Preferences, utility and decision making

Christos Dimitrakakis

April 29, 2014

Goals of this lecture

Utility

- Understand the concept of preferences.
- See how utility can be used to formalize preferences.
- Show how we can combine utility and probability to deal with decision making under uncertainty.

The decision-theoretic foundations of artificial intelligence.

- Probability: how likely things are?
- Utility: which things do we want?

Interpretations of probability

- Objective: inherent randomness.
- Frequentist: long-term averages.
- Algorithmic: program complexity.
- Subjective: uncertainty.

Interpretations of utility

- Monetary.
- Psychological.
- “true” value of things?

Rewards

- We are going to receive a reward r from a set $\mathcal{R}$ of possible rewards.
- We prefer some rewards to others.

Example 1 (Possible sets of rewards $\mathcal{R}$)

- $\mathcal{R}$ is a set of tickets to different musical events.
- $\mathcal{R}$ is a set of financial commodities.

Preferences

Example 2 (Musical event tickets)

- Case 1: $\mathcal{R}$ are tickets to different music events at the same time, at equally good halls with equally good seats and the same price. Here preferences simply coincide with the preferences for a certain type of music or an artist.
- Case 2: $\mathcal{R}$ are tickets to different events at different times, at different quality halls with different quality seats and different prices. Here, preferences may depend on all the factors.

Example 3 (Route selection)

- $\mathcal{R}$ contains two routes, one short and one long, of the same quality.
- $\mathcal{R}$ contains two routes, one short and one long, but the long route is more scenic.

Preferences among rewards

Preferences

Let $a, b \in R$.

- Do you prefer a to b? Write $a \succ^* b$.
- Do you like a less than b? Write $a \prec^* b$.
- Do you like a as much as b? Write $a \approx^* b$.

We also use $\succsim^*$ and $\precsim^*$ for I like at least as much as and for I don't like any more than

Properties of the preference relations.

(i) For any $a, b \in R$, one of the following holds: $a \succ^* b$, $a \prec^* b$, $a \approx^* b$.

(ii) If $a, b, c \in R$ are such that $a \precsim^* b$ and $b \precsim^* c$, then $a \precsim^* c$.

Is transitivity a reasonable assumption?

Is transitivity a reasonable assumption?

Consider $r = (a, b)$, such that:

- $r \succ^* r'$ if $a > a'$ and $|b - b'| < \epsilon$
- $r \succ^* r'$ if $b >> b'$.

When we cannot select rewards directly

- In most problems, we cannot just choose which reward to receive.

When we cannot select rewards directly

- In most problems, we cannot just choose which reward to receive.
- We can only specify a distribution on rewards from a limited number of choices.

When we cannot select rewards directly

- In most problems, we cannot just choose which reward to receive.
- We can only specify a distribution on rewards from a limited number of choices.

Example 4 (Route selection)

- Each reward $r \in R$ is the time it takes to travel from A to B.
- We prefer shorter times.
- There are two routes, P_1, P_2.
- Route P_1 takes 10 minutes when the road is clear, but 30 minutes when the traffic is heavy. The probability of heavy traffic on P_1 is q_1.
- Route P_2 takes 15 minutes when the road is clear, but 25 minutes when the traffic is heavy. The probability of heavy traffic on P_2 is q_2.

Exercise 1

Say $q_1 = q_2 = 0.5$. Which route would you prefer?

Preferences among probability distributions

Preferences

Let P_1, P_2 be two distributions on $(R, \mathcal{F}_R)$.

- Do prefer P_1 to P_2? Write $P_1 \succ^* P_2$.
- Do you like P_1 less than P_2? Write $P_1 \prec^* P_2$.
- Do you like P_1 as much as P_2? Write $P_1 \eqsim^* P_2$.

We also use $\succsim^*$ and $\precsim^*$ in the usual sense.

Utility

In order to assign preferences to probability distributions, we use the concept of utility.

Utility

Definition 5 (Utility)

The utility is a function $U : R \to \mathbb{R}$, such that for all $a, b \in R$

$$a \succsim^* b \quad \text{iff} \quad U(a) \geq U(b), \tag{2.1}$$

Utility

Definition 5 (Utility)

The utility is a function $U : R \to \mathbb{R}$, such that for all $a, b \in R$

$$a \succsim^* b \quad \text{iff} \quad U(a) \geq U(b), \tag{2.1}$$

Definition 6 (Expected utility)

The expected utility of a distribution P on $\mathcal{R}$ is:

$$\mathbb{E}_P(U) = \sum_{r \in R} U(r)P(r) \tag{2.2}$$

Utility

Definition 5 (Utility)

The utility is a function $U : R \to \mathbb{R}$, such that for all $a, b \in R$

$$a \succsim^* b \quad \text{iff} \quad U(a) \geq U(b), \tag{2.1}$$

Definition 6 (Expected utility)

The expected utility of a distribution P on $\mathcal{R}$ is:

$$\mathbb{E}_P(U) = \sum_{r \in R} U(r) P(r) \tag{2.2}$$

Assumption 1 (The expected utility hypothesis)

The utility of P is equal to the expected utility of the reward under P. Consequently,

$$P \succsim^* Q \quad \textit{iff} \quad \mathbb{E}_P(U) \geq \mathbb{E}_Q(U). \tag{2.3}$$

Utility

Definition 5 (Utility)

The utility is a function $U : R \to \mathbb{R}$, such that for all $a, b \in R$

$$a \succsim^* b \quad \text{iff} \quad U(a) \geq U(b), \tag{2.1}$$

Definition 6 (Expected utility)

The expected utility of a distribution P on $\mathcal{R}$ is:

$$\mathbb{E}_P(U) = \int_R U(r) \, \mathrm{d}P(r) \tag{2.2}$$

Assumption 1 (The expected utility hypothesis)

The utility of P is equal to the expected utility of the reward under P. Consequently,

$$P \succsim^* Q \quad \textit{iff} \quad \mathbb{E}_P(U) \geq \mathbb{E}_Q(U). \tag{2.3}$$

Example 7

r	U(r)	P	Q
did not enter	0	1	0
paid 1 CU and lost	-1	0	0.99
paid 1 CU and won 10	9	0	0.01

Table : A simple gambling problem

	P	Q
$\mathbb{E}(U \mid \cdot)$	0	-0.9

Table : Expected utility for the gambling problem

Monetary rewards

Example 8

Choose between the following two gambles:

- A The reward is 500,000 with certainty.
- B The reward is 2,500,000 with probability 0.10. It is 500,000 with probability 0.89, and 0 with probability 0.01.

Monetary rewards

Example 8

Choose between the following two gambles:

- A The reward is 500,000 with probability 0.11, or 0 with probability 0.89.
- B The reward is: 2,500,000 with probability 0.1, or 0 with probability 0.9.

Monetary rewards

Example 8

Choose between the following two gambles:

- A The reward is 500,000 with certainty.
- B The reward is 2,500,000 with probability 0.10. It is 500,000 with probability 0.89, and 0 with probability 0.01.

Example 9

Choose between the following two gambles:

- A The reward is 500,000 with probability 0.11, or 0 with probability 0.89.
- B The reward is: 2,500,000 with probability 0.1, or 0 with probability 0.9.

Exercise 2 (Is the following statement true or false?)

For any finite U, if gamble A is preferred in the first example, gamble A must also be preferred in the second example.

The St. Petersburg Paradox

A simple game [Bernoulli, 1713]

- A fair coin is tossed until a head is obtained.
- If the first head is obtained on the n-th toss, our reward will be 2^n currency units.

The St. Petersburg Paradox

A simple game [Bernoulli, 1713]

- A fair coin is tossed until a head is obtained.
- If the first head is obtained on the n-th toss, our reward will be 2^n currency units.

How much are you willing to pay, to play this game once?

A 0

B 1-2

C Between 2 and 10?

D Between 10 and 1000?

E More than 1000?

The St. Petersburg Paradox

A simple game [Bernoulli, 1713]

- A fair coin is tossed until a head is obtained.
- If the first head is obtained on the n-th toss, our reward will be 2^n currency units.

- The probability to stop at round n is 2^{-n}.

The St. Petersburg Paradox

A simple game [Bernoulli, 1713]

- A fair coin is tossed until a head is obtained.
- If the first head is obtained on the n-th toss, our reward will be 2^n currency units.

- The probability to stop at round n is 2^{-n}.
- Thus, the expected monetary gain of the game is

$$\sum_{n=1}^{\infty} 2^n 2^{-n} = \infty.$$

The St. Petersburg Paradox

A simple game [Bernoulli, 1713]

- A fair coin is tossed until a head is obtained.
- If the first head is obtained on the n-th toss, our reward will be 2^n currency units.

- The probability to stop at round n is 2^{-n}.
- Thus, the expected monetary gain of the game is

$$\sum_{n=1}^{\infty} 2^n 2^{-n} = \infty.$$

- If your utility function were linear you'd be willing to pay any amount to play.

Concave versus convex functions

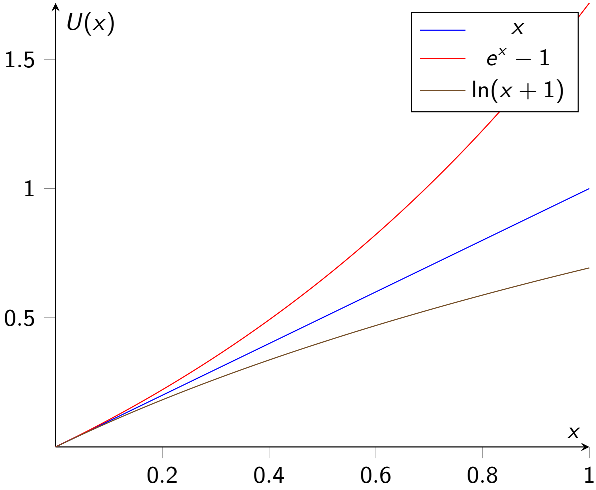

Convex functions

Definition 10

A function g is convex on A if, for any points $x, y \in A$, and any $\alpha \in [0, 1]$:

$$\alpha g(x) + (1 - \alpha)g(y) \geq g[\alpha x + (1 - \alpha)y]$$

Theorem 11 (Jensen's inequality)

If g is convex on $\mathcal{S}$ and $x \in \mathcal{S}$ with measure $P(A) = 1$ and $\mathbb{E}(x)$ and $\mathbb{E}[g(x)]$ exist, then:

$$\mathbb{E}[g(x)] \geq g[\mathbb{E}(x)]. \tag{2.4}$$

Example 12

If the utility function is convex, then we choose a gamble giving a random gain x rather than one giving a fixed gain $\mathbb{E}(x)$. Thus, a convex utility function implies risk-taking. An example function is

$$f(x) = e^x.$$

Concave functions

Definition 13

A function g is concave on $\mathcal{S}$ if, for any points $x, y \in \mathcal{S}$, and any $\alpha \in [0, 1]$:

$$\alpha g(x) + (1 - \alpha)g(y) \leq g[\alpha x + (1 - \alpha)y]$$

Example 14

If the utility function is concave, then we choose a gamble giving a fixed gain $\mathbb{E}[X]$ rather than one giving a random gain X. Consequently, a concave utility function implies risk aversion. An example concave function is

$$f(x) = \ln x.$$

St. Petersburg paradox - continued

Exercise 3

We have establised that if your utility for money us $U(x) = x$, and the coin was fair, then you would be willing to pay any amount to play the game.

- *Assume that the coin is fair, but your utility for money is $U(x) = \ln x$. How much would you now be willing to pay to play the game? Hint: Calculate the expected utility of playing.*
- *Assume that the coin is not fair, but only comes head with probability 0.4, and your utility is linear: $U(x) = x$. How much would you be willing to play now?*

Example 15 (Insurance)

The act of buying insurance can be related to concavity of our utility function. Let d be the insurance cost, h our insurance cover and ϵ the probability of needing the cover.

Exercise 4

Assume that $U(x) = \ln(C + x)$. C can be seen as the amount of credit that we can sustain before becoming ruined.

- *If $\epsilon > 0$, $h > C$, how high a premium d are we willing to pay?*
- *What if $h = (1 - p)C$, with $p \in (0, 1)$?*

Example 15 (Insurance)

The act of buying insurance can be related to concavity of our utility function. Let d be the insurance cost, h our insurance cover and ϵ the probability of needing the cover. Then we are going to buy insurance if the utility of losing d with certainty is greater than the utility of losing $-h$ with probability ϵ.

Exercise 4

Assume that $U(x) = \ln(C + x)$. C can be seen as the amount of credit that we can sustain before becoming ruined.

- *If $\epsilon > 0$, $h > C$, how high a premium d are we willing to pay?*
- *What if $h = (1 - p)C$, with $p \in (0, 1)$?*

Example 15 (Insurance)

The act of buying insurance can be related to concavity of our utility function. Let d be the insurance cost, h our insurance cover and ϵ the probability of needing the cover. Then we are going to buy insurance if the utility of losing d with certainty is greater than the utility of losing $-h$ with probability ϵ.

$$U(-d) > \epsilon U(-h) + (1-\epsilon)U(0). \tag{2.5}$$

Exercise 4

Assume that $U(x) = \ln(C + x)$. C can be seen as the amount of credit that we can sustain before becoming ruined.

- *If $\epsilon > 0$, $h > C$, how high a premium d are we willing to pay?*
- *What if $h = (1-p)C$, with $p \in (0,1)$?*

Example 15 (Insurance)

The act of buying insurance can be related to concavity of our utility function. Let d be the insurance cost, h our insurance cover and ϵ the probability of needing the cover. Then we are going to buy insurance if the utility of losing d with certainty is greater than the utility of losing $-h$ with probability ϵ.

$$U(-d) > \epsilon U(-h) + (1-\epsilon)U(0). \tag{2.5}$$

The company has a linear utility, and fixes the premium d high enough for

$$d > \epsilon h. \tag{2.6}$$

Exercise 4

Assume that $U(x) = \ln(C + x)$. C can be seen as the amount of credit that we can sustain before becoming ruined.

- *If $\epsilon > 0$, $h > C$, how high a premium d are we willing to pay?*
- *What if $h = (1-p)C$, with $p \in (0, 1)$?*

Summary

- We can subjectively indicate which events we think are more likely.
- Using relative likelihoods, we can define a subjective probability P for all events.
- Similarly, we can subjectively indicate preferences for rewards.
- We can determine a utility function for all rewards.
- Hypothesis: we prefer the probability distribution (over rewards) with the highest expected utility.
- Concave utility functions imply risk aversion (and convex, risk-taking).

[1] Morris H. DeGroot. *Optimal Statistical Decisions*. John Wiley & Sons, 1970.

[2] Milton Friedman and Leonard J. Savage. The expected-utility hypothesis and the measurability of utility. *The Journal of Political Economy*, 60(6):463, 1952.

[3] Leonard J. Savage. *The Foundations of Statistics*. Dover Publications, 1972.