robust to that uncertainty. That is, if we vary the uncertainty, do we still end up with the same recommended alternative?

If we do, we feel confident that this uncertainty doesn't impact our decision. And so it's safe for us to ignore it. If it does, then we have to think about making decisions under uncertainty perhaps. And Mike is going to talk about those concepts in the last two modules in this course.

So with that, let's go back to the Rolling Thunder Prairie example. You've worked through and simplified it as much as you can to this point. So now you're actually going to use the smart method to do the quantitative trade-offs.

Now, you might want to develop your own sense of what your weights are. Or you might want to try different weights to see how different weights influence the outcome. And once you finish the example, come back to the video. And we'll work through it together.

Hopefully, you've had a chance to go through the skill check for this module. And using the simplified consequence table that you developed in the first part of the skill check, you can now actually use that to practice the smart technique.

This is my spreadsheet. And what I have here is I have my five remaining objectives. And I have my three remaining alternatives after I've simplified as much I can. And I've put into these cells the different actual values for each of the alternatives under each objective, the real values.

So this is the real cost of the alternative, the expected number of neighbor complaints and so on. So these are the real original measurable attributes.

Now what I'm going to do here is I'm going to translate those original scores into the normalized scores. So the idea here is that I translate the alternative that does the worst on a given objective will get a score of zero. The alternative that does the best will get a score of one. So I've filled in some of these. But let's do a couple together for the practice.

The spring burn alternative under the costs objective, I want to translate this score of 10,000 into a normalized score. So if I scroll over here a little bit, I can see the formulas I'm going to use. I want to normalize given that I want to minimize the objectives. So I'm going to use this value right here.

So I'll put that equation into this cell here. So it's going to be one minus the value. By the value, we mean the actual score so this value right here. If I click in that cell, it'll give me the value C8. And then I subtract from that the minimum across this objective.

And so what I'll do is I'll use the formula min in Excel. And then I'll use C8 colon E8. And that formula, min parentheses C8 colon E8, will give me the minimum value across those cells.

And then I divide that by the max. So I want the max here. So I'll use this Excel formula for max, C8 colon E8, and subtract from that the min, again, C8 colon E8. Now, what I see is that the cost of $10,000 translates into a normalized value of 0.357, which is somewhere between the most expensive, which gets a zero, and the cheapest of the alternatives, which gets a one.

Let's do one more formula here in this cell. Now, in this cell we're trying to maximize the positive effects on listed plants. So we're going to, because we're trying to maximize, we're going to use this formula here. So let's put that in.

That's going to just be the value in this cell here. We click in the cell minus the min, again, from C10, in this case, to E10 divided by the max from C10 to colon E10 minus the min from C10 colon E10.

And I get a one there, which seems sensible because this score of 10 is the highest for the three alternatives over this objective. So I should be getting a one there. So that looks good.

I now have my normalized consequence table. And I'm going to scroll up. Now, what I have here is my consequence table where I'm actually multiplying by the weights. Well, now in my case, I've started here with equal weights across the five objectives. So they've each got a weight of 20%.

So what I need to do now is I need to multiply these normalized scores by the weights. So I'm going to multiply the score for my cost objective that's normalized multiplied by the weight for that objective. So I multiply C16 by C24.

Now I do the same thing for the mowing alternative. I multiply the cost of that alternative, the normalized cost, multiplied, again, by the weight for that objective. And I do the same thing here, the normalized cost multiplied by the weight.

What I see is that I have these normalized weighted scores for each of my objectives under each alternative. So now I've put everything on the same scale. And I've weighted that scale. So it's reasonable for me to sum the weighted scores together across the alternatives.

And given these weights, what I see is that the grazing alternative is the preferred alternative. It scores quite a bit higher than the spring burn and the mowing alternative. And we can see that here in this table.

We can then maybe do some sensitivity analysis. So let's try doing some sensitivity analysis in the table. So we'll go back up to the original table here.

Say we actually have some uncertainty about the effects on listed plants. So maybe we think that, in fact, the effect on listed plants is much lower than we originally thought it would be for the spring burn. So maybe we think it could be as high as 10. But maybe it could be actually as low as five.

So we're going to change this from a 10 to a five temporarily. And we wonder whether that should change our selection of our alternatives. And in fact, it doesn't. In this case, our grazing alternative is still preferred.

And that seems pretty sensible in this case because actually the grazing alternative was already preferred despite the fact that it did not do very well on the effects on listed plants. Our decision we would say is robust to the uncertainty we have in the effect of the spring burn on our listed plants. It could be as low as five and as high as 10 and our decision would still be the same.

What if we actually had some uncertainty in say the effects of the grazing on butterflies? So maybe we thought that this actually could have a very poor effect on butterflies, grazing. Maybe we think it could be as low as 0.01. Well, will that change our decision?

Now in this case, it actually did change our decision. If things are as bad as that for our effects on butterflies, we would actually take a different decision. We would decide to do mowing rather than grazing. So if we had uncertainty to this degree, we might actually consider trying to do a study to see if we could resolve that uncertainty before we went forward.

You had a chance to apply the smart method. And then you also had a chance to think about sensitivity analysis. So sensitivity analysis is what we use when we have uncertainty in our multiple objective problem. And the key question we're asking is, depending on the uncertainty and depending on the resolved value of the uncertainty, would it change our decision?