Operations Research - Lecture 11

Olariu E. Florentin

December 15, 2025

Introduction

- Decision analysis provides a framework and methodology for rational decision making when the outcomes are uncertain.
- *Decision-maker*: a person or a group of persons.
- *Data for decision analysis*:
  - Available courses of action (*alternatives*).
  - *Payoffs* for all possible outcomes.
  - Subjective *probabilities* for all possible random events.
- *Criteria* used for making decisions reflect the attitude towards risk, ranging from optimism to pessimism.
- Decision analysis is a powerful tool for determining an optimal course of action.

- A *question* addressed with decision analysis is whether to make the needed decision *immediately* or to first do some testing (at some expense) to reduce the level of uncertainty about the outcome of the decision.
- *Option* that can be incorporated into the analysis: *experimentation* to obtain better estimates of the probabilities of the possible states of nature.
- Decision analysis divides decision making between the cases of *with experimentation* and *without experimentation*.
- Decision trees are an useful visual tool for analyzing the decision process when a series of decisions needs to be made.
- *Utility theory* provides a way of incorporating the decision maker's attitude toward risk into the analysis.

Practical application of Decision Analysis:

- General Motors: more than 40 major decision analysis projects;
- Chevron uses decision analysis in all major decisions of the company;
- National Weather Service developed a plan for responding to flood (forecasts and warnings).

Kinds of decision making in the face of great uncertainty:

- A manufacturer introducing a new product into a market place.
- A financial firm investing in securities.
- A pharmaceutical company evaluating the potential of new drug candidates.
- A government contractor bidding on a new contract.
- An agricultural firm selecting the mix of crops and livestock for the upcoming season.
- An oil company deciding whether to drill for oil in a particular location.

Decision analysis

- has been used to solve the problem of closing military basis;
- has been used to determine optimal strategies for purchase of long-term care insurance as a function of age, wealth, and risk tolerance;
- is the core of cost-effectiveness analyses in public health and medicine;
- has been used to evaluate alternatives for radioactive-waste repositories (in US and UK).

Ronald A. Howard, professor of Management Science and Engineering at Stanford University is credited with originating the term in 1964.

The Goferbroke Company (GC) owns a tract of land that may contain oil. A consulting geologist has reported to management that she believes there is 1 chance in 4 of oil. Because of this prospect, another oil company has offered to purchase the land for \$90.000. However, GC is considering holding the land in order to drill for oil itself. The cost of drilling is \$100.000. If oil is found, the resulting expected revenue will be \$800.000, so the company's expected profit (after deducting the cost of drilling) will be \$700.000. A loss of \$100.000 (the drilling cost) will be incurred if the land is dry (no oil).

	State of nature	
Alternatives	Oil	Dry
Drill for oil	700	-100
Sell the land	90	90
Prior probability	0.25	0.75

- ***The Maximin Payoff Criterion***
  For each possible decision alternative, find the minimum payoff over all possible states of nature. Next, find the maximum of these minimum payoffs. Choose the alternative whose minimum payoff gives this maximum.
  For GC the optimal choice is to sell the land.
- ***The Maximum Likelihood Criterion***
  Identify the most likely state of nature (the one with the largest prior probability). For this state of nature, find the decision alternative with the maximum payoff. Choose this decision alternative.
  For GC the choice is to sell the land.

- *Bayes' Decision Rule*

  Using the best available estimates of the probabilities of the respective states of nature (currently the prior probabilities), calculate the expected value of the payoff (expected payoff) for each of the possible decision alternatives. Choose the decision alternative with the maximum expected payoff.

  For GC the optimal choice is to drill for oil because

$$\mathbb{E}[Payoff\,(drill)] = 100 > 90 = \mathbb{E}[Payoff\,(sell)]$$

- *Sensitivity Analysis with Bayes' Decision Rule*

  - prior probabilities: interval
  - payoffs.

p - prior probability of oil; the decision is very *sensitive* to p.

$$\mathbb{E}[Payoff(drill)] = 700p - 100(1-p) = 800p - 100.$$

$$\mathbb{E}[Payoff(sell)] = 90$$

The crossover point is $p = 0.2375$.

Conclusion: Sell the land if $p < 0.2375$ and drill for oil otherwise.

The key question when trying to refine the true value of p is whether it is smaller or larger than 0.2375.

Posterior probabilities: Improved estimates of prior probabilities. *Important questions*:

1. How to derive the posterior probabilities?
2. How to decide whether it is worthwhile to conduct experimentation (taking into account the cost of experimentation)?

Notations:

- n: number of possible states of nature
- $P(State = S_i)$: prior probability that true state of nature is S_i, for $i = \overline{1, n}$.
- $Finding$: finding from experimentation (a random variable).
- f_j: one possible value of the random value $Finding$, $j = \overline{1, m}$.
- $P(State = S_i | Finding = f_j)$: posterior *probability* that the true state of nature is S_i, given that $Finding = f_j$, for $i = \overline{1, n}, j = \overline{1, m}$.

Decision making with experimentation

Basic question: given $P(State = S_i)$ and $P(Finding = f_j | State = S_i)$ for $i = \overline{1,n}, j = \overline{1,m}$, what is

$$P(State = S_i | Finding = f_j) = ?$$

From Bayes' theorem

$$P(State = S_i | Finding = f_j) =$$

$$= \frac{P(Finding = f_j | State = S_i) \cdot P(State = S_i)}{\sum_{k=1}^{n} P(Finding = f_j | State = S_k) \cdot P(State = S_k)}.$$

Example continues. For GC, *possible findings* of a seismic survey are:
USS (Unfavorable Seismic Soundings): oil is fairly unlikely;
FSS (Favorable Seismic Soundings): oil is fairly likely.

Based on *past experience:*

$$P(USS|State = Oil) = 0.4 \Rightarrow P(FSS|State = Oil) = 0.6;$$

$$P(USS|State = Dry) = 0.8 \Rightarrow P(FSS|State = Dry) = 0.2.$$

If the finding of the seismic survey is USS then

$$P(State = Oil|Finding = USS) = 0.4 \cdot 0.25/(0.4 \cdot 0.25 + 0.8 \cdot 0.75) = 1/7$$

$$\Rightarrow P(State = Dry|Finding = USS) = 6/7.$$

If the seismic survey gives FSS then

$$P(State = Oil|Finding = FSS) = 0.6 \cdot 0.25/(0.6 \cdot 0.25 + 0.2 \cdot 0.75) = 1/2$$

$$\Rightarrow P(State = Dry|Finding = FSS) = 1/2.$$

The *probability tree diagram* in figure 1 shows a nice way of organizing these calculations in an intuitive manner.

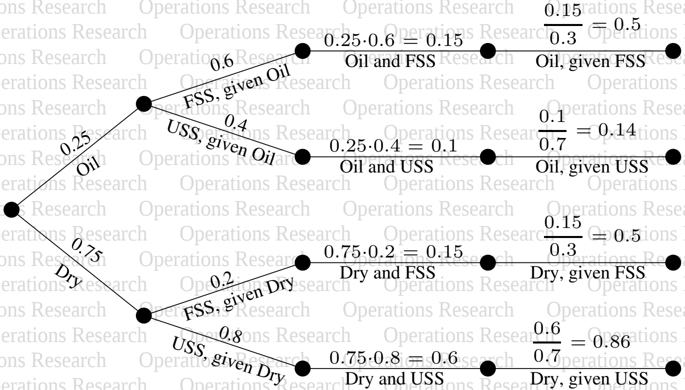

Figure: The probability tree diagram for GC example.

Suppose the cost of experimentation is 30. Apply Bayes' decision rule with the posterior probabilities replacing the prior probabilities. Expected payoffs if finding is USS are:

$$\mathbb{E}[Payoff(drill|Finding = USS)] = 1/7 \cdot 700 + 6/7 \cdot (-100) - 30 = -15.7,$$

$$\mathbb{E}[Payoff(sell|Finding = USS] = 1/7 \cdot 90 + 6/7 \cdot 90 - 30 = 60.$$

Expected payoffs if finding is FSS are:

$$\mathbb{E}[Payoff(drill|Finding = FSS)] = 1/2 \cdot 700 + 1/2 \cdot (-100) - 30 = 270,$$

$$\mathbb{E}[Payoff(sell|Finding = FSS] = 1/2 \cdot 90 + 1/2 \cdot 90 - 30 = 60.$$

Since the objective is to maximize the expected payoff, these results yield the optimal policy shown below:

Finding from seismic survey	Optimal alternative	Expected payoff excluding cost of survey	Expected payoff including cost of survey
USS	Sell the land	90	60
FSS	Drill for oil	300	270

We present two complementary methods for evaluating the experimentation potential value:

- Expected Value of Perfect Information (EVPI) provides an *upper bound* of the potential value of the experiment. If the upper bound is less than the cost of the experiment, the experiment may be forgone.

  Expected payoff with perfect information (EPPI) is obtained by weighting the maximum payoff for each state of nature by the prior probability of that state of nature.

$$EVPI = EPPI - EPWE,$$

where $EPWE$ is the *expected payoff without experimentation*.

Example continues

$$EPPI = 0.25 \cdot 700 + 0.75 \cdot 90 = 242.5,$$

$$EPWE = 0.25 \cdot 700 + 0.75 \cdot (-100) = 100,$$

$$EVPI = 242.5 - 100 = 142.5$$

Conclusion: since $142.5 > 30$ it may be worthwhile to proceed with the seismic survey. To find out for sure, we now go to the second method of evaluating the potential benefit of experimentation.

- *Expected Value of Experimentation (EVE)* is the difference between *expected payoff with experimentation (EPE)* and expected payoff without experimentation:

$$EVE = EPE - EPWE,$$

where expected payoff with experimentation is

$$EPE = \sum_{j=1}^{m} P(Finding = f_j) \cdot \mathbb{E}[payoff \,|\, Finding = f_j].$$

Example continues: $P(USS) = 0.7$, $P(FSS) = 0.3$ - obtained by summation from data in Figure 1. From the third column of the previous table we have:

$$\mathbb{E}(Payoff \,|\, Finding = USS) = 90,$$

$$\mathbb{E}(Payoff \,|\, Finding = FSS) = 300.$$

Thus,

$$EPE = 0.7 \cdot 90 + 0.3 \cdot 300 = 153, EVE = 153 - 100 = 53.$$

Conclusion: since $53 > 30$ the seismic experiment should be done.

- A *decision node*, represented by a square, indicates that a decision needs to be made at that point in the process.
- An *event node* (or chance node), represented by a circle, indicates that a random event occurs at that point.
- *Branches* have labels: numbers (that are not inserted in parentheses) inserted under or over the branches represent payoffs, while number in parentheses represent prior or posterior probabilities, or probabilities of findings.
- At each terminal node is attached the sum of all numbers along the path from the initial node.

Constructing the Decision Tree

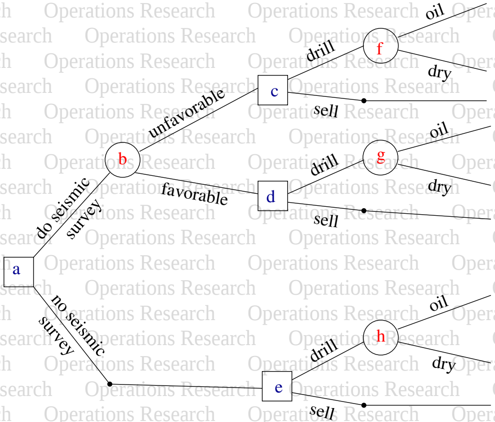

Figure: The decision tree (before including any numbers) for GC example.

- *Step 1.* Start at the right side of the decision tree and move left one column at a time. For each column, perform either Step 2 or Step 3 depending upon whether the nodes in that column are event nodes or decision nodes.
- *Step 2.* For each event node, calculate its *expected payoff* by multiplying the expected payoff of each branch by the probability of that branch and then adding up these products. Record this expected payoff in boldface above the node, and designate this quantity as also being the expected payoff of the branch leading to this node.
- *Step 3.* For each decision node, compare the expected payoffs of its branches and choose the alternative whose branch has the largest expected payoff. Record this largest expected payoff in boldface above the decision node and record the choice on the decision tree by inserting a barrier through each rejected branch.

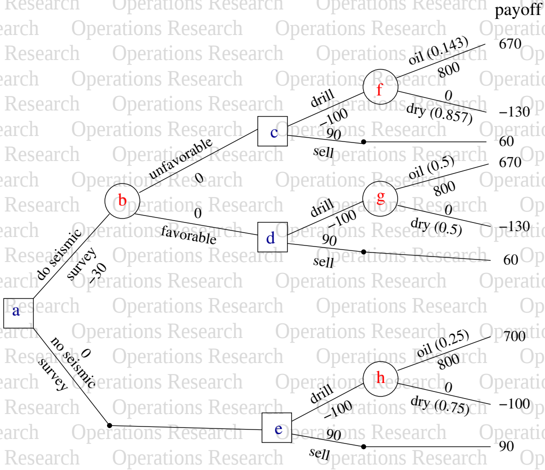

Figure: The decision tree (after adding the probabilities and payoffs).

- For any decision tree, this *backward induction procedure* always will lead to the *optimal policy* (or policies) after the probabilities are computed for the branches emanating from a chance node.
- The decision maker now can read the tree from left to right to see the actual progression of events.
- Bayes' decision rule says to follow only the open paths from left to right to achieve the largest possible expected payoff.

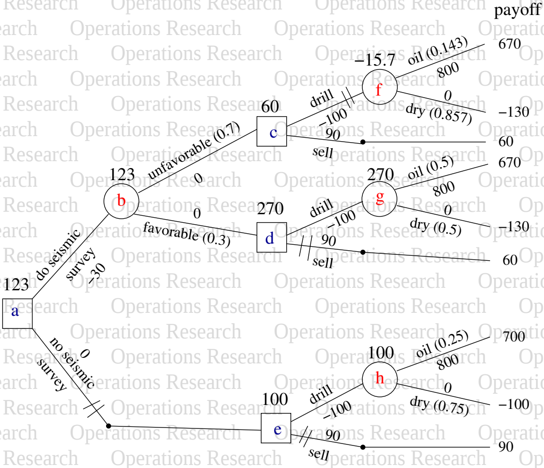

Figure: The final decision tree that records the analysis.

Example continues *Optimal policy:* Do the seismic survey.

If the result is unfavorable, sell the land.

If the result is favorable, drill for oil.

The expected payoff (including the cost of the seismic survey) is 123 ($123,000).

- *Risk-averse* individuals have *decreasing marginal utility for money* (the slope of their utility function *decreases* as the amount of money increases).
- *Risk-seekers* have *increasing marginal utility for money* (the slope of their utility function *increases* as the amount of money increases).
- *Risk-neutral* individuals prize money at their face value ($u(M) = M$). Such an individual's utility for money is simply proportional to the amount of money involved.

The fact that different people have different utility functions for money has an important implication for decision making in the face of uncertainty.

- It is possible to exhibit a mixture of these kinds of behavior depending on the amount of money at stake.
- Although some people appear to be risk-neutral when only small amounts of money are involved, it is unusual to be truly risk-neutral with very large amounts.
- An individual's attitude toward risk also may be different when dealing with one's personal finances than when making decisions on behalf of an organization.

When an *utility function for money* is incorporated into a decision analysis approach to a problem, the utility function must be constructed to fit the preferences and values of the decision maker (a single individual or a group of people) involved.

Fundamental Property: The decision maker's utility function for money has the property that the decision maker is *indifferent* between two alternative courses of actions if the two alternatives have the same *expected utility*.

- When the decision maker's utility function for money is used to measure the relative worth of the various possible monetary outcomes, Bayes' decision rule replaces monetary payoffs by the corresponding utilities. Therefore, the optimal action (or series of actions) is the one which *maximizes the expected utility*.

- Utility functions can sometimes still be constructed when some of or all the important consequences of the alternative courses of action are *not monetary*. Under these circumstances, it is important to incorporate value judgments into the decision process.
- *One way to construct* $u(M)$ is to ask the decision maker to repeatedly make a difficult decision about which probability would make him/her indifferent between two alternatives. Then using pairs of points that relate possible payoffs and the corresponding utilities, a smooth curve can be drawn through these points.

Another approach for estimating $u(M)$ is to assume that the utility function has a certain mathematical form, and then adjust this form to fit the decision maker's attitude toward risk as closely as possible.

- The *exponential utility function*: $u(M) = R(1 - e^{-M/R})$, where R is the decision maker's *risk tolerance*. This function is designed to fit a *risk-averse* individual.
- A great aversion to risk corresponds to a small value of R (which gives a much more gradual bend in the curve).
- A drawback of the exponential utility function is that it assumes a constant aversion to risk (a fixed value of R), regardless of how much (or how little) money the decision maker currently has.

- In situations where the consequences of the potential losses are not as severe, assuming an exponential utility function may provide a reasonable approximation.
- An easy way of estimating the appropriate value of R is by asking the decision maker to choose a number that would make him/her indifferent between the following two alternatives:
  - A 50:50 gamble where he/she would gain R dollars with probability 0.5 and lose $R/2$ dollars with probability 0.5.
  - Neither gain nor lose anything.

- The procedure for using a decision tree to analyze a problem now is identical to that described previously *except* for substituting utilities for monetary payoffs.
- Therefore, the value obtained to evaluate each node of the tree now is the *expected utility* there rather than the expected (monetary) payoff.
- Consequently, the optimal decisions selected by Bayes 'decision rule maximize the expected utility for the overall problem.

Bibliography

- Hillier, F. S., G. J. Lieberman, *Introduction to Operations Research*, McGraw-Hill, 7th edition, 2001.
- Taha, H. A., *Operations Research: An Introduction*, Prentice Hall International, 2007.