US Fish and Wildlife Service | Making Decisions Under Risk _Part 2_

[MUSIC PLAYING]

So how do we deal with risk? How do we deal with uncertain outcomes and making decisions when we have that kind of situation? Well, there's a number of strategies that have been developed in the field of decision analysis over the last, I'd say, 60 years. And I'm going to give you a little example of all of these. And these really differ all in how the decision maker wants to approach risk.

The first one is optimization where what you're trying to do is maximize or minimize the expected value of your objective. The actual performance may be more or less owing to the uncertainty. But what you're trying to do is seek the highest expected value.

So that's one approach. Sometimes we call that an expected value decision maker. They're just going to take the action that has the highest average value over the sort of uncertain events that could occur.

Another approach that's quite common is called mini-max or maxi-min. And what you're doing is you're either guarding against the worst loss or optimizing a minimum gain. So you can minimize your maximized loss. Or you maximize your minimum gain.

So it's kind of like a worst-case scenario sort of thing. You're trying to guard against the worst-case scenario. And I'll show you an example of that in a second.

There's another strategy that's called satisficing and robustness. And what this does is focus on attaining a minimum performance requirement. So you say, well, what I really want to do is this minimum level. If I could attain that, I'll be happy. Anything above that doesn't really matter to me. I [? want to ?] have this minimum performance requirement.

Beyond that, what I really want to do is seek a solution that guarantees that minimum performance over the greatest range of uncertainty. So what I'll do is try to find a solution that robustly gives me my satisfactory goal. I'll show you an example of that.

And then the last technique that's used is called utility. And that's one that we'll spend a fair amount of time on in the next few minutes. And what you do is you translate your return into a utility scale that expresses the risk attitude of the decision maker. And there's been a lot of theoretical work done on utility and how you can use the idea of utility to express the risk adverseness or the risk seeking value of the decision maker.

The point is that all of these approaches differ essentially in the decision maker's attitude towards risk. Now let's look at a short little example that demonstrate all of these at once.

So I have a little table here that has four choices. A decision maker has to choose between-- let's say they're for envelopes, A, B, C, or D. And there's four possible outcomes. The outcomes are a gain of zero, a gain of 10, a gain of 20, or a gain of 30.

So you might think of this as there's some amount of money in each of these envelopes, I guess. And there's a probability of it being $0 or $10 or $20 or $30. And you need to decide which envelope to choose.

So for instance, envelope A, there's a 25% chance that there's nothing in it, a 25% chance it's $10, a 25% chance it's $20, and a 25% chance that it's $30. Envelope D, there's a 20% chance it's nothing and an 80% chance that it's $20. So there's some uncertainty, right?

You don't know what's in the envelope before you choose it. And there's different probabilities. Well, how are you going to choose which envelope is best? Let's look at a couple different strategies.

The expected value, the column that says EV, is the expected value. So that's the average payout of the envelope. So envelope A, because there's an even chance of $0, $10, $20, or $30, the expected value is $15. In envelope B, there's a 25% chance it's $10, a 25% chance it's $30, and a 50% chance it's $20. The weighted average of those is $20. The weighted average for envelope C is $23. The weighted average for envelope D is $16.

If we're an expected value decision maker, if we're trying to optimize, then we might choose envelope C because it's got the highest expected value of $23. Suppose on the other hand, we want to take a maxi-min approach. We want to look at the minimum amount that'll be in an envelope and choose the envelope that's got the maximum minimum amount.

Because it sounds silly to say it, let's show. Envelope A could contain $0. The worst case for envelope A is that it's got $0. OK, so in this column in this table that says Min, the minimum amount you could get from envelope A is $0.

In envelope B, there's no chance of it being zero. The minimum amount you could have is $10. So that's what's in the minimum column. C and D have a nonzero probability of nothing in the envelope. And so their minimum value, the worst case, is $0.

So what we might want to do is choose the best worst case. All right, so the best worst case is envelope B. So another way to say this is in envelope B, you're guaranteed to get at least $10. And that's the highest guaranteed minimum of any of the envelopes.

OK, so you're guarding sort of against the worst case in this particular strategy. That's reasonable. Maybe that does express the risk attitude of a particular decision maker. People might be quite different on this.

Let's look at the robust satisficing. Let's say that there is a music album that you want to download from iTunes. And it's going to be $15.

What you really care about is having at least $15. And otherwise, you don't much care. Because if you have $15, then you can get this album. So you might chose the envelope that has the highest probability of giving you at least $15.

And so this last column, the P greater than 15, this is the probability that the amount that you get will be at least $15. So in envelope A, there's a 25% chance that you'll get $20 and a 25% chance you'll get $30. So that's a 50% chance that you'll get something greater than $15.

Envelope B, there's a 75% chance. Envelope C, there's a 70% chance. And in envelope D, there's an 80% chance that you'll get greater than $15. And so if you want to robustly reach your satisfactory goal of $15, you should choose envelope D because it's got an 80% chance of reaching that goal.

So the point here is that these are all legitimate ways to approach risk. They just express different risk attitudes. And that's part of the value judgment of the decision maker.

Let's look at risk attitude another way, considering the following wager. Suppose you win $500 with probability of 0.5 or lose $500 with probability of 0.5. Would you pay to get out of this wager?

Suppose you were forced into this wager. Somebody held you at gunpoint and say, all right, I'm going to flip a coin. And if it's heads, you win $500. And if it's tails, you lose $500.

This is a strange way for somebody to rob you or not rob you. But would you pay this gunman to say, yeah, let me out, let me out of this wager? And if so, how much?

Or another way to think about this is maybe you'd pay to get into this wager. Would you buy a ticket that had this kind of gamble associated with it? And how much would you pay for that ticket?

So I'd like you to do is pick one of these, whether you would pay to get into this wager or pay to get out of this wager. And I want you to draw a decision tree that would represent this question, this decision. OK, so you've got some space in your book on page I-5 where you could draw a decision tree that might represent the decision you're faced with in one of these two settings. And so why don't you take a few minutes, pause the video, and sketch out a decision tree that represents this decision. And then think a little bit about how you feel about that particular wager, whether you'd pay to get in or out of it.

I've asked you to draw a decision tree that represented this sort of strange kind of wager. Let's look at the first one. Would you pay to get out of this wager?

There's two branches here. If you pay to get out of this wager, and if the answer is yes, I'd pay to get out of this wager, there's a sure outcome. Your return on this is negative x. Whatever amount you would pay to get out of this wager. Let's say if it was $5 or $50 or $100, whatever you would pay. That represents a loss to you. It's a known loss, right?

There's no uncertainty about it. You pay that money. You're out of the wager. That's it.

If you don't pay to get out of the wager, then the coin gets flipped. It's heads or tails. And you either win $500 or lose $500. And so these are the possible payouts.

Now, it turns out that the way we set this up is there's a 50-50 chance of winning or losing if you're in the wager. And so the expected value if you pay to get out of the wager is negative x. If you stay in, the expected value is $0.

Now an expected value decision maker would just stay in, right? The expected value is $0. So it wouldn't make sense to pay any dollars to get out of the wager because that's a guaranteed loss verses an expected value that's even.

But it might be quite rational for you to pay to get out of this wager, largely because you're concerned about that $500 loss. And you want to avoid, in particular, that $500 loss. And it may be worth taking a smaller loss, a guarantee of say $50, to guard against the uncertain loss of $500. So that would be the logic behind paying to get out of this wager.

What about paying to get in? Let's take a look at that decision tree. That decision tree looks a little bit different.

If somebody is offering you the opportunity to be in here, that you're not forced into it, you'd have to pay to get into this wager. If you don't pay to get in, then you're not out any dollars. And you don't face any risk. So the expected outcome is $0, right?

But if you pay to get in, then when you flip the coin, if you win, you win the $500 less whatever you paid to get in. So $500 minus y, let's say. If you lose, you lose the $500 and you lose the amount that you paid to get in. So say you paid $50 to get in. You paid $50 for this ticket. Then if you win, then your gain is $450. If you lose, your loss is $550.

So we do the probabilities. The probability is 50-50 either way. The expected value, the expected loss is whatever you paid for the ticket. The expected loss is the $50, or whatever you paid for that opportunity to be in that wager, versus an expected win, or loss of $0, if you don't buy this ticket. Well, under what circumstances would it make sense to do this?

Well if you're really interested in that $500 gain, you may be willing to pay money for that ticket on the chance that you'll get that gain. And you're not so worried about the loss. So you're really seeking the opportunity for that win. And you're willing to take something that's got an expected value that's negative for the off chance of that win.

So the point here is that whether you'd pay to get in or pay to get out of that gamble expresses something about your risk attitude. And so there's some terms that come out of decision analysis that refer to risk attitude. The terms are risk-averse and risk-seeking.

So a person who is risk-averse, if you would trade a gamble for a sure amount that's less than the expected value of the gamble. So in other words, you would choose the sure amount that has a lower expected value rather than the gamble that has the higher expected value. Well an example of this, that virtually all of us do, is buying insurance.

If you buy car insurance or you buy health insurance or you buy homeowner's insurance, what you're doing is you're trading a gamble. If you don't own the car insurance and you get in an accident, then you could be out a lot of money. There's a gamble. There's an uncertain event there.

You might be fine. You might not get into an accident. But if you do get into an accident, then you might be out a lot of money. And so that's a gamble, right?

But if we buy the insurance, then essentially what we're saying is that the financial outcome at least is certain. I'm not going to be out any money if I get an accident because the insurance covers it. But I've had to pay the premium on the insurance. And you can be assured that the expected value of buying insurance is lower than the expected value of not buying insurance.

In a sense, the insurance companies are charging you a higher premium then the expected loss, the loss that they expect from you. Otherwise, insurance companies wouldn't be in business. So we're trading a gamble for a sure amount that's got a value that's less than the gamble. We're being risk-averse. We're guarding against a downside risk.

A risk seeker is somebody who would trade a sure amount for a gamble that has a smaller expected value, but the chance of a larger payout. So when you buy lottery tickets, again, you can be assured that this is a losing proposition when you buy a lottery ticket in the sense that the expected value of that lottery ticket is lower than the expected value of not owning that lottery ticket. Because whatever you pay, the $2 or $5 you pay for that lottery ticket, the expected value of that lottery ticket is quite a bit less than the amount that you paid for it. But you're trading a sure amount, that $5 that you used for that lottery ticket. You're trading a sure amount for a gamble now.

The lottery ticket's a gamble. And the expected value of the gamble is less than the sure amount that you paid for it. Well, essentially what you're doing is you're looking at the upside potential here. And you're seeking that risk.

Well interestingly, say what you will. But it's not irrational for somebody to buy insurance or somebody to buy lottery tickets. In fact, it's not even irrational for the same person to buy insurance and also to buy lottery tickets because the context and the scale in which those things are happening is quite different. The point is, again, that the risk attitude is a values judgment on the part of the particular decision maker.

And I think with regard to the decisions we make for natural resource management, we need to be able to think quite clearly about what kind of risk attitude we're taking and what values that represents. So let's look at a case that-- it's a simple case. It's a little bit made up. But at least it's sort of in a conservation context.

Suppose you are the manager of a fish hatchery. And this fish hatchery produces fry that are used for some stocking program for endangered species. Suppose somebody comes along and says, look, we've got this new technology that you could put into your hatchery.

It's new. It hasn't been tested thoroughly. But it might work. It really might enhance production at your facility.

So what's the decision that you're faced with? Well, if you don't adopt this technology, you know that you've been able to produce 40,000 fry from your hatchery. And you have every expectation that if you don't change the technology in the hatchery, you will continue to be able to produce fry at that rate.

On the other hand, you think that if you adopt this technology, and it works, then there's a chance you'll increase your production to 70,000. But if it doesn't work, and it disrupts all of the fine tuning that you've done over the years to get this hatchery working really well, then your production might drop to 10,000 fry. Should you adopt the technology?

You're making a decision in the face of uncertainty. And notice that also, these two branches, one of them is a gamble. Adopting the technology is a gamble because you're not sure of the outcome. Not adopting the technology is a sure thing, the way we've drawn this at least, because you know what the production is likely to be. So you're choice is between a sure thing and a gamble. And how do you make that decision?

Well, let's add some detail to this. Let's suppose that the probability that this technology works is 80%. The probability it doesn't work is 20%.

Well, let's think about the expected values here. Well, the expected value of not adopting the technology is the 40,000. We know what the sure outcome is of not adopting the technology.

If we adopt the technology, we take the weighted average of 70,000 and 10,000 weighted by these probabilities, the probability of success. 80% times 70,000 plus 20% times 10,000, that gives us an expected value of 58,000. So if we were an expected value decision maker, all we cared about was the expected performance, the expected number of fry that's produced. The expected production's higher if we take this gamble. So an expected value decision maker, who is risk-neutral, would take the gamble.

What about if we're not risk-neutral? Let's think a little bit about this concept of utility. So the idea of utility is maybe we don't care about the outcomes in a linear way.

All right, so you've got a graph on page I-7 of your notebook that shows the hatchery production on the x-axis. And that goes between 10,000 and 70,000, which is the range of production that we're thinking about. And on the y-axis is something called utility. And this is meant to represent how much we care about the outcome.

Well, if we're risk-neutral, then we care about the outcome linearly in relation to hatchery production. And so the production of 40,000, since that's halfway between 10,000 and 70,000, it's also halfway between a utility of 0 and 1. So that's got a utility of 50%. And so if we're risk neutral, we just think we can average the hatchery production. And that's the same thing as averaging sort of a linear utility function to get our answer.

Well, let's think about this another way. Let's say that 40,000 represents a pretty satisfactory production for you. That's meeting all the needs of the recovery plan. Everybody that wants fry is getting some fry to reintroduce into the wild.

And so that 40,000 works pretty well. Sure a production of 70,000 would be a little bit better, but not that much better. Production of 10,000 fry, it would be terrible because that wouldn't meet any of your recovery plan goals, et cetera.

And so we might say that on some sort of utilities scale that makes a values judgment about how much you care about those outcomes, we might say that on a scale of 0 to 1, where that 10,000 production is 0 and 70,000 production is 1, maybe the 40,000 production is 0.9. It's almost all the way to 1, but not quite. It's 0.9. It's really quite high.

Then we would draw a curve. This risk-averse curve that's concave down. And you'll need to sketch this in in your notebook. That's concave down and represents how we would value different outcomes in a way that's not linearly related to hatchery production. OK, so that might be a risk-averse decision maker.

Let's think, is there a setting that you might be risk-seeking? Well let's suppose on the flip side that this 40,000 isn't meeting your management goals at all. It's not enough to achieve recovery. It's not enough to fill the needs of all the partners that want to have some fry to reintroduce into the wild.

But 70,000 would be. So you're not really happy with the 40,000 production. And going down to 10,000, ah, what does it matter? We're not meeting our goals anyway. So that's not a huge loss. But getting the 70,000 would be an extraordinary gain because the implications are much, much better.

So in that case, we might draw a very different curve. We might say that the utility of production of 40,000 is only 0.1. You know, it's really quite small. It's not much different than the utility of 10,000 fry in production. But the utility of 70,000 is 1. It's very high.

And so we've got a concave up curve now, a risk-seeking curve that represents a very different value stance with regard to risk associated with these different outcomes. So we can use this idea of utility, this idea of putting a nonlinear value on these outcomes. And we can use that in the kind of calculations we did in the decision tree to show us how to integrate our risk stance into our analysis in the face of uncertainty.

A risk-averse decision maker would look at the 70,000, the 10,000, and the 40,000. And in this case, OK, the 70,000 is always going to have utility of 1. We're scaling it so that's the best value. It's got a utility of 1.

The 10,000 is the worst value. That's always going to have the utility of 0. But in this case now, the risk-averse decision maker that had the concave down risk utility curve, the utility placed on 40,000 is 0.9.

So now instead of taking the expected value of production, let's take the expected utility for this decision tree. So the expected utility for the upper curve, for the gamble, is 80% times the utility of 1 plus 20% times a utility of 0. That's an expected utility of 0.8.

Whereas on the bottom, the sure thing, that has a utility of 0.9. The expected utility is 0.9. In this case, what would the risk-averse decision maker do?

Well, they would choose not to adopt the technology because they're worried about the loss of production going down to 10,000 fry. OK, so that's how we could think about this from a risk-averse setting. So the point here with risk-averse utility is that your trading gamble with an expected value of 58,000, right?

Remember that the gamble of adopting the technology had an expected value of 58,000. You're trading a gamble, with an expected value of 58,000, for a sure thing with an expected value of 40,000. You're trading a gamble that's got a higher expected value for a sure thing with a lower expected value. That's the hallmark of a risk-averse decision because you're guarding against the uncertainty associated with the gamble.

Suppose this technology is actually unproven. So let's flip the probabilities around. The same sort of setup except now there's an 80% chance that the new technology won't work and only a 20% chance that it will work.

So if we do the expected value calculations, the expected value of adopting the new technology is only 22,000 because it's 80% times 10,000 plus 20% times 70,000. That comes out to 22,000 expected production. Whereas if we don't adopt the technology, again, we have the expected value of 40,000. In fact, it's a sure thing that the production is going to be 40,000.

Now in this case, a risk-neutral decision maker would say, no way. I'm not going to adopt this technology because the expected production is lower than if I don't do anything. But what about the risk-seeking decision maker?

So in a risk-seeing setting, the utilities are 1 and 0 on the 70,000 and 10,000 as before. But the utility on the 40,000 we said is 0.1. Remember this decision maker's really unsatisfied with the 40,000 production because it's not meeting any of the recovery plan goals and doesn't really distinguish the 40,000 from a 10,000 production. Both of them are essentially failing to meet what's needed. But 70,000 would be great.

So here if we take the expected utilities, the expected utility up top is 0.2 times 1 plus 0.8 times 0. That's an expected utility of 0.2. But the expected utility of the 40,000 fry is only 0.1.

So here's a case where the risk-seeking decision maker would choose the gamble over the sure thing, even though the gamble has a lower expected value than the sure thing. So you're trading a sure thing with the value of 40,000 for a gamble with an expected value of 22,000. You're choosing the gamble. And it's got a lower expected value.

Why would you do that? Because you're seeking the upside gain of the 70,000 fry. You're seeking that outcome that meets your goals, even though it's got a low probability.

That might be a rational thing to do in certain settings. OK, so the point is that we can use this idea of utility to capture the risk attitude of the decision maker. So this idea of utility is a pretty powerful tool because it's a tool to really engage with the decision makers and think about what their risk attitude is. And I think that's helpful. I think we need tools for dealing with risk, dealing with uncertainty, and helping decision makers think about the values judgments that are associated with uncertain outcomes.

There's another kind of utility that's actually used quite a bit. And I think it's important for us to think about. And this is sometimes called temporal discounting.

The idea is there's a special form of utility that's used to deal with returns that occur at different times. So for instance, we could ask the question how would you value a $50 gain today versus a $55 gain in 10 years? So I could say to you, you've got a choice. I'm either going to give you $50 today. Or I'm going to give you $55 in 10 years.

And let's suppose you trust me. And you actually think I'll deliver on that promise. OK, which would you choose?

Well, in order to think about that, I mean you might actually turn down the $55 and say, no I want the $50 today. And what you're really doing is you're thinking about how the utility of that gain changes depending on when that gain occurs. So the timing of the gain is an important part of how you value it and how you assign utility to it.

Often we discount future returns. In economic applications, I mean there's all kinds of theory. I mean this is very, very deeply ingrained in economic decision analysis. In economic applications, we calculate the current value of future returns by applying usually the expected inflation rate.

So we could say, well, what's the inflation rate expected to be if I had $50 today and invested it at the current inflation rate? What would it be worth in 10 years? And that's how I can gauge.

And if it's going to be worth more than $55 in 10 years, then it makes sense for me to take $50 today rather than $50 to $5 in 10 years. So we're basically trying to find a way to translate, to make what's

$50 now equivalent to, in dollars, in 10 years. OK, so in economic applications, we often use the inflation rate.

There is certainly some discounting that we do for future returns in conservation settings. One caution that we certainly have to have is as humans, I think there is evidence in the psychological literature that when we make intuitive decisions, we have a really strong tendency to over discount the future. It's easy for us to discount the future quite heavily and not think about outcomes in the future and focus on more immediate outcomes. And maybe there's reasons for that. I mean maybe there's some evolutionary aspects to this about how we evolved and how we dealt with uncertainty about the timing of when gains might occur.

But we have to be aware of that. We have to be aware of that tendency and think, kind of deliberately, about how we should be discounting the future. Should we be discounting the future for conservation applications?

I mean there is some really interesting and difficult, challenging philosophical issues here. But compare these two things. Suppose you've got one scenario in which a species goes extinct in 50 years versus another scenario where it goes extinct in 500 years.

Do you value those two scenarios differently? In both cases, the species is going to go extinct. But it goes extinct at a different time. Does that matter to you? And how much would you value those differently?

And I think some of that's a real struggle. And I think, particularly when we're thinking about-- I think climate change really gives us some challenging aspects here. Because we're starting to think about changes that are occurring that could be occurring to the ecosystems over 10 years, 50 years, 75 years, 100 years, 200 years, 500 years. And we need to think about how those different outcomes at different time scales are important to us.

So there's some challenging issues here. And we need to think about this temporal discounting. And really what it has to do with is how we're valuing resources and returns over different time periods, the utility that we placed on them. Now a little bit of caution is needed, however, particularly when applying discounting in some conservation settings.

If you discount the future too heavily, for example, in harvest management situations, you would favor harvesting everything right now. So for example, suppose you're harvesting a deer population. And let's say that deer population is expected to grow at, let's say, 10% per year, 10% or 12% per year. That's the intrinsic growth rate of this deer population.

If you discount future returns at a rate greater than 12%, what that would suggest is you're better off harvesting all the deer now. Kill them all now. Take them all now. That's going to give you the highest return when you take that discounting into effect.

And so I think we've got to be careful there because we might have some additional conservation goals about long-term persistence of both the deer population and hunting opportunity that are important. And so we have to be aware of how we're applying discounting in some of those kind of conservation settings.

So in summary with regard to decisions in the face of uncertainty, we need to consider the decision maker's attitude towards risk. This amounts to valuing outcomes in a way that's not linearly related to the gain or the loss. And when we do that, the expected utility-- that's one method we can use. And that can capture the decision maker's risk attitude. So there's tool for that.

The point of this all is to expose you to the idea that there's tools for making decisions in the face of uncertainty, when you have to live with that uncertainty. Now what I've talked about so far, that assumes we have to live with the uncertainty. And what we're going to do in the next module is talk about the case where maybe we can reduce that uncertainty. And if we have the opportunity to reduce that uncertainty, how do we fold that into our decision making? So more on that in the next module.

[MUSIC PLAYING]