

Our Journey



■ Information for Decision

■ Formalise

- Utility Theory, Expected Utility and Bayesian Decision Theory
- Examples
 - Decision Problems
 - Decision Processes
 - Decision Systems
 - Markets



Recap

- We introduced Graphical Models
- We talked about Inference in Graphical Models
- But...
- That's just information processing
- What do we do with that information?
- Use it in Decision Making
- Any system must make a decision to be useful, even if that decision is just what information to provide.



Utility Theory

- Utility indicates the subject value of an outcome or state of affairs to an individual.
- Usually use U to denote it.
- Utility is a function of many things.
- E.g. $U(\text{egg, eggcup}) > U(\text{egg, no eggcup}) + U(\text{no egg, eggcup})$.
- Utility simply indicates preference orderings.
 - It is an real valued ordinal quantity – only the order matters, not the explicit numbers.
 - Monotonic transformations of utilities are equivalent.
- Utility is actually about choice: maximum utility.



Expected Utility Theory

- But what about utility under uncertainty?
- Think about uncertainty as many possible worlds.
- Prefer lottery that increases chances of preferred things happening.
- **von Neumann-Morgenstern axioms**
 - Expected Utility Theory
 - $\max E p(x)(U(x))$



From Wikipedia

Main article: Von Neumann–Morgenstern utility theorem

The von Neumann-Morgenstern axioms [\[edit\]](#)

There are four axioms^[5] of the expected utility theory that define a *rational* decision maker. They are completeness, transitivity, independence and continuity.

Completeness assumes that an individual has well defined preferences and can always decide between any two alternatives.

- Axiom (Completeness): For every A and B either $A \succeq B$ or $A \preceq B$.

This means that the individual either prefers A to B, or is indifferent between A and B, or prefers B to A.

Transitivity assumes that, as an individual decides according to the completeness axiom, the individual also decides consistently.

- Axiom (Transitivity): For every A, B and C with $A \succeq B$ and $B \succeq C$ we must have $A \succeq C$.

Independence also pertains to well-defined preferences and assumes that two gambles mixed with a third one maintain the same preference order as when the two are presented independently of the third one. The independence axiom is the most controversial one.

- Axiom (Independence): Let A, B, and C be three lotteries with $A \succeq B$, and let $t \in (0, 1]$; then $tA + (1 - t)C \succeq tB + (1 - t)C$.

Continuity assumes that when there are three lotteries (A, B and C) and the individual prefers A to B and B to C, then there should be a possible combination of A and C in which the individual is then indifferent between this mix and the lottery B.

- Axiom (Continuity): Let A, B and C be lotteries with $A \succeq B \succeq C$; then there exists a probability p such that B is equally good as $pA + (1 - p)C$.

If all these axioms are satisfied, then the individual is said to be rational and the preferences can be represented by a utility function, i.e. one can assign numbers (utilities) to each outcome of the lottery such that choosing the best lottery according to the preference \succeq amounts to choosing the lottery with the highest expected utility. This result is called the [von Neumann–Morgenstern utility representation theorem](#).



Example

Expected Utility

You are asked if you wish to take a bet on the outcome of tossing a fair coin. If you bet and win, you gain £100. If you bet and lose, you lose £200. If you don't bet, the cost to you is zero.

$$\begin{aligned}U(\text{win, bet}) &= 100 & U(\text{lose, bet}) &= -200 \\U(\text{win, no bet}) &= 0 & U(\text{lose, no bet}) &= 0\end{aligned}$$

Our expected winnings/losses are:

$$\begin{aligned}U(\text{bet}) &= p(\text{win}) \times U(\text{win, bet}) + p(\text{lose}) \times U(\text{lose, bet}) \\&= 0.5 \times 100 - 0.5 \times 200 = -50\end{aligned}$$

$$U(\text{no bet}) = 0$$

Based on taking the decision which maximises expected utility, we would therefore be advised not to bet.



Example

Utility of Money

You have £1,000,000 in your bank account. You are asked if you would like to participate in a fair coin tossing bet in which, if you win, your bank account will become £1,000,000,000. However, if you lose, your bank account will contain only £1000. Should you take the bet?

$$U(\text{bet}) = 0.5 \times 1,000,000,000 + 0.5 \times 1000 = 500,000,500.00$$

$$U(\text{no bet}) = 1,000,000$$

Based on expected utility, we are therefore advised to take the bet.



Example

Utility of Money

You have £1,000,000 in your bank account. You are asked if you would like to participate in a fair coin tossing bet in which, if you win, your bank account will become £1,000,000,000. However, if you lose, your bank account will contain only £1000. Should you take the bet?

$$U(\text{bet}) = 0.5 \times 1,000,000,000 + 0.5 \times 1000 = 500,000,500.00$$


$$U(\text{no bet}) = 1,000,000$$

Based on expected utility, we are therefore advised to take the bet.



Concavity of Utility

- Millionaires are unlikely to be willing to risk losing almost everything in order to become a billionaire.
- Utility of money is not just quantity of money.
- Utility is concave. Risk is penalised.



Subjective Expected Decision Theory

- Or Savage Decision Theory casts DT in terms of subjective decision making.
- The Expectation is wrt Posterior Belief.
- Actions \rightarrow Outcomes \rightarrow Utilities \rightarrow Decisions.
- Actually can augment our Bayesian Networks to include action and reward nodes.



Example

Should I do a PhD?

Consider a decision whether or not to do PhD as part of our education (E). Taking a PhD incurs costs, U_C both in terms of fees, but also in terms of lost income. However, if we have a PhD, we are more likely to win a Nobel Prize (P), which would certainly be likely to boost our Income (I), subsequently benefitting our finances (U_B). The ordering is (excluding empty sets)

$$E^* \prec \{I, P\}$$

$\text{dom}(E) = (\text{do PhD, no PhD})$, $\text{dom}(I) = (\text{low, average, high})$,
 $\text{dom}(P) = (\text{prize, no prize})$.

$$p(\text{win Nobel prize}|\text{no PhD}) = 0.0000001$$

$$p(\text{win Nobel prize}|\text{do PhD}) = 0.001$$

$$p(\text{low}|\text{do PhD, no prize}) = 0.1$$

$$p(\text{low}|\text{no PhD, no prize}) = 0.2$$

$$p(\text{low}|\text{do PhD, prize}) = 0.01$$

$$p(\text{low}|\text{no PhD, prize}) = 0.01$$

$$p(\text{average}|\text{do PhD, no prize}) = 0.5$$

$$p(\text{average}|\text{no PhD, no prize}) = 0.6$$

$$p(\text{average}|\text{do PhD, prize}) = 0.04$$

$$p(\text{average}|\text{no PhD, prize}) = 0.04$$

$$p(\text{high}|\text{do PhD, no prize}) = 0.4$$

$$p(\text{high}|\text{no PhD, no prize}) = 0.2$$

$$p(\text{high}|\text{do PhD, prize}) = 0.95$$

$$p(\text{high}|\text{no PhD, prize}) = 0.95$$

The utilities are


$$U_C(\text{do PhD}) = -50000, \quad U_C(\text{no PhD}) = 0,$$

$$U_B(\text{low}) = 100000, \quad U_B(\text{average}) = 200000, \quad U_B(\text{high}) = 500000$$



Example

Should I do a PhD?



- Note: Strictly there should only be one utility node, as in general utilities don't additively decompose

The expected utility of Education is

$$U(E) = \sum_{I,P} p(I|E,P)p(P|E) [U_C(E) + U_B(I)]$$


so that $U(\text{do phd}) = 260174.000$, whilst not taking a PhD is

$U(\text{no phd}) = 240000.0244$, making it on average beneficial to do a PhD.




Markov Decision Processes

- Processes Through Time with actions and payoffs



Markov Decision Processes



We want to make that decision d_1 that will lead to maximal expected total utility

$$U(d_1|x_1) \equiv \sum_{x_2} \max_{d_2} \sum_{x_3} \max_{d_3} \sum_{x_4} \dots \max_{d_{T-1}} \sum_{x_T} p(x_{2:T}|x_1, d_{1:T-1})U(x_{1:T})$$

$$U(x_{1:T}) = u(x_2) + u(x_3) + \dots + u(x_T)$$

Our task is to compute $U(d_1|x_1)$ for each state of d_1 and then choose that state with maximal expected total utility. To carry out the summations and maximisations efficiently, we can use a simple message passing approach.



Stop Point

- Questions?

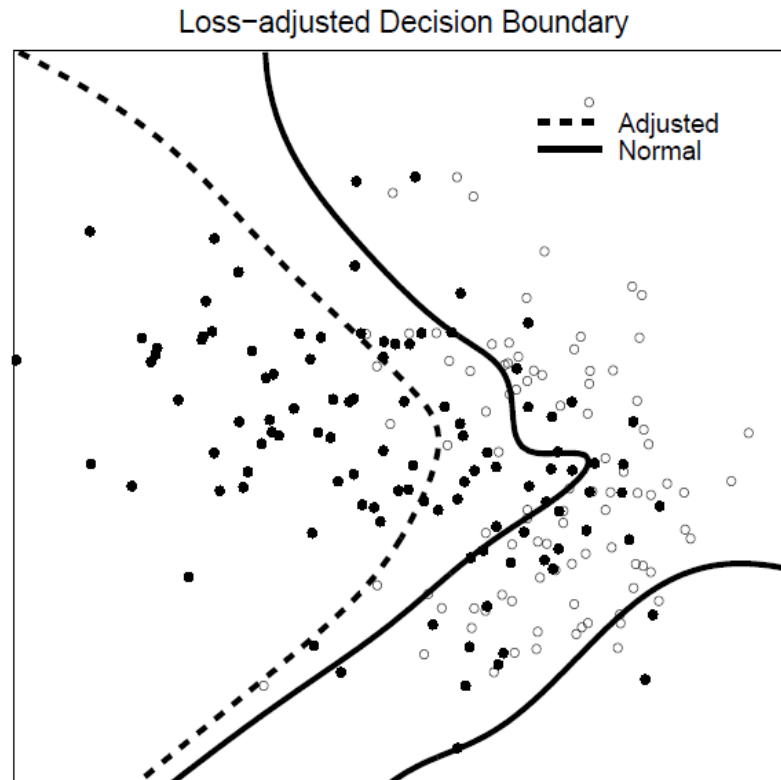


Decision in Machine Learning

- Standard Decisions in Machine Learning:
 - Classification
 - Choosing a path
 - Choosing a model
 - Labelling
 - Defining thresholds
 - Accepting a scientific theory
- These should all take account of the utility
- Often refer to loss (negative utility).

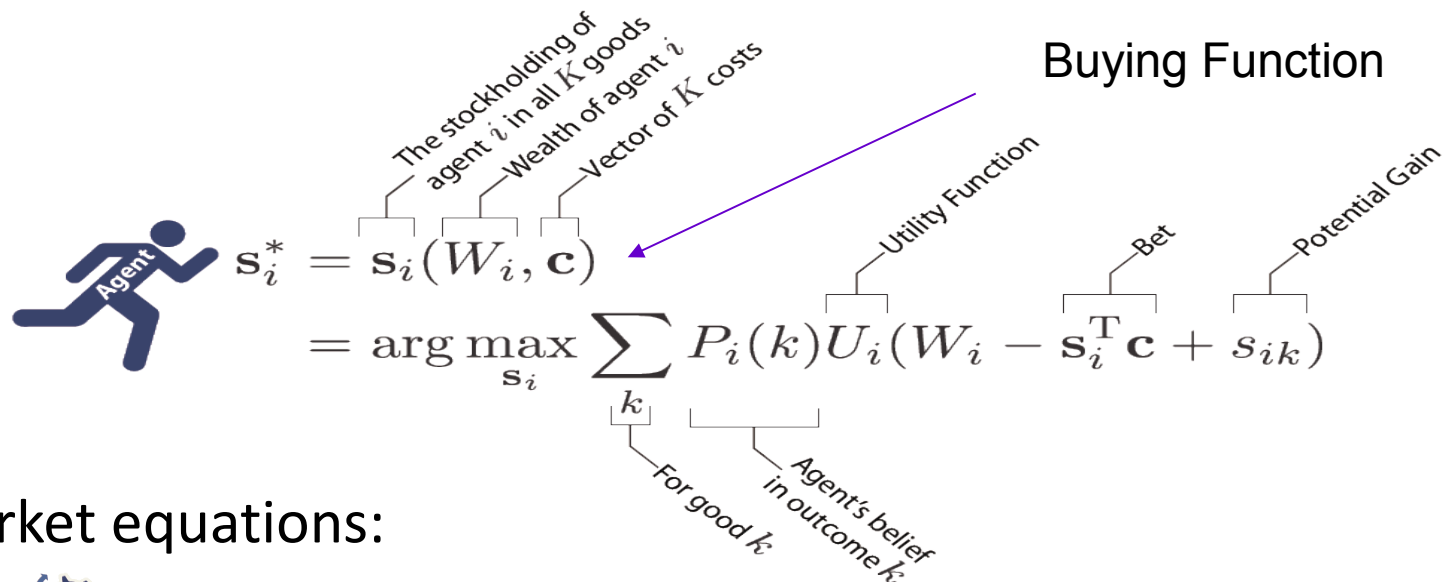


■ E.g.



Games

- Game Theory formulations involve expected utility:
 - Mixed strategies, and outcomes.
- Markets:
 - Market Trading and Expected Utility
 - E.g. Prediction Markets



Market equations:



$$\sum_{i=1}^{N_A} s_{ik}^* = \sum_{i=1}^{N_A} W_i \quad \forall k$$

N_A is number of agents
Amount won equals amount invested



Conclusions

- Decisions are about communication.
- They say how we decide what to say or do, and what not to.
- So Decision Making (and Decision Theory) is a necessary endpoint of all useful inference.
- Note the Bayesian view is that it is an endpoint: get the probabilities and then make the decision in the end.
- In reality any time you do optimization you are making a decision. Include the loss function.

