

Decision Theory

There are four types of criteria that we will look at.

Expected Value (Realist)

Compute the expected value under each action and then pick the action with the largest expected value. This is the only method of the four that incorporates the probabilities of the states of nature. The expected value criterion is also called the Bayesian principle.

Maximax (Optimist)

The maximax looks at the best that could happen under each action and then chooses the action with the largest value. They assume that they will get the most possible and then they take the action with the best best case scenario. The **maximum** of the **maximums** or the “best of the best”. This is the lotto player; they see large payoffs and ignore the probabilities.

Maximin (Pessimist)

The maximin person looks at the worst that could happen under each action and then choose the action with the largest payoff. They assume that the worst that can happen will, and then they take the action with the best worst case scenario. The **maximum** of the **minimums** or the “best of the worst”. This is the person who puts their money into a savings account because they could lose money at the stock market.

Minimax (Opportunist)

Minimax decision making is based on opportunistic loss. They are the kind that look back after the state of nature has occurred and say “Now that I know what happened, if I had only picked this other action instead of the one I actually did, I could have done better”. So, to make their decision (before the event occurs), they create an opportunistic loss (or regret) table. Then they take the **minimum** of the **maximum**. That sounds backwards, but remember, this is a *loss* table. This similar to the maximin principle in theory; they want the best of the worst losses.

Example: A bicycle shop

Zed and Adrian and run a small bicycle shop called “Z to A Bicycles”. They must order bicycles for the coming season. Orders for the bicycles must be placed in quantities of twenty (20). The cost per bicycle is \$70 if they order 20, \$67 if they order 40, \$65 if they order 60, and \$64 if they order 80. The bicycles will be sold for \$100 each. Any bicycles left over at the end of the season can be sold (for certain) at \$45 each. If Zed and Adrian run out of bicycles during the season, then they will suffer a loss of “goodwill” among their customers. They estimate this goodwill loss to be \$5 per customer who was unable to buy a bicycle. Zed and Adrian estimate that the demand for bicycles this season will be 10, 30, 50, or 70 bicycles with probabilities of 0.2, 0.4, 0.3, and 0.1 respectively.

Actions

There are four actions available to Zed and Adrian. They have to decide which of the actions is the best one under each criteria.

1. Buy 20 bicycles
2. Buy 40 bicycles
3. Buy 60 bicycles
4. Buy 80 bicycles

Zed and Adrian have control over which action they choose. That is the whole point of decision theory - deciding which action to take.

States of Nature

There are four possible states of nature. A state of nature is an outcome.

1. The demand is 10 bicycles
2. The demand is 30 bicycles
3. The demand is 50 bicycles
4. The demand is 70 bicycles

Zed and Adrian have no control over which state of nature will occur. They can only plan and make the best decision based on the appropriate decision criteria.

Payoff Table

After deciding on each action and state of nature, create a payoff table. The numbers in parentheses for each state of nature represent the probability of that state occurring.

State of Nature	Action			
	Buy 20	Buy 40	Buy 60	Buy 80
Demand 10 (0.2)	50	-330	-650	-970
Demand 30 (0.4)	550	770	450	130
Demand 50 (0.3)	450	1270	1550	1230
Demand 70 (0.1)	350	1170	2050	2330

Ok, the question on your mind is probably “How the [expletive deleted] did you come up with those numbers?”. Let’s take a look at a couple of examples.

Demand is 50, buy 60:

They bought 60 at \$65 each for \$3900. That is -\$3900 since that is money they spent. Now, they sell 50 bicycles at \$100 each for \$5000. They had 10 bicycles left over at the end of the season, and they sold those at \$45 each of \$450. That makes $\$5000 + 450 - 3900 = \1550 .

Demand is 70, buy 40:

They bought 40 at \$67 each for \$2680. That is a negative \$2680 since that is money they spent. Now, they sell 40 bicycles (that’s all they had) at \$100 each for \$4000. The other 30 customers that wanted a bicycle, but couldn’t get one, left mad and Zed and Adrian lost \$5 in goodwill for each of them. That’s 30 customers at -\$5 each or -\$150. That makes $\$4000 - 2680 - 150 = \1170 .

Opportunistic Loss Table

The opportunistic loss (regret) table is calculated from the payoff table. It is only needed for the minimax criteria, but let's go ahead and calculate it now while we're thinking about it.

The maximum payoffs under each state of nature are shown in **bold** in the payoff table above. For example, the best that Zed and Adrian could do if the demand was 30 bicycles is to make \$770. Each element in the opportunistic loss table is found taking each state of nature, one at a time, and subtracting each payoff from the largest payoff for that state of nature. In the way we have the table written above, we would subtract each number in the row from the largest number in the row.

State of Nature	Action			
	Buy 20	Buy 40	Buy 60	Buy 80
Demand 10	0	380	700	1020
Demand 30	220	0	320	640
Demand 50	1100	280	0	320
Demand 70	1980	1160	280	0

Remember that the numbers in this table are **losses** and so the smaller the number, the better.

Expected Value Criterion

Compute the expected value for each action.

For each action, do the following: Multiply the payoff by the probability of that payoff occurring. Then add those values together. Matrix multiplication works really well for this as it multiplies pairs of numbers together and adds them. If you place the probabilities into a 1x4 matrix and use the 4x4 matrix shown above, then you can multiply the matrices to get a 1x4 matrix with the expected value for each action.

Here is an example of the "Buy 60" action if you wish to do it by hand.

$$0.2(-650) + 0.4(450) + 0.3(1550) + 0.1(2050) = 720$$

The expected values for buying 20, 40, 60, and 80 bicycles are \$400, 740, 720, and 460 respectively. Since the best that you could expect to do is \$740, you would buy 40 bicycles.

Maximax Criterion

The maximax criterion is much easier to do than the expected value. You simply look at the best you could do under each action (the largest number in each column). You then take the best (largest) of these.

The largest payoff if you buy 20, 40, 60, and 80 bicycles are \$550, 1270, 2050, and 2330 respectively. Since the largest of those is \$2330, you would buy 80 bicycles.

Maximin Criterion

The maximin criterion is as easy to do as the maximax. Except instead of taking the largest number under each action, you take the smallest payoff under each action (smallest number in each column). You then take the best (largest of these).

The smallest payoff if you buy 20, 40, 60, and 80 bicycles are \$50, -330, -650, and -970 respectively. Since the largest of those is \$50, you would buy 20 bicycles.

Minimax Criterion

Be sure to use the opportunistic loss (regret) table for the minimax criterion. You take the largest loss under each action (largest number in each column). You then take the smallest of these (it is loss, afterall).

The largest losses if you buy 20, 40, 60, and 80 bicycles are \$1980, 1160, 700, and 1020 respectively. Since the smallest of those is \$700, you would buy 60 bicycles.

Putting it all together.

Here is a table that summarizes each criteria and the best decision.

	Action				
Criterion	Buy 20	Buy 40	Buy 60	Buy 80	Best Action
Expected Value	400	740	720	460	Buy 40
Maximax	550	1270	2050	2330	Buy 80
Maximin	50	-330	-650	-970	Buy 20
Minimax	1980	1160	700	1020	Buy 60

Practice Problem

Since there aren't any problems of this kind in the text, work this problem out, and then you can check your answer with the instructor.

Finicky's Jewelers sells watches for \$50 each. During the next month, they estimate that they will sell 15, 25, 35, or 45 watches with respective probabilities of 0.35, 0.25, 0.20, and ... (figure it out). They can only buy watches in lots of ten from their dealer. 10, 20, 30, 40, and 50 watches cost \$40, 39, 37, 36, and 34 per watch respectively. Every month, Finicky's has a clearance sale and will get rid of any unsold watches for \$24 (watches are only in style for a month and so they have to buy the latest model each month). Any customer that comes in during the month to buy a watch, but is unable to, costs Finicky's \$6 in lost goodwill.

Find the best action under each of the four decision criteria.