

Decision Theory
 © 2006 Samuel L. Baker
 Assignment 13 is on page 11.

Decision theory is about how to decide among alternatives when you face uncertainty about what will happen after you choose.

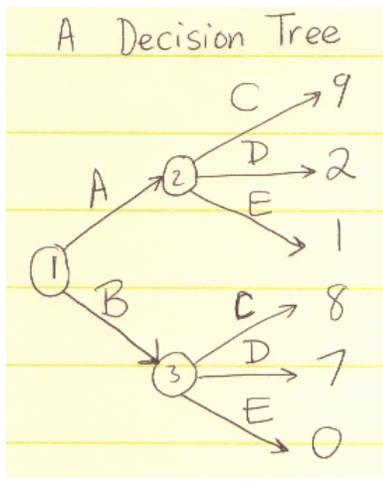
To apply mathematics to making decisions, you have to be able to put numerical values on your possible outcomes. Better outcomes have higher numbers; worse outcomes have lower numbers. This is most straightforward to do for business decisions for which the outcome is the amount of money that you will gain or lose.

If each choice leads directly to an outcome, the decision is easy: You make the choice that gives you the outcome with the highest number.

Uncertainty makes decisions interesting. Consider the following table, called a payoff table, adapted from a classic textbook¹:

		What the world does		
		C	D	E
What you do	A	9	2	1
	B	8	7	0

You have a choice of two strategies, A and B. After you choose A or B, the world chooses C, D, or E. The cells with numbers tell you how much you gain for each combination of your strategy and the world's strategy. For example, if you do A and the world does C, you get 9.



You do not know for sure what will happen after you make your choice. Choosing A might win you as much as 9, or you may have to settle for just 2 or 1, depending on what happens after you make your choice. Choosing B might win you as much as 8, or you may have to settle for 7 or 0, again depending on what happens after you make your choice.

Another way to look at this situation is using a decision tree. In the tree diagram, each decision is at a numbered branching point. At (1), you pick A or B. That moves you to either (2) or (3). The world then picks C, D, or E. The result is the payoff to the right.

Decision tree diagrams can have more branchings. For example, suppose you pick A and the world picks C. You could then get another

¹ William J. Baumol, *Economic Theory and Operations Analysis, Second Edition*, Prentice-Hall, 1965. “Adapted” means that I got the idea from Baumol, but changed the numbers.

decision, followed by the world getting another decision. The C arrow would point to another numbered decision point. More arrows would come off of each of that point. The world would have another set of choices, back and forth until finally you get to the payoffs.

Decision Theory tries to give you guidance about what to do in this kind of situation. How you decide what to do depends on:

1. What your needs are, such as how well you can afford a bad outcome and how much a good outcome would mean to you.
2. If you have any idea how to anticipate what the world will do.

Before we go further, I do want to point out that in real life it could take a lot of work just to get to where you could write out a payoff table like the one above. You would have to analyze the costs and benefits of each combination of your actions and what the world does. As with the critical path analysis that we did, a lot of thought and work goes into just getting the problem set up so that you can apply the mathematical technique. A lot of the benefit of using the technique is that it requires you to think about your project or your choices in an organized way.

Let us look at some possible ways to decide.

Expected value as a criterion

If you can assign probabilities to the world's choices, you can convert this uncertainty problem into a risk problem. ("Risk," in economics jargon, applies situations in which you know the probabilities of the various outcomes.) Then you can choose the strategy that gives you the highest expected value.

To get the expected value, multiply the value of each outcome by the probability of that outcome. Then add up the products.

For more explanation of risk and expected value, please see this on-line tutorial:
<http://hspm.sph.sc.edu/COURSES/ECON/Risk/Risk.html>

To apply the expected value method, we need to assign probabilities to the world's strategies, C, D, and E? What numbers should we use? I have no idea! We will have to make something up.

In practice, if you know something about the situation, you may be able to come up with some reasonable probabilities for C, D, and E. Any reasonable guess is better than nothing.

There are some who argue that, if you have no idea, you should give an equal probability to each of the world's strategies. Here, that would mean giving each one a probability of one-third. That seems bogus to me. For one thing, it depends on what strategies for the world you choose to put in your model. What if strategy E is really two possibilities, E_1 and E_2 , that are fairly similar? Should you then give C and D probabilities of one-fourth, and the two E's a probability of two-fourths ($1/2$)? It is common to imagine the best that can happen, the worst that can happen, and something in-between. C, D, and E might be the best, in-between, and worst. That is reasonable as a guide to planning, but it doesn't make the three possibilities equally likely.

As we will see later, it is possible to analyze this situation without assigning probabilities, but that is not ideal either. After reading those sections, you may decide that a poor effort to assign a probabilities is better than no effort. If you can research the history and come up with reasonable probabilities, you should generally try them and see what decision they imply. Let us do that here.

Let us imagine that you have researched history and consulted experts. You conclude that the world's strategies C, D, and E, are pretty much equally likely. Each has a probability of $1/3$ of occurring. (Hmm... Same as the bogus method, but let us do it anyway.)

We then have this calculation of the expected value of each strategy:

$$\begin{array}{l} \text{Strategy A expected value:} \quad 4 \quad = 1/3 \times 9 \quad + 1/3 \times 2 \quad + 1/3 \times 1 \\ \text{Strategy B expected value:} \quad 5 \quad = 1/3 \times 8 \quad + 1/3 \times 7 \quad + 1/3 \times 0 \end{array}$$

Strategy B is the winner. Its expected value is 5, compared with A's expected value of 4.

A sensitivity analysis can be done. This means finding out how much we have to change the probabilities to make the decision switch from A to B. If B is better over a big range of possible probabilities, that might make us feel better about choosing B. We could then say that our choice to do B is not "sensitive" to our choice of what probabilities to assign to C, D, and E.

Let's do some algebra:

The expected value of doing A is $9c + 2d + 1e$, where c is the probability that C will happen, d is the probability that D will happen, and e is the probability that E will happen. 9, 2, and 1 are the payoffs if C, D, or E happen.

The expected value of doing B is $8c + 7d + 0e$. 8, 7, and 0 are the payoffs in C, D, or E happen.

A has a higher expected value if $9c + 2d + 1e > 8c + 7d$. (The $0e$ drops out, because it's 0.)

The possibilities for the world are C, D, or E. Their probabilities have to add up to 1, so $c + d + e = 1$.

We can solve that last equation for e , to get $e = 1 - c - d$.

Substituting that for e in our inequality gives us $9c + 2d + 1 - c - d > 8c + 7d$.

Combining like terms gives us $8c + d + 1 > 8c + 7d$. This reduces to $1 > 6d$, or $d < 1/6$.

We find that A is the better strategy if and only if the probability of D is less than 0.1666... (Usually the numbers won't work out to give you an answer that is this simple.)

This is our sensitivity analysis: If you are comfortable with giving D a probability of at least one out of six, your better strategy is B. You can say, equivalently, that B is better so long as the combined probability of C and E is less than $5/6$. (This is because the probabilities of all three add up to 1.)

We assumed when we started that the probability of outcome D is $1/3$. That is reasonably far from $1/6$,

which is our tipping point between choosing B and choosing A. If our tipping point between choosing B and choosing A were closer to 1/3, we would have said that our choice of B instead of A was “sensitive” to the probabilities we assigned. As it is, we can say that our decision to go with B is not sensitive to small changes in the probabilities we assigned.

The expected value method is sometimes called the Bayes method, after the Reverend Thomas Bayes (1702-1761), who advocated using intuitively determined probabilities explicitly in statistical analysis. Bayesian statistics is named for him.

Expected value with risk aversion

Often, in business decisions, the payoffs are in money terms. We can imagine that, in our example, the “9” for what happens if we do A and the world does C represents that we would have a net gain of \$9 (or \$9 million). Money and value may not be the same thing, though, even when all the benefits and costs are money.

My economics tutorial at <http://hspm.sph.sc.edu/COURSES/ECON/RiskA/RiskA.html> discusses risk aversion. This is the idea that most people do not like risks, especially risks of big losses. They are willing to pay money to an insurance company to have the company take on the risk for them. They will do this even if the insurance premium is greater than the expected value of their loss.

The tutorial presents the Von Neumann-Morgenstern theory to explain this aversion to risk. That theory says that the value of extra money diminishes as you get more and more of it. It is like saying that \$9 is not really worth – to you – 9 times as much 1 dollar. Nine dollars may only be worth three times as much as \$1.

If 9 dollars is really only worth three times as much as \$1, we can extrapolate from that and say that the actual value to you of any amount of money is the square root of the number of dollars. That is because the square root of 9 is 3. By taking the square root of 9, we get the value to you – 3 – of \$9.

If we assume that the square root is how we convert money amounts into actual values to us, we take the square root of all the numbers in the payoff table and get:

Payoff Table with Square Roots of the Money Values

		What the world does		
		C	D	E
What you do	A	3	1.414	1
	B	2.828	2.646	0

This changes the relative values of the outcomes. Adjusted this way, the high numbers are not so high anymore.

If we assume that the probabilities of C, D, and E are all equal, at 1/3, we find this:

$$\text{The expected value of A is } 3.000 \cdot 0.3333 + 1.414 \cdot 0.3333 + 1 \cdot 0.3333 = 1.805$$

The expected value of B is $2.828 \cdot 0.3333 + 2.646 \cdot 0.3333 + 0 \cdot 0.3333 = 1.825$

It's close, but B is still your preferred strategy.

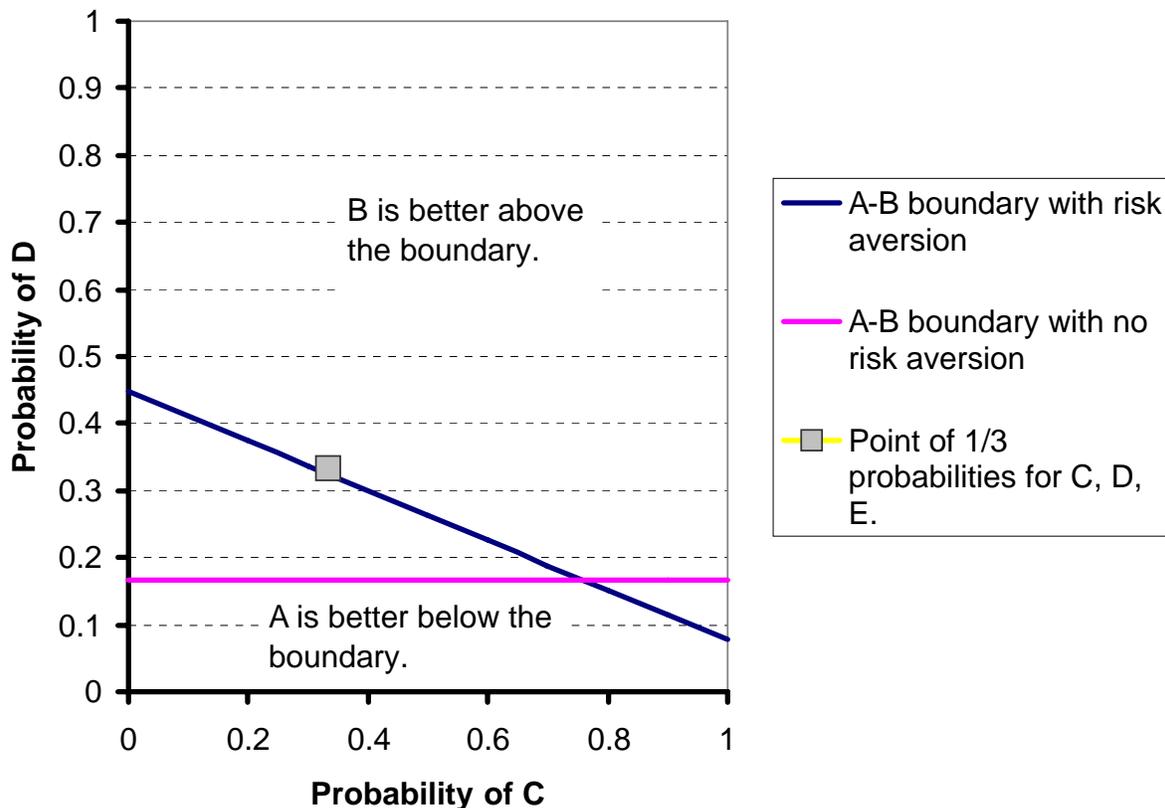
The fact that it is so close should make you suspect that this time the decision is very sensitive to our choice of probabilities.

If we do the algebra as before we get this:

A is better if $3.000c + 1.414d + 1e > 2.828c + 2.646d$. The numbers c, d, and e are the probabilities of C, D, and E, respectively.

As before, $c + d + e = 1$. Substituting $1 - c - d$ for e and simplifying, we wind up with $1 > 0.828c + 2.232d$. If this inequality is true, A is better. Otherwise, B is better.

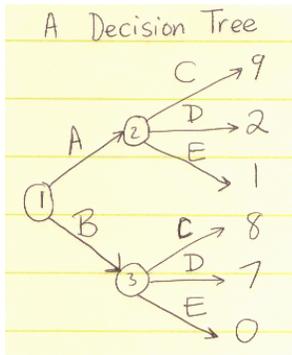
The graph below shows this inequality as a line, the boundary between the A-is-better region and the B-is-better region. For comparison, the graph also shows the boundary that we had before we brought in the risk aversion idea. Then, our equal-probability point was pretty far from the boundary. Now, the equal-probability point is very near the boundary. If we reduce the probability of C or D by just a little, which also means that if we increase the probability of E by a little, we move into the region where A is better. Our decision is now very sensitive to our assumptions about probabilities.



The Maximin criterion

If you do not wish to assign probabilities to what the world does, you can try applying ideas from Game Theory. One such idea is to treat the world as if it is a crafty opponent that will use the strategy that is least rewarding to you. To protect yourself, you must choose the strategy whose worst outcome is least bad. This is the Maximin criterion. You find the strategy that has the maximal (highest) minimum (least good) outcome. The best of the worst.

		What the world does		
		C	D	E
What you do	A	9	2	1
	B	8	7	0



The worst that can happen if you choose A is that the world will choose E and you will get 1. The worst that can happen if you choose B is that the world will choose E and you will get 0. A is better – it has the maximin payoff.

Unless the world really is out to get you, the maximin criterion is a highly pessimistic way to choose. Cowardly, even. It is equivalent to assuming that E will happen with a probability of 1.

The maximin criterion ignores most of the information in the table of payoffs. B is much better than A if the world happens to choose D. The payoffs in the bottom row could be 80, 70, 0 or even 800, 700, and 0. That would be a huge advantage for B if the world chooses C or D. The maximin criterion would have you choose A anyway. It says you must ignore those big differences in outcome and just look at how A gets you 1 more than B if the world chooses E. There is no up-side gain to the maximin strategy.

Another way to think of the maximin is that it is risk aversion carried to an extreme. If the 9, 8, 7, and 2 were really worth to you no more than the 1 in the upper right corner, you would definitely choose strategy A.

The Maximax criterion

Nobody uses this, but it is worth thinking about because it is the opposite of the maximin criterion. Maximax means best of the best. It says to choose the strategy that gets you a chance to get the best possible outcome, no matter how bad the other outcomes are. The world will be your friend, and will choose the strategy that gets you that best outcome.

The maximax strategy in our example is A. That gives you the possibility of making 9 if the world chooses C.

I should not say that this criterion is never used. Players of contract bridge do use this sometimes. The declarer may deduce that the only way to make the contract is if the cards lie a certain way. He or she will then play the hand as if the cards lie that way. The scoring system in duplicate bridge tournaments

encourages this kind of gamble.

If you are not in a bridge tournament, the maximax has the same disadvantage as the maximin, except in reverse. The maximax ignores all the possible bad outcomes. It tells you to go for the best, no matter how little the advantage over the second best and no matter how bad the worse possibilities are.

You could also say that if you use maximax, you are the opposite of risk-averse. You are risk-preferring.

The Minimax Regret criterion

There have been a number of attempts to find a way between the maximin and maximax methods. Here is one, proposed by the economist Leonard J. Savage (1917-1971). His idea was to adapt the notion of opportunity cost and look at “regret,” meaning how much you lose if you make a choice and it turns out that another choice would have been better.

The minimax regret method might also be called the avoid-looking-like-an-idiot-later method.

Here is our original payoff table again:

Payoff Table

		What the world does		
		C	D	E
What you do	A	9	2	1
	B	8	7	0

We use this to create a regret table. The reasoning goes like this:

Suppose you pick strategy A and the world picks strategy C. You get 9. You have no regrets. You do not wish you had chosen B instead.

Similarly, suppose you pick strategy B and the world picks strategy D. You get 7. Here, too, you have no regrets. You do not wish you had chosen A instead.

And again, suppose you pick strategy A and the world picks strategy E. You get 1. You have no regrets here either. You do not wish you had chosen B instead.

The first step in making the regret table is thus: Look for the highest number in each column of the payoff table. Put a 0 in the corresponding place in the regret table.

Regret Table – 1st step

		What the world does		
		C	D	E
What you do	A	0		0
	B		0	

Now for the other cells in the regret table.

The payoff table says that if you choose B and the world chooses C, you get 8. You have some regret. You would have been better off by 1 (9 instead of 8) if you had chosen A. Your regret is 1. That goes in the lower left cell of the regret table.

The payoff table says that if you choose A and the world chooses D, you get 2. You have a lot of regret. You would have been better off by 5 (7 instead of 2) if you had chosen B instead. Your regret is 5. That goes in the top center cell of the regret table.

The payoff table says that if you choose B and the world chooses E, you get 0. You have some regret. You would have been better off by 1 (1 instead of 0) if you had chosen A. Your regret is 1. That goes in the lower right cell of the regret table.

The regret is the outcome you actually got, subtracted from the outcome you would have gotten if you had known that the world was going to do what it did.

Regret Table

		What the world does		
		C	D	E
What you do	A	0	5	0
	B	1	0	1

Now we apply what amounts to the maximin criterion, except that, in a regret table, all the numbers are bad and big numbers are worse. So we find the minimax, the strategy that has the smallest maximum regret number. Strategy B is that strategy. The worst regret we can have with strategy B is 1. If we choose strategy A, we might suffer a regret of 5. Strategy B is the best at protecting us against looking stupid later.

The minimax regret criterion gives reasonable-looking advice in this example, because the regret in column D is so much more than the regrets in the other columns. Even so, this method has a drawback similar to the other methods (maximin and maximax) in that it ignores numbers other than the very worst (or best) one.

Suppose, for example, this were the payoff table:

Different Payoff Table

		What the world does		
		C	D	E
What you do	A	9	2	4
	B	4	7	0

The corresponding regret table would have been:

A Different Regret Table

		What the world does		
		C	D	E
What you do	A	0	5	0
	B	4	0	4

The minimax regret criterion would still pick strategy B, because strategy B can cost you 4, but strategy A can cost you 5. However, with this regret table, strategy B does not look so good intuitively. To pick B, you have to assume that there is a 100% chance that the world will pick strategy D. What if you pick B and the world picks C or E? You will look almost as stupid as if you had chosen A and the world had chosen D. Picking B instead of A gains you only a little looking-stupid protection, while costing you what could be a better chance of looking good, with 0 regret. If you think that the three strategies for the world are close to equally likely, picking strategy B gives you a two-thirds chance of looking stupid after the fact.

What is shrewd about this criterion is that it recognizes that, as an executive, you are an agent for somebody else, the stockholders or the institution that you are working for. Your performance may be judged by Monday-morning quarterbacks who will ask why you weren't smart enough to foresee what was going to happen. Bear in mind, though, that you are not being a good agent if you use the minimax regret strategy. You are putting your own interest ahead of the institution's interest.

Conclusion

There is no one right way to make decisions in the face of uncertainty. In practice, you have to make judgements about what you really value and how well you tolerate risk. For instance, if you are working for a large pharmaceuticals company, they tolerate risk very well, and they understand that many projects do not pan out. The expected value method might be appropriate.

If you are gambling with your own money, or working at a small institution and “betting the farm” on this decision, the expected-value-with-risk aversion or even the maximin criterion may be appropriate. If all you care about is seizing a chance to strike it rich, you might want the maximax criterion. If you are working for Monday-morning quarterbacks who will whine that you should have anticipated what happened, even though there is no way you could have, then you might choose the minimax regret method.

Your turn:

Assignment 13

Consider this payoff table, also from Baumol's book. (He credits it to John Milner.)

Payoff Table

	What the world does				
		<i>E</i>	<i>F</i>	<i>G</i>	<i>H</i>
What you do	<i>A</i>	2	2	0	1
	<i>B</i>	1	1	1	1
	<i>C</i>	0	4	0	0
	<i>D</i>	1	3	0	0

1. Which strategy do you choose if you use the expected value criterion and you assume that *E*, *F*, *G*, and *H* are equally likely, each with a probability of 0.25?
2. Which strategy do you chose if you use the maximin criterion, so that your worse outcome is as high as it can be?
3. Which strategy do you chose if you use the maximax criterion, giving yourself a chance to get the best possible outcome?
4. Which strategy do you chose if you use the minimax regret criterion, protecting yourself as best you can against looking stupid later?

(A note about calculating the regret when you have more than two strategies to choose from: The regret is the difference between what you got and the best you could have done, given what the world did. Your regret is measured relative to the highest number in each column. For example, if you choose strategy *A* and the world chooses *F*, your regret is 2, because you could have gotten 4 instead of 2 by choosing *C* instead of *A*.)

5. Suppose you were making this decision where you are working now. (Alternatively, imagine that you are working at the job that you aspire to.) Which of these criteria would you use? Explain why.