Decision Theory

VK
Room: M1.30
knightva@cf.ac.uk
www.vincent-knight.com

Last updated October 3, 2012.

Overview

Non-Probabilistic Decision Making

Decision Analysis

Decision Analysis refers to a set of methodologies based on expected values, maximin, and related criteria that are used to select the best alternative when a decision maker is faced with uncertainty.

What is the best decision?

Suppose we own a large plot of land that may contain oil. The present value of the land is 10 million. It costs 20 million to drill for oil. If oil is found then we expect to earn 80 million, but if no oil is found then we can only sell the land for 5 million.

We can either sell the land or drill for oil. There are two possible states of the land (oily or dry), so there are four possible outcomes.

What is the best decision?

Four possible outcomes:

	Oily	Dry
Drill	80-20=60	5-20=-15
Sell	10	10

≥ 3 approaches to solving this problem.

MaxMax Payoff approach

This approach finds the *best* case scenario. For each decision we identify the best outcome (maximum payoff) over all possible states. We then find the *max*imum of these *max*imum payoffs.

	Oily	Dry	Max
Drill	60	-15	60
Sell	10	10	10
Max			60

MaxMin Payoff approach

This approach finds the *best worst* case scenario. For each decision we identify the worst outcome (minimum payoff) over all possible states. We then find the *max*imum of these *min*imum payoffs.

	Oily	Dry	Min
Drill	60	-15	-15
Sell	10	10	10
Max			10

MinMax Regret approach

This approach first identifies the regret relevant to each decision. For each decision we identify the distance from the best possible decision for a particular outcome. We then find the *min*imum of these *max*imum regrets.

	Oily	Dry	Oily Regret	Dry Regret	Max Regret
Drill	60	-15	0	25	25
Sell	10	10	50	0	50
Min					25

Non-Probabilistic Decision Making - Summary

3 approaches:

- Maxmax (risk-seeking)
- Maxmin (risk-averse)
- Minmax regret (risk-neutral)

Non-Probabilistic Decision Making - Summary

3 approaches:

- Maxmax (risk-seeking)
- Maxmin (risk-averse)
- Minmax regret (risk-neutral)

None of these approaches take in to account the likelihood of an outcome!

Decision Making Under Uncertainty

Extra Information

Comparison with nearby land tells us that the chance (prior probability) of finding oil is 0.6.

Maximum likelihood Payoff approach

This approach finds the *most likely* state. For the most likely state we identify the best decision (maximum payoff). In this approach we ignore all the states that are less likely to happen.

	Oily	Dry	Most likely
Drill	80-20=60	5-20=-15	60
Sell	10	10	10
Probability	.6	.4	

Maximum expected value Conclusion

This approach (also called the Bayes' decision rule approach) finds the *best most likely* payoff. We compute the expected payoffs for all the possible decisions. We then choose the one with the maximum expected payoff.

	Oily	Dry	Expected Payoff
Drill	80-20=60	5-20=-15	.6(60)+.4(-15)=30
Sell	10	10	.6(10)+.4(10)=10
Probability	.6	.4	

A new strategy

Suppose that some oil-drilling company has investigated the land and suggests the following proposal. The company will provide a drilling service at a lower cost, with the condition that if there is oil then the company gets $\frac{3}{8}$ of the profit but will not charge at all for the cost of drilling. If there is no oil found, then we have to pay 2 million for the drilling.

New pay off matrix

	Oily	Dry
Drill	80-20=60	5-20=-15
Sell	10	10
Accept Offer	$(0.625)80 = 50$	5-2=3
Probability	.6	.4

Choosing the best approach

	Oily	Dry	Max	Min	Max Regret	ML	EP
Drill	60	-15	60	-15	25	60	30
Sell	10	10	10	10	50	10	10
Accept Offer	50	3	50	3	10	50	31.2
Probability	.6	.4					

Conclusion

- If the MaxMax approach is used then the optimal decision is to Drill. (high risk)
- If the MaxMin approach is used then the optimal decision is to sell the land directly. (risk averse)
- If the MinMax Regret approach is used then the optimal decision is to accept the offer. (risk neutral)
- If the maximum likelihood approach is used then the optimal decision is to drill without accepting the offer. (high risk)
- If we apply the maximum expected value approach, the optimal decision is to drill whilst accepting the offer. (risk neutral)

Bayes' Theorem

Cause of accidents

Consider the following two events:

- A: There is an obstacle on the road.
- B: I have an accident.

Cause of accidents

Consider the following two events:

- A: There is an obstacle on the road.
- B: I have an accident.

Let $P(A) = .3$, $P(B) = .5$ and $P(B|A) = \frac{2}{3}$.

Cause of accidents

Consider the following two events:

- A: There is an obstacle on the road.
- B: I have an accident.

Let $P(A) = .3$, $P(B) = .5$ and $P(B|A) = \frac{2}{3}$.

If I've had an accident what is the chance that it was caused by an obstacle on the road?

Cause of accidents

Consider the following two events:

- A: There is an obstacle on the road.
- B: I have an accident.

Let $P(A) = .3$, $P(B) = .5$ and $P(B|A) = \frac{2}{3}$.

If I've had an accident what is the chance that it was caused by an obstacle on the road?

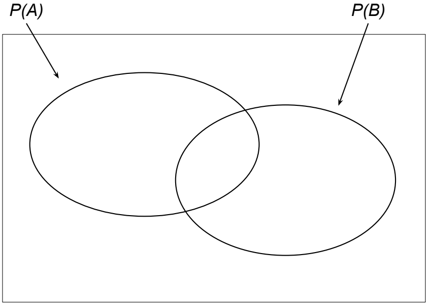

Cause of accidents

Consider the following two events:

- A: There is an obstacle on the road.
- B: I have an accident.

Let $P(A) = .3$, $P(B) = .5$ and $P(B|A) = \frac{2}{3}$.

If I've had an accident what is the chance that it was caused by an obstacle on the road?

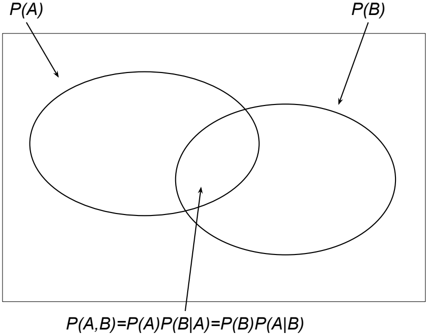

Cause of accidents

Consider the following two events:

- A: There is an obstacle on the road.
- B: I have an accident.

Let $P(A) = .3$, $P(B) = .5$ and $P(B|A) = \frac{2}{3}$.

If I've had an accident what is the chance that it was caused by an obstacle on the road?

$$P(A)P(B|A) = P(B)P(A|B)$$

$$\Leftrightarrow$$

$$P(A|B) = \frac{P(A)P(B|A)}{P(B)} = \frac{.3\frac{2}{3}}{.5} = \frac{2}{5}$$

Bayes' Theorem

The difficulty of decision making lies with the "uncertainty of state". Very often experiments can be done to improve our prior estimates of the probabilities of each state. The improved estimates are called *posterior probabilities*. A useful tool for calculating posterior probabilities is *Bayes' Theorem*:

Theorem

Suppose that $A_1, A_2, \ldots, A_n$ *are mutually exclusive events and the union of these events is the entire sample space, then for any other event* B*,*

$$\begin{aligned} P(A_i \mid B) &= \frac{P(A_i)P(B \mid A_i)}{P(B)} \\ &= \frac{P(A_i)P(B \mid A_i)}{P(A_1)P(B \mid A_1)+\cdots+P(A_n)P(B \mid A_n)} \end{aligned}$$

Example

In a nuclear plant, there are 3 main causes of accidents:

- Human error (H):

$$P(H) = .02, \; P(E \mid H) = .01$$

- Mechanical error (M):

$$P(M) = .01, \; P(E \mid M) = .05$$

- Natural disaster (D):

$$P(D) = .001, \; P(E \mid D) = .1$$

The nuclear plant has exploded. What is the probability that it is due to a mechanical error?

Solution

Recall:

$$P(H) = .02,\ P(M) = .01,\ P(D) = .001$$

and

$$P(E \mid H) = .01,\ P(E \mid M) = .05,\ P(E \mid D) = .1$$

thus:

$$\begin{aligned} P(M \mid E) &= \frac{P(M)P(E \mid M)}{P(H)P(E \mid H)+P(M)P(E \mid M)+P(D)P(E \mid D)} \\ &= \frac{.05\times.01}{.01\times.02+.05\times.01+.1\times.001} \end{aligned}$$

Decision Trees

Decision trees are used to calculate expected payoffs for different decisions. We will illustrate this with the following problem. A company is considering an investment in a certain stock. The company assesses that the stock has a 60% chance of going up, in which case they can make a profit of 20. If the stock goes down they will lose 20.

The company has the option of paying a financial advisor C thousand for an assessment of the stock (we will allow C to vary). The advisor is known to be 80% successful at forecasting a stock increase and 70% successful at forecasting a stock decrease. We will draw the decision tree for this problem and describe how it is constructed.

Step 1: Set out Decisions and Variable outcomes

Firstly we set out the possible decisions and variable outcomes. Decisions are denoted by square boxes and variable outcomes by circles.

Step 1: Set out Decisions and Variable outcomes

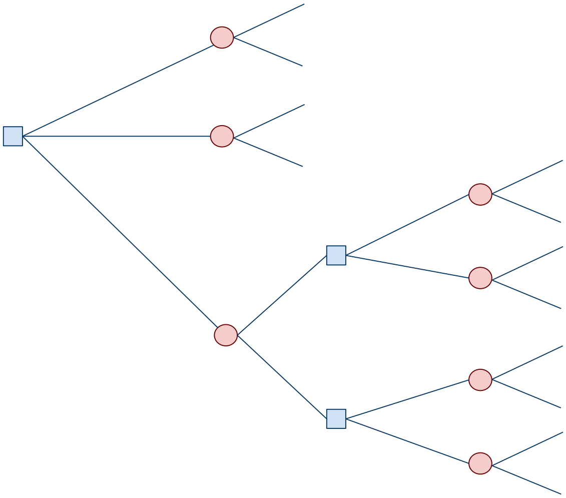

Step 1: Set out Decisions and Variable outcomes

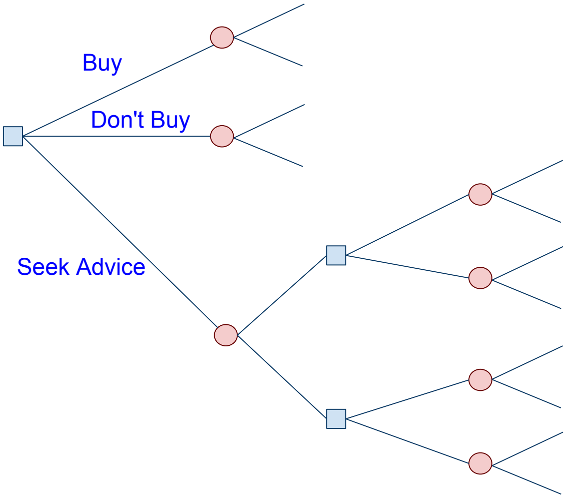

Step 1: Set out Decisions and Variable outcomes

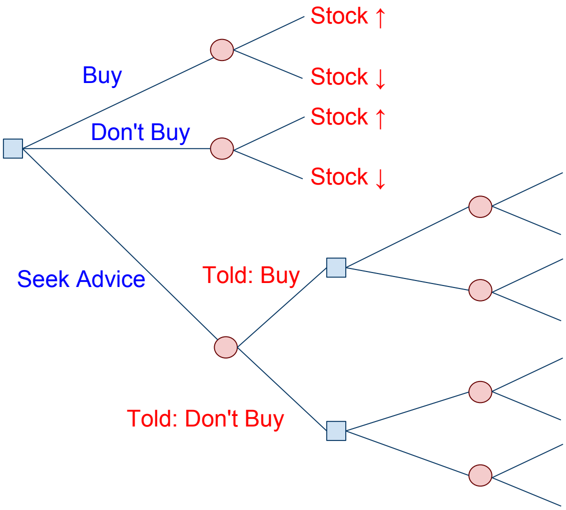

Step 1: Set out Decisions and Variable outcomes

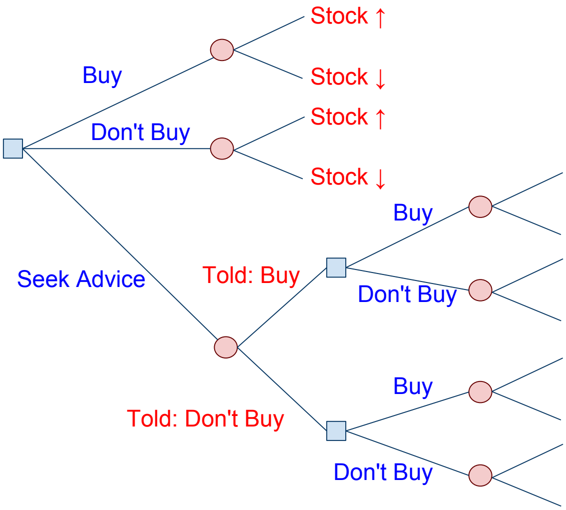

Step 1: Set out Decisions and Variable outcomes

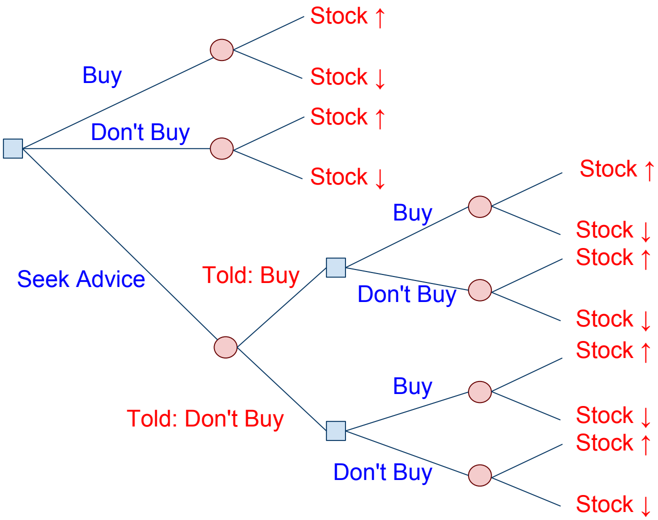

Step 2: Calculate the Returns

Next we work out the return for each possible combination of decisions and outcomes.

Step 2: Calculate the Returns

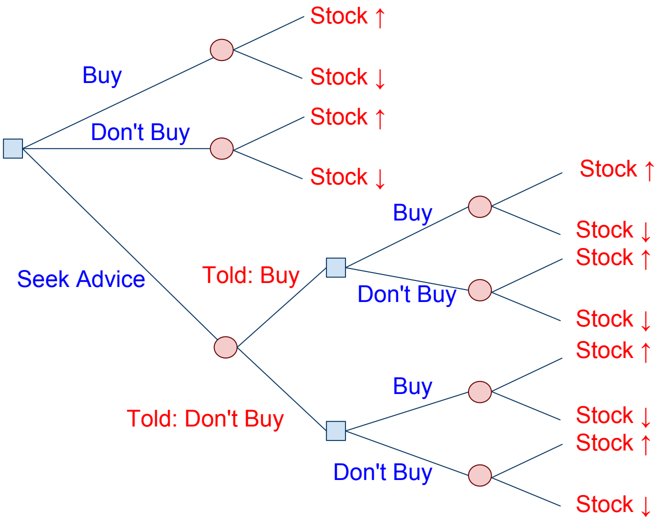

Step 2: Calculate the Returns

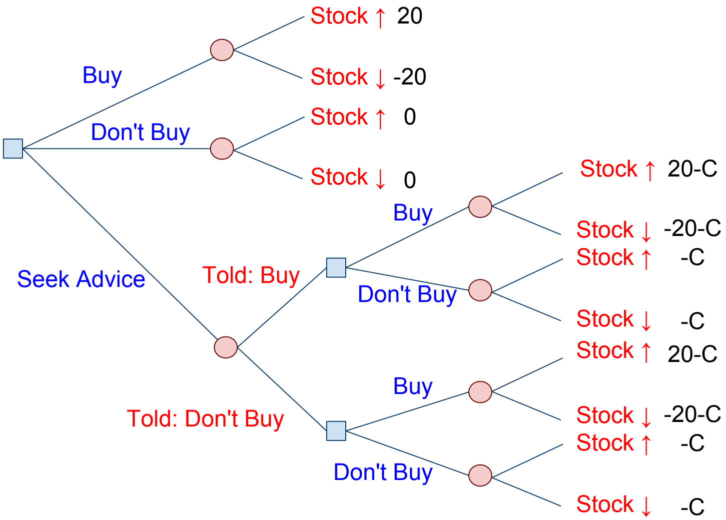

Step 3: Probabilities of Outcomes

Next we work out the the probabilities of each variable outcome ( ). This is done from left to right. Note that all probabilities are conditional on outcomes to the left.

Let U be the event *Stock Goes Up*, D the event *Stock Goes Down*, Y the event *Told: Buy* and N the event *Told: Don't Buy*. We have:

$$P(U) = .6,\; P(D) = .4,\; P(Y \mid U) = .8\,, P(N \mid D) = .7$$

To complete the decision tree we need to know $P(Y),\; P(N),\; P(U \mid Y)$ and $P(D \mid N)$:

$$P(Y) = P(Y \mid U)P(U) + P(Y \mid D)P(D) = .6$$
$$P(N) = 1 - P(Y) = .4$$

Step 3: Probabilities of Outcomes

We use Bayes' Theorem to calculate $P(D \mid Y)$ and $P(U \mid N)$:

$$P(U \mid Y) = .8$$
$$P(D \mid Y) = .2$$
$$P(U \mid N) = .3$$
$$P(D \mid N) = .7$$

Step 3: Probabilities of Outcomes

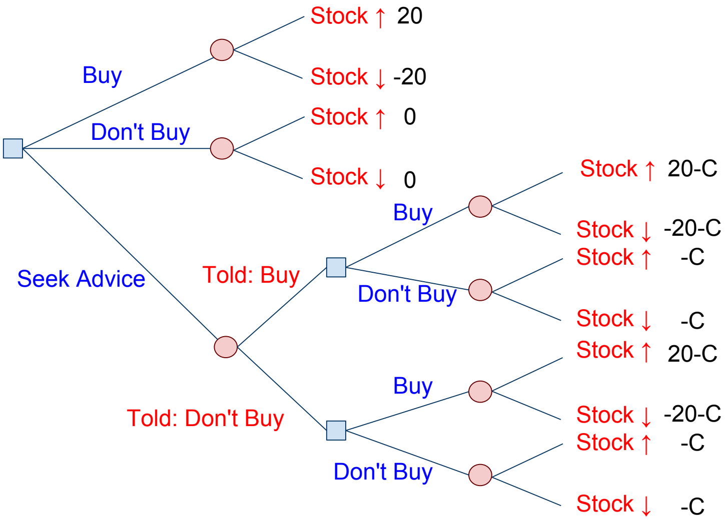

Step 3: Probabilities of Outcomes

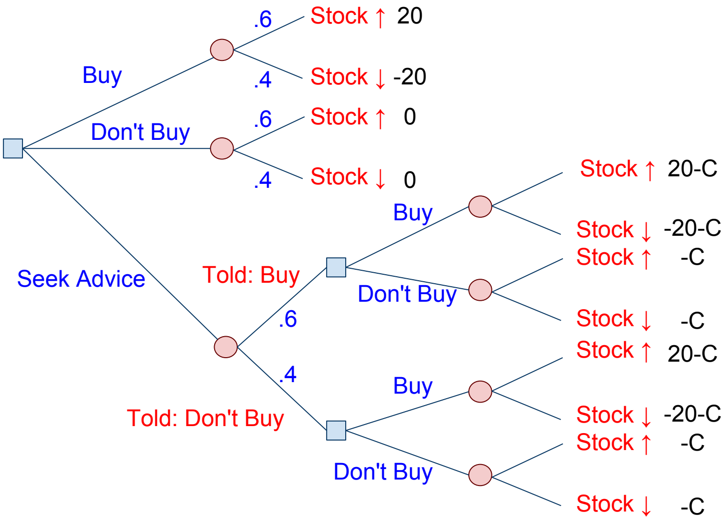

Step 3: Probabilities of Outcomes

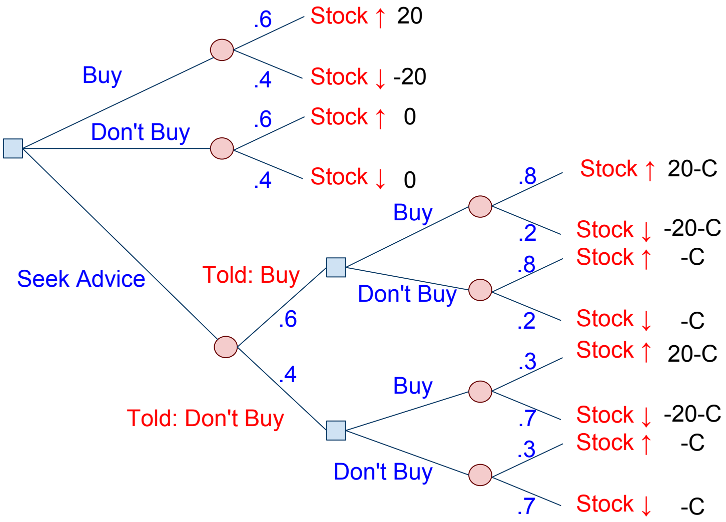

Step 4: Expected Returns

Using the conditional probabilities we can calculate the expected return at each options node ( ), working from right to left.

By choosing the decisions that maximise the expected return, we can also work out the expected return at each decision node ( ).

Step 4: Expected Returns

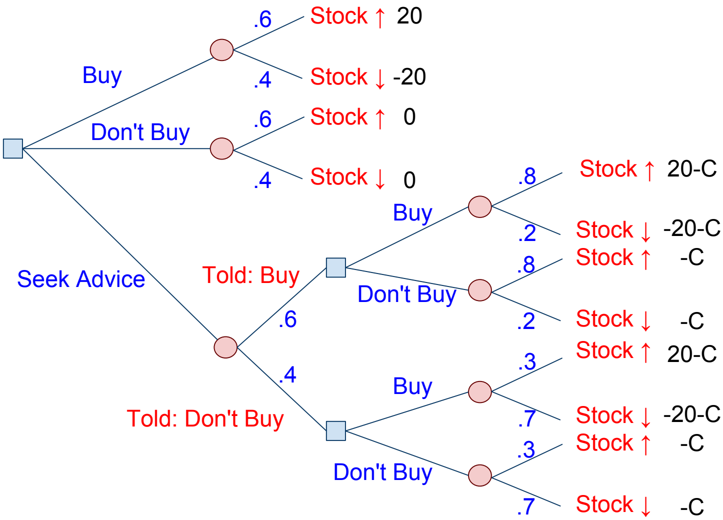

Step 4: Expected Returns

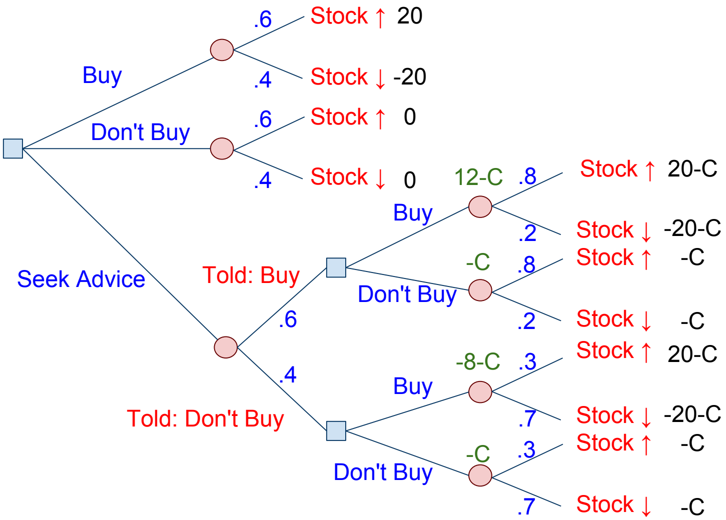

Step 4: Expected Returns

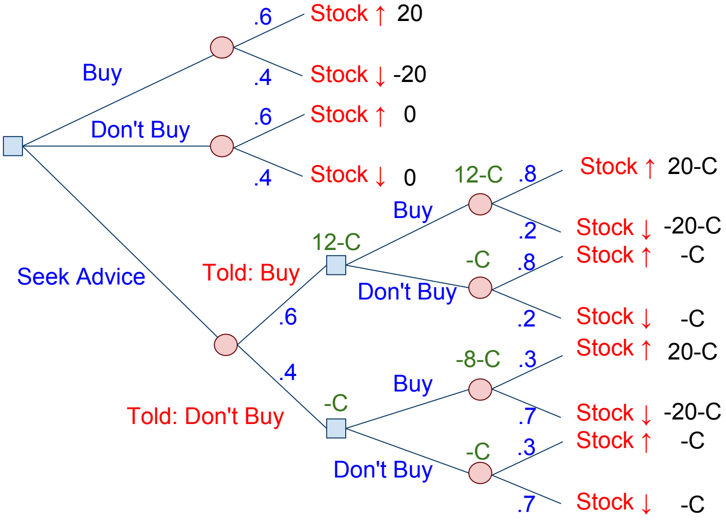

Step 4: Expected Returns

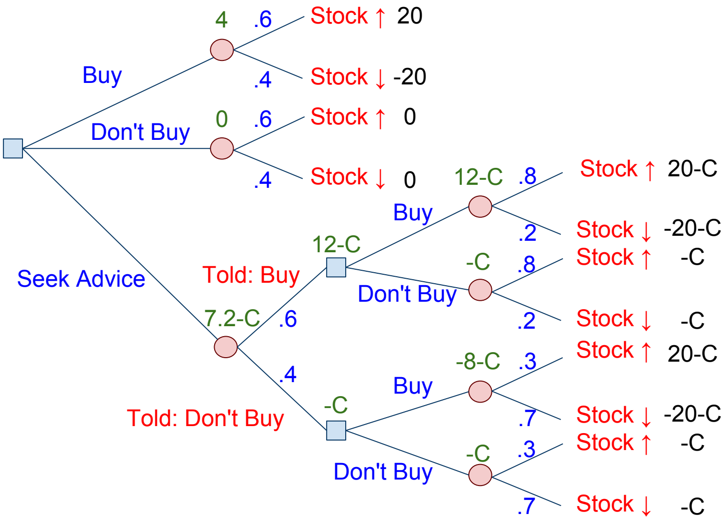

Example

Consider the following example:

- Decide to not toss a coin and receive 4.
- Toss an unbiased coin :
  - heads: receive 10.
  - tails: receive nothing.

Example

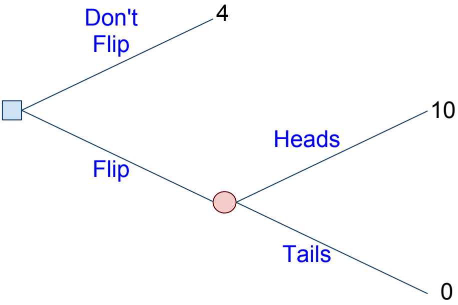

Example

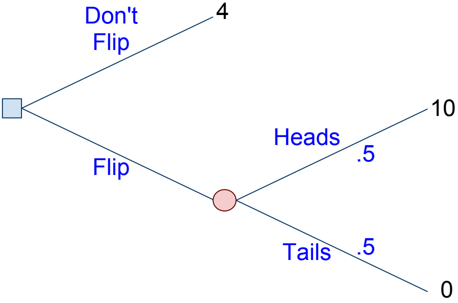

Example

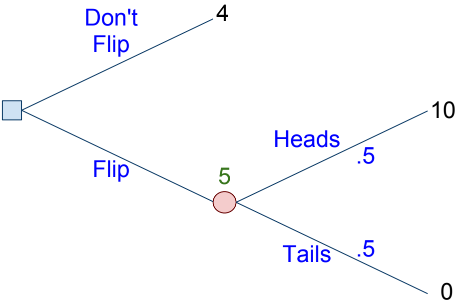

Utility functions

Utility Functions

We would like to reflect different opinions about the value of certain gain versus uncertain gain, but still be able to make use of decision trees. We need a way of transforming absolute gain to an appropriate scale that reflects the decision maker's preference. This scale is called a *utility function*.

Example

For example consider the utility function: $u : \mathbb{R} \to \mathbb{R}$:

$$u(x) = \sqrt{x}$$

Example

For example consider the utility function: $u : \mathbb{R} \to \mathbb{R}$:

$$u(x) = \sqrt{x}$$

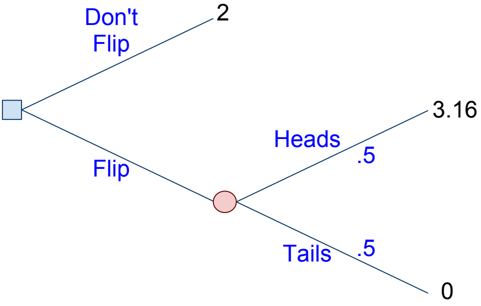

Example

For example consider the utility function: $u : \mathbb{R} \to \mathbb{R}$:

$$u(x) = \sqrt{x}$$

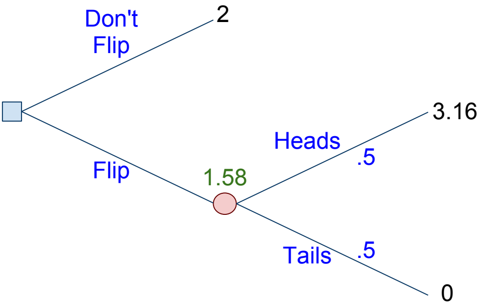

Utility Functions

3 categories of utility functions:

- **Risk-averse**: the utility function is risk-averse if it is concave:

$$u(x) = \sqrt{x}$$

- **Risk-seeking**: the utility function is risk-seeking if it is convex:

$$u(x) = x^2$$

- **Risk-neutral**: the utility function is risk-neutral if is is linear:

$$u(x) = x$$

Utility Functions

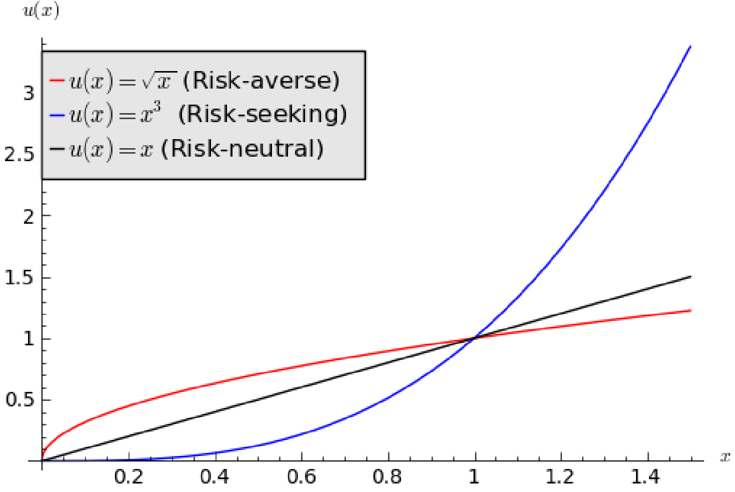

Choice of utility function

When a utility function is used for decision analysis, the utility function must be constructed to fit the preferences and values of the decision maker.

Munduruku Example

In "Alex's Adventures in Numberland" by **Alex Bellos** an account of the **Munduruku** tribe is given.

- Indigenous tribe of 7000.
- Brazilian Amazon.
- Language:
    - No tense.
    - No plurals.
    - No words for numbers > 5.

Numerical Experiment

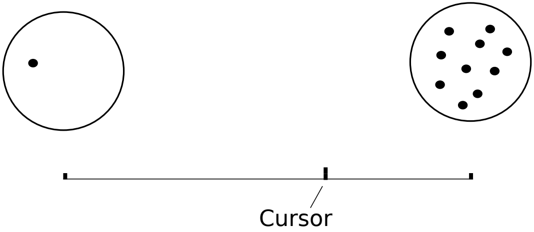

Results

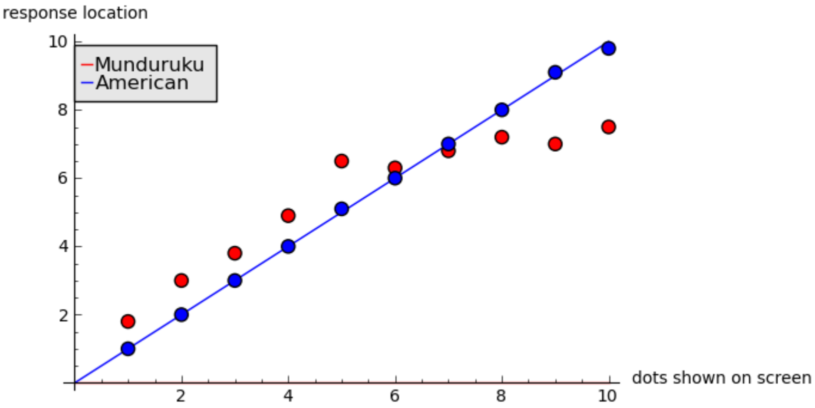

response location
10
8
6
4
2
Munduruku
American
2
4
6
8
10
dots shown on screen