

DECISION MAKING UNDER UNCERTAINTY: STATISTICAL DECISION THEORY

Introduction

I'd like to start today's lecture with a reminder about something I said a long time ago when we finished our survey of population viability analysis. Population viability analysis is best seen not as a way of garnering precise predictions about the fate of a population but as a way of ensuring that all relevant life-history variables have been considered, that they have been considered efficiently, and that we have a reasonable sense of the trajectory that the population is likely to follow if current trends continue. It provides a way of structuring our thinking about the problem.

That's precisely the way I think we should regard the approach to decision making that I'm about to describe. One of the most difficult tasks facing conservation biologists, as I have emphasized repeatedly, is that decisions must often, perhaps usually, be made in the face of woefully inadequate data. There is no way to freeze the status quo in a time capsule while we study the situation and come up with a solution. A few years ago one of the small projects in this course involved a critique of the recovery plan for Furbish's lousewort (*Pedicularis furbishiae*). Many of the students assigned that project expressed dissatisfaction at the amount of data on which the advice was based, even though it was based on more *10 years* of field work. Not only was the amount of information about life history and population dynamics much greater than typically available for endangered plants then, it's still vastly more than is available for most endangered plants now. Even though you might like to have more data before writing a recovery plan, you have no choice but to write it with the data that's available at the time.

Similarly, if you're working for the Fish & Wildlife Service, The Nature Conservancy, or a local land trust and you have an opportunity to buy or accept the gift of a piece of land, you won't have a lot of time before you have to make a decision on whether or not to accept it. If you're lucky you may have a chance to walk the property a couple of times and consult some museums, herbaria, or databases about the communities and species present. You won't have the option of waiting until you have all of the information you'd like to have. We'll see this illustrated even more vividly when we discuss attempts to determine global

conservation priorities. Our ability to set priorities is limited by the data available.

Does this mean that we should just give up and go home? No, of course not. It does mean that we should take advantage of any techniques at our disposal that can help us make the best decisions possible, given the large amount of uncertainty with which we'll have to live. That's what the techniques I'm going to describe are intended to do. But they are valuable less because they provide numbers to guide our decisions than because they help us to structure our thinking about decisions in a way that helps to ensure that we don't ignore any important features of the problem.

The Framework of Statistical Decision Theory

How do we reconcile uncertainties with the need for a decision *now*, recalling that a decision not to act until more information is available is still a decision? Risk analysis and statistical decision theory can provide some guidelines.

We've already encountered the most basic idea in decision theory in our discussion of the practical problems facing conservation biologists. Namely, the recognition from statistics that there are two types of errors we can make in evaluating an hypothesis:

- We may say that something is happening when it isn't (Type I), or
- We may say that something isn't happening when it is (Type II).

In classical statistics, an experimenter must decide beforehand the acceptable level of Type I error and design an experiment with an acceptable level of Type II error. The decision on whether to accept or reject the null hypothesis then depends *only* on the level of Type I error the experimenter is willing to accept, e.g., the conventional level of 5%.

An alternative approach, known as Bayesian statistics, treats the probability quite differently.¹ Prior beliefs are combined with the likelihood of the data to produce posterior probabilities about the hypothesis. More importantly, Bayesians treat the decision of whether to accept or reject the null hypothesis as a real *decision*, i.e., as an action that has consequences. Moreover, they recognize that the cost of being wrong depends on how you're wrong. It is probably worse, for example, to decide that there is no evidence of a worldwide decline in amphibian populations when they are declining than to decide that there there is evidence for a worldwide decline when they are not declining. We have to remember that in conservation planning, you don't have the option of *not* deciding. At most you have the option of deciding *not to decide*, which is still a decision. Saying

¹Those of you who know me will be astonished that it's taken me this long to get to a discussion of Bayesian approaches.

“I don’t know whether species X poses a significant threat to conservation values in area Y. Let me study the problem for five years and get back to you.”

is equivalent to saying

“I don’t know whether species X poses a significant threat to conservation values in area y. *But I know that any threat it poses is mild enough that I can afford to investigate the problem for five years before doing anything to reverse the threat.*”

Decision Trees

The most basic technique in decision theory is the decision tree. Take, as an example, the question of whether to regulate air pollution to improve forest health.²

Suppose you are asked to provide an opinion on whether or not air pollution should be regulated to improve forest health.³ Three responses are possible:

1. Pollution *is* affecting forest health and should be regulated.
2. Pollution is *not* affecting forest health and should not be regulated .
3. The evidence is inconclusive (so I need some grant money to study the problem, and I’ll get back to you).

Of course, regardless of your opinion about whether pollution is affecting forest health, pollution either *is* having an effect or it is *not* having an effect. Even if you regard the evidence as inconclusive, pollution either is or is not having an effect.⁴

In addition, regulators will either decide to regulate air pollution or not. We hope that their decision is informed by the advice we give them, but their decision on whether to regulate or not depends on whether they think that pollution is affecting forest health,⁵ and just as our opinion can be wrong, their decision can be wrong too.⁶ Notice that there are 12 possible outcomes: 3 answers we might give to the question, 2 possible states of the world, and 2 possible decisions by the regulator.

These possibilities and their consequences can be summarized in a decision tree (Figure 1).

²There are obviously *other* reasons why we might want to regulate air quality as well, but we’re only going to consider this aspect of the problem.

³This example is taken from [1].

⁴The effect may not be particularly large or important, but it’s unlikely that there is *no* effect. A complete discussion would phrase the question differently and ask whether pollution is having a large enough effect

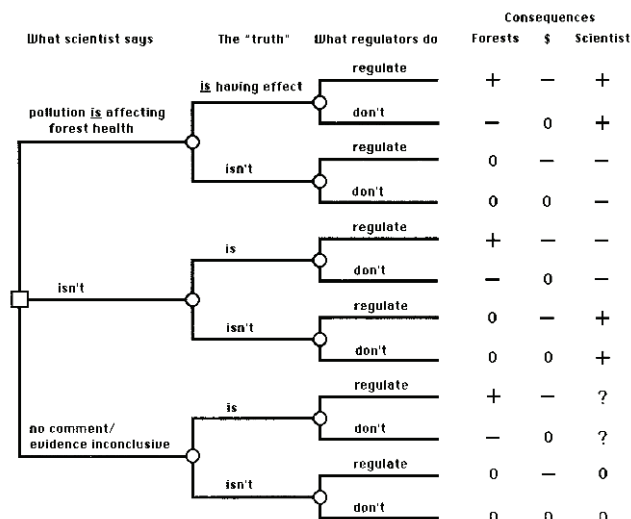


Figure 1: Decision tree for analysis of regulation on air quality intended to improve forest-health.

- The squares represent points in the tree over which you have control as a decision-maker, i.e., they represent the decisions you can make. In this case our decision is what to tell the person who asked for our opinion.
- The circles represent points in the tree over which you have no control.
- Along each branch of the tree, all possibilities are considered.

The consequences of following each branch of the tree can be evaluated for several different variables. Notice that the payoffs to the scientist and the forest are not the same. More importantly, they payoffs are measured in different units. They're not easy to compare.⁷

on forest health to make it worth regulating. To keep things simple, we'll ignore that subtlety.

⁵For simplicity we'll ignore the possibility that their decision is influenced by other factors, like economic impact.

⁶Although we hope that we're less likely to be wrong than they are.

⁷This is a point we'll come back to after Thanksgiving break.

Statistical Decision Theory

For many purposes, simply laying out all the alternatives and their consequences is sufficient to arrive at a reasonable solution. The preponderance of effects in one direction may be so overwhelming that you don't need to do anything more. In other cases, it may be difficult to decide on a course of action based on the tree skeleton and consequences alone.

In such cases it may be useful to do the analysis quantitatively by assigning probabilities to various outcomes along the decision tree. Numerical values may also be assigned to the various outcomes, and a mathematical criterion may be chosen to identify the "best" decision.

One of the most difficult parts of doing a formal, quantitative analysis of the decision tree is arriving at probability estimates for the various outcomes. After all, the whole reason we're even exploring this approach is that we don't have as much information as we'd like.

- This is one place where you can use some interesting techniques to elicit *subjective* probabilities from "gut" reactions. In that way, you take advantage of the intuitive "feel" many will have for the situation and explore their consequences. In the absence of solid data, we can even use guesses. This approach guarantees that we are at least consistent in how we use those subjective impressions.
- We'll also talk about a couple of alternatives that don't require us to estimate these probabilities before introducing a couple of less structured approaches.

Of course, if you have a formal population viability analysis with persistence probabilities specified over some time period, so much the better. This approach allows them to be incorporated directly.

The Procedure

1. Define a *gain function* or a *loss function*. This is the thing that measures the value or cost of a particular outcome. The value function may be qualitative—allows persistence, doesn't allow persistence—or it may be quantitative—the probability of extinction in 100 years is p or the present value of the population expected to exist 100 years from now is x .⁸

⁸For those of you who've had a little economics (and especially for those of you who haven't), let me remind you that if you're using dollars as your measure, it's not enough to calculate how much something is worth 100 years from now. Investment decisions are generally, and reasonably, based on the amount of money you'd have to invest now at some specified interest rate (the *discounting* rate) to have the amount you're interested in at the end of that period. To have \$1,000,000 100 years from now, I would need \$10,051.84 today, assuming a 5% discount rate (the magic of compound interest). Thus, the present value of \$1,000,000 in 2100 is \$10,051.84 today. We'll return to these ideas in the next couple of lectures.




Figure 2: Decision tree for analysis of rhino conservation.

- Associate utility values with each possible outcome in your decision tree. Utilities are a measure of how valuable (or how costly) a particular outcome is. Costs are often measured in dollars (or other currencies), but properly speaking the money is used merely as a metric to make different choices comparable. The dollar value *per se* isn't important. It's what the dollar value represents. To make this concrete, let's take a look at the decision tree Maguire and her co-authors [2] present for management of the Sumatran rhino (Figure 2). In this case, we have two possible *loss* functions: (1) an estimate of the probability of extinction associated with the particular management alternative and (2) an estimate of how much it would cost to implement each of the management alternatives.
- Having associated values with each possible outcome, you associate probabilities with each branch on the tree that's outside your control. These probabilities should reflect the probability of a particular outcome, given the events that have preceded it in the tree. Thus, the probabilities of particular events may depend on their context, as they should if we really think our management decision is going to have an effect.
- Now that you have values associated with each possible outcome and probabilities associated with each branch of the tree, it's possible to calculate the *expected* values associated with each possible decision. The expected value of a decision is the

weighted average of all outcomes associated with that decision, where the weights are the probabilities associated with each step in the decision tree.

So for example, if we were to intervene by establishing a captive breeding program, the diagram tells us that

- There is an 80% chance that the captive breeding program will succeed, in which case there's zero chance of extinction in 30 years.
- There is a 20% chance that the captive breeding program will fail, in which case there's a 95% chance of extinction in 30 years.
- The *expected* probability of extinction is

$$(0.8)(0.0) + (0.2)(0.95) = 0.19$$

This compares with a 0.53 expected probability of extinction if reserves are simply expanded.

If we take the probability of extinction as our loss function, we're done. We simply choose the alternative that has the lowest probability of extinction, namely captive breeding, which has $pE = 0.19$.

Eliciting a Utility Function

Recall, however, that there's another loss function included in this figure, one with a dollar sign at the top. Different management options will require different amounts of money. Because resources, including money, people, time, and energy, are limited, we need to consider the possibility that choosing the option with the lowest probability of extinction may cause us to lose opportunities to do other things. Simply fencing reserves, for example, would provide a 45% expected probability of survival. Is it really *worth* spending six times as much money for a captive breeding program that would increase the expected probability of survival to 81%?

One way of answering that question is simply to think about the decision as if we were gambling.⁹ To make it simple, let's suppose that our two choices are:

1. Spend \$0.6M to fence reserve and have a 45% chance of survival.
2. Spend \$3.7M on a captive breeding program and have an 81% chance of survival.

⁹Which we are, in some sense.

Let me change the wording of those options a little:

1. Buy a lottery ticket for \$6 that gives you a 45% chance of winning a million bucks.
2. Buy a lottery ticket for \$37 that gives you an 81% chance of winning a million bucks.

Which one of those would you choose? Well, the *expected* return from the first option is

$$0.45 * 1000000 - 6 \approx 450000$$

While the expected return from the second option is

$$0.81 * 1000000 - 37 \approx 810000 \quad .$$

Not too hard to pick, is it? The second option is clearly better. What if instead of that lottery ticket winning you a million bucks, though, it won you only fifty. Then the return from the first option is

$$0.45 * 50 - 6 = 16.5$$

and the return from the second option is

$$0.81 * 50 - 37 = 3.5 \quad .$$

Again, not too hard to pick, but now we'd pick the first option instead of the second one.

What does this have to do with rhinos? Just an illustration of the simple point that it's not only how much it *costs* to save rhinos that matters, but also how much it *benefits* us to save them. The challenge that we'll talk more about after Thanksgiving break is that we can measure costs fairly easily in dollars and cents (or Euros or pounds or Rand or ...). It's typically a lot harder to measure the benefits.

One approach is to use what economists call a *utility function*. The utility function is intended to describe how much a unit of something is "worth."¹⁰ I've already illustrated the idea with the two extremes in terms of lottery outcomes, but let's see how we might go about finding one for the rhino example.

- A zero probability of extinction (the desired outcome) is, by definition, associated with a utility value of 1.
- Certain extinction is associated with a utility value of 0.

¹⁰Economists are often accused of trying to put a dollar value on everything, and we'll talk about valuation methods after Thanksgiving break. It's important to realize, however, that dollars are really just a convenient shorthand for economists. What they are usually trying to do is to maximize expected utility.

- All intermediate probabilities of extinction are associated with a utility value between 0 and 1, inclusive, and the utility function should be non-increasing.
- How much is a 50% probability of extinction worth on this scale?
- How much is a 25% probability of extinction worth on this scale?
- How much is a 75% probability of extinction worth on this scale?
- Interpolate a curve. (1999 results $U(p) = 0.99 + 0.64p - 3.77p^2 + 2.13p^3$).
- Calculate utility of

– status quo

$$\begin{aligned} (0.1)(U(0.95)) + (0.9)(U(0.85)) &= (0.1)(0.022) + (0.9)(0.12) \\ &= 0.11 \end{aligned}$$

– new reserve

$$\begin{aligned} (0.6)(U(0.9)) + (0.4)(U(0.4)) &= (0.6)(0.065) + (0.4)(0.78) \\ &= 0.35 \end{aligned}$$

– captive breeding

$$\begin{aligned} (0.8)(U(0)) + (0.2)(U(0.95)) &= (0.8)(1) + (0.2) * (0.022) \\ &= 0.80 \end{aligned}$$

– expand reserve

$$\begin{aligned} (0.1)(U(0.9)) + (0.2)(U(0.9)) + (0.7)(U(0.37)) &= (0.1)(0.065) + (0.2)(0.065) + \\ &\quad (0.7)(0.82) \\ &= 0.59 \end{aligned}$$

Sensitivity Analysis

As you can see, coming up with a utility function isn't easy, especially if there are many people involved. The value I place on a particular outcome is likely to be different from yours. The actions themselves may change my personal utility function.¹¹ Still, coming

¹¹I put a lot more value in Bach and Mozart now than I did in junior high, for example.

up with a utility function may be a useful exercise to get us to think carefully about the alternatives and their consequences.

There is still the danger that we will be captured by the apparent precision of the numerical results. One way of combating this, and sometimes of producing agreement where there was none, is through a sensitivity analysis. The idea is exactly the same as we saw with sensitivity and elasticity coefficients in demographic matrices. Which aspects of the decision tree and the utility function have a big impact on the decision.

- Use the existing utility function

- change probability that captive breeding is successful to 0.5

$$(0.5)(U(0)) + (0.5)(U(0.95)) = 0.511$$

- change probability that new reserve prevents timber harvest to 0.8

$$(0.05)(U(0.9)) + (0.15)(U(0.9)) + (0.8)(U(0.37)) = 0.67$$

- **Exercise:** Use Maguire's probabilities change the utility function and calculate utility of captive breeding versus new reserve.

Sensitivity analysis allows us to explore how different things would have to be for us to change our minds about the decision. Indicates how robust our decision is and points out areas where uncertainty is particularly problematic.

An example to try on your own

Question: Should logging be halted immediately in all old-growth forests of the Pacific Northwest?

- Define the possible outcomes
- Define the some gain/loss functions, e.g., pE for northern spotted owl, economic impact on local communities
- Define the intermediate states that affect the probabilities of the various outcomes
- Define the branch-specific probabilities of particular event sequences
- Calculate the expected gains and losses
- Translate the expected gains and losses into utilities.
- Explore sensitivity of outcome to various parameters.

	No damage	Low damage	High damage
Stringent control	9	10	11
Moderate control	1	5	14
Lax control	0	10	30

Table 1: Total costs associated with different combinations of regulation and forest damage (modified [3]).

Alternatives to expected utility

Suppose we don't want to (or can't) come up with reasonable probabilities for our decision tree. Then we can't use all that nice mathematical machinery I just described. Are we dead in the water?¹² No. There are a couple of alternatives we can try.

Let's return to our example of whether to regulate air pollution.¹³ Suppose that there could be either no damage, low damage, or high damage to forest trees and there could be either lax pollution control, moderate control, or stringent control. That makes nine possible combinations of levels of regulation and damage. With each of those combinations there's a cost: the cost of regulation and the costs associated with damage to the forest. Those costs are summarized in Table 1.

If we knew the probability of each outcome given a choice of control measures, we could calculate the expected cost associated with each choice (low, moderate, or stringent control) and choose the one with the lowest expected cost. But that's not the only choice we could make.

The mini-max cost choice

We could decide that we want to pick the alternative that makes our maximum cost as small as possible, i.e., the one that *minimizes* our *maximum* cost. Looking back at Table 1, it's easy to see that (a) the *maximum* loss associated with each choice occurs when there is high damage and (b) the *minimum* loss in that column is associated with stringent pollution controls. Therefore, the minimax choice would be to enact stringent pollution controls.

As Polasky et al. [3] point out, the mini-max choice is *very* conservative.¹⁴ It puts all of

¹²Would I be asking this question if the answer were "Yes"?

¹³This example is adapted from [3].

¹⁴Because I've changed their example a bit, what they call a maxi-min choice is the same as the mini-max choice as I've described it here. Sorry for the confusion, but I find it easier to make costs positive than to leave them negative. That's the cause of the difference.

	No damage	Low damage	High damage
Stringent control	$9 - 0 = 9$	$10 - 5 = 5$	$11 - 11 = 0$
Moderate control	$1 - 0 = 1$	$5 - 5 = 0$	$14 - 11 = 3$
Lax control	$0 - 0 = 0$	$10 - 5 = 5$	$30 - 11 = 19$

Table 2: “Regrets” associated with different combinations of regulation and forest damage (based on Table 1).

the weight for the decision on the worst possible outcome.

The mini-max regret choice

An alternative is to calculate the “regret” associated with making the “wrong” choice under each possible outcome. For example, if we knew that the actual outcome would be low damage, the best choice would be to enact moderate pollution controls, because it’s the least expensive. We’d lose 5 units if we chose either lax or stringent regulation. The “regret” matrix is shown in Table 2. Each entry is calculated by identifying the smallest number in a column and subtracting it from all of the numbers in that column.

Now we might pick the alternative that makes our maximum “regret” as small as possible, i.e., the one that *minimizes* our *maximum* “regret.” Looking at Table 2, it’s easy to see that (a) the maximum regret for the stringent control choice is 9, for the moderate control choice is 3, and for the lax control choice is 19 and that (b) the choice that minimizes the maximum “regret” is moderate control.

It’s vital to note that neither mini-max on cost nor mini-max on “regret” can be said to be the “right” choice. They balance tradeoffs in different ways related to how risk averse you might be. It’s also important to notice that even if you *can* estimate probabilities and calculate the expected cost associated with each action, you might still prefer a mini-max choice based on either cost or “regret”.

Less structured approaches

While I think there’s a lot to be gained from the systematic attention to detail that structured approaches provide, Polasky et al. [3] describe less structured alternatives that may also be useful in some circumstances.

- **Thresholds approach:** Rather than trying to choose explicitly among alternatives it may be sufficient to identify boundaries that define acceptable from unacceptable or

safe from dangerous outcomes. If the we are approaching a threshold that would lead to drastic changes, e.g., loss of the stratospheric ozone layer, it may lead to a policy response. One danger of a focus on thresholds is that it may lead to the impression that values below the threshold are harmless or without risk.

- **Scenario planning:** There may be a host of variables that are difficult to quantify or unravel. Nonetheless, it may be possible to develop plausible scenarios that describe different possible outcomes, e.g., the different emission scenarios envisioned by the IPCC. If the consequences of different scenarios are clear, that may be sufficient to facilitate decision making. One danger is that those who identify the scenarios may fail to identify important possibilities. Scenario planning requires a great deal of creativity, and it's difficult to be sure that all relevant scenarios have been considered.
- **Resilience thinking:** In my mind, resilience thinking could be considered a subset of the thresholds approach: first, identify the critical thresholds a system might cross; second, develop early warning systems to identify when the system is approaching those thresholds; and third, develop plans that will either pull the system back from those thresholds or allow it to function appropriately under the new regime.

Conclusions

- Decision theory provides a useful framework to explore alternatives.
 - It forces us to recognize that deciding not to take action is just as much a decision as deciding which action to take.
 - It forces us to recognize that we may err either by taking an unnecessary action or by failing to take a necessary action.
- It helps us formalize and categorize our thinking to make sure that we have considered all relevant possibilities.
- Quantitative analyses must be viewed as explorations of possibilities, not hard predictions, but
 - the process of quantification may help us to clarify our thinking, and
 - it provides us a way of assessing which parts of the decision tree have a particularly large impact on the outcome and of determining how robust our preferred course of action is to other possibilities. In complex problems this may not be evident from the outset.

- It makes clear that many of the political conflicts involving environmental and conservation decision-making arise because different participants in the process place different values on different aspects of the outcome. That's why we're going to try to talk a bit about values in the last few lectures of this course.

References

- [1] L A Maguire. Risk analysis for conservation biologists. *Conservation Biology*, 5:123–125, 1991.
- [2] Lynn A Maguire, Ulysses S Seal, and Peter F Brussard. Manging critically endangered species: the Sumatran rhino as a case study. In Michael E Soul, editor, *Viable Populations for Conservation*, pages 141–158. Cambridge Univ. Press, Cambridge, 1987.
- [3] Stephen Polasky, Stephen R Carpenter, Carl Folke, and Bonnie Keeler. Decision-making under great uncertainty: environmental management in an era of global change. *Trends in Ecology & Evolution*, In Press,, 2011.

Creative Commons License

These notes are licensed under the Creative Commons Attribution-NonCommercial-ShareAlike License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-sa/3.0/> or send a letter to Creative Commons, 559 Nathan Abbott Way, Stanford, California 94305, USA.