Decision Methods and Models

Master's Degree in Computer Science

Roberto Cordone

DI - Università degli Studi di Milano

Schedule:	Thursday 16.30 - 18.30 in Aula Magna (CS department)
	Friday 12.30 - 14.30 in classroom 301
Office hours:	on appointment
E-mail:	roberto.cordone@unimi.it
Web page:	https://homes.di.unimi.it/cordone/courses/2024-mmd/2024-mmd.html
Ariel site:	https://myariel.unimi.it/course/view.php?id=4467

Decision-making under risk

We assume

- a preference relation Π that is a weak order
  with a known consistent value function $u(f)$ (replaced by a cost f)
- an uncertain environment: $|\Omega| > 1$ with probabilistic information
- a single decision-maker: $|D| = 1 \Rightarrow \Pi_d$ reduces to Π

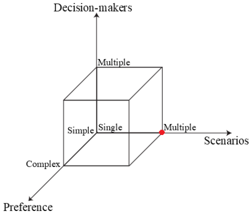

We aim to overcome the limits of the expected value criterium

- introducing a number of desired properties (axioms)
- building a choice criterium that by construction satisfies them

After normalisation, there will be exactly one

Stochastic utility theory

We have seen that the expected value criterium has problems, related to

- decision-makers having strict "preferences" between alternatives of equal expected value
- decision-makers having different "preferences" between the same pairs of alternatives
- paradoxes related to small probabilities and large impacts

Strictly speaking, preferences refer to impacts f and f',
but we refer here to a practical choice between alternatives x and x',
based on the impact vectors $f(x, \cdot)$ and $f(x', \cdot)$

The Austrian-Hungarian mathematicians Janos Von Neumann and Oskar Morgenstern proposed a constructive way out of these problems

Lotteries

A finite simple lottery is a pair $\ell\,(f,\pi) = (f(\omega), \pi(\omega))$ where

- $f(\omega) : \Omega \to \mathbb{R}$ is a random variable
- $\pi(\omega) : \Omega \to [0,1]$ is a probability function
- $\Omega = \left\{\omega^{(1)}, \ldots, \omega^{(r)}\right\}$ is a finite sample set

The extension to infinite sets is possible, but not considered here

Consequently, the set of all finite simple lotteries on Ω will be

$$L_{F,\Omega} = F^{|\Omega|} \times \mathcal{P}\,(\Omega)$$

In a decision problem in conditions of risk,
each alternative $x \in X$ corresponds to a lottery $\ell(x) \in L_X \subset L_{F,\Omega}$

$$x \leftarrow \ell(x) = (f(x,\omega_1), \pi(\omega_1)) \oplus \ldots \oplus (f(x,\omega_r), \pi(\omega_r))$$

where we skip the terms of zero probability for the sake of brevity

This notation is nonstandard

Special lotteries

Some particular cases compact the notation even further

- degenerate lottery, where a single scenario has probability 1

$$\ell_f = (f, 1)$$

- binary lottery, where two scenarios have positive probability

$$\ell_{f,\alpha,f'} = (f, \alpha) \oplus (f', 1-\alpha)$$

Compound lotteries

The simple lottery set $L_{F,\Omega}$ extends by recursion to compound lotteries, where the impact obtained in each scenario is allowed to be a lottery

The notation abuse is to avoid multiplying symbols

Compound lotteries model

- lotteries taking place in subsequent phases
- decisions taken before a sequence of uncertain events

They admit a graphical tree representation with

- uncertain events on the internal nodes
- deterministic impacts on the leaves
- conditional probabilities on the arcs
- probabilities summing to 1 on the arcs going out of each node

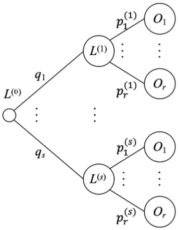

Preference relations on lotteries

A preference between lotteries is a binary relation Π on the lottery set

$$\Pi \subset 2^{L_{F,\Omega} \times L_{F,\Omega}}$$

A stochastic utility function $u : L_{F,\Omega} \to \mathbb{R}$ is consistent with Π when

$$\ell \preceq \ell' \Leftrightarrow u(\ell) \geq u(\ell') \qquad \text{for all } \ell, \ell' \in L$$

Von Neumann and Morgenstern

1. assume the existence of a preference relation Π on lotteries
2. impose suitable conditions on Π
3. build a stochastic utility function $u(\ell)$ that is consistent with Π
4. reduce the decision problem in conditions of risk to

$$\max u(\ell(x))$$
$$x \in X \Leftrightarrow \ell(x) \in L_X$$

where feasible region X corresponds to the feasible lotteries $L_X \subset L_{F,\Omega}$

Stochastic utility axioms

The preference relation Π on $L_{F,\Omega}$ enjoys the following properties

1. weak ordering: it is reflexive, complete and transitive
2. monotony: lotteries giving larger probabilities to better outcomes are preferable:

$$\alpha \geq \beta \Leftrightarrow (\ell, \alpha) \oplus (\ell', 1-\alpha) \preceq (\ell, \beta) \oplus (\ell', 1-\beta) \text{ for all } \ell \preceq \ell'$$

3. continuity: any intermediate impact between two lotteries admits an equivalent compound lottery with the two given ones as outcomes:

$$\ell \preceq f \preceq \ell' \Rightarrow \exists \alpha \in [0;1] : f \sim (\ell, \alpha) \oplus (\ell', 1-\alpha)$$

Changing the probabilities modifies the preference continuously

No impact remains uncovered

Stochastic utility axioms

4. independence (or substitution): the preference between two lotteries does not change combining them with the same lottery with the same probability:

$$\ell \preceq \ell' \Leftrightarrow (\ell, \alpha) \oplus (\ell'', 1-\alpha) \preceq (\ell', \alpha) \oplus (\ell'', 1-\alpha) \text{ for all } \alpha \in (0; 1]$$

5. reduction: a compound lottery is indifferent to a simple lottery with
   - the same final impacts
   - probabilities given by the laws of
     - conditional probabilities: multiply along each path
     - total probabilities: sum on disjoint paths

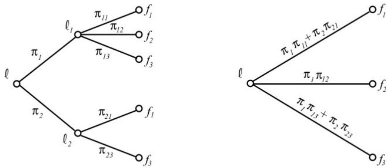

Impacts and probabilities are relevant, the lottery structure is not

The Von Neumann-Morgenstern theorem

The axioms look overall reasonable, even if some have been criticised

In particular, the independence axiom is often violated in practice

Theorem

Let F and Ω be finite sets, and Π a preference relation on lotteries. Let Π be not fully indifferent on F and satisfy the five axioms. Then, there exists a single function $u(\ell)$ that is consistent with Π, equal to 0 in the worst impacts and to 1 in the best

If all impacts are indifferent, just set $u(\ell)$ uniform

In the other cases, the proof is constructive

1. set the utility of extreme impacts: since F is finite and Π is a weak order, there is a best impact f° and and a worst impact $f^\dagger$, with

$$f^\circ \prec f^\dagger$$

Otherwise, all impacts are indifferent

Now, set $u(f^\circ, 1) = 1$ and $u(f^\dagger, 1) = 0$

Example

Consider the following decision problem in conditions of risk

$f(x,\omega)$	$\omega^{(1)}$	$\omega^{(2)}$	$\omega^{(3)}$
$x^{(1)}$	10	20	50
$x^{(2)}$	50	10	30

$\pi(\omega)$	0.25	0.50	0.25

Suppose that

- the impact set is $F = \{10, 20, 30, 50\}$ (finite)
- the impacts are benefits (no full indifference)

The theorem's assumptions are satisfied

The first step amounts to

- finding the worst and best impact: $f^{\dagger} = 10$ and $f^{\circ} = 50$
- imposing the normalised extreme values of utility:

$$u(10) = 0 \qquad u(50) = 1$$

The Von Neumann-Morgenstern theorem

2. set the utility of all degenerate lotteries: since f is intermediate between $f^\dagger$ and f°, by continuity

$$f^\circ \preceq f \preceq f^\dagger \Rightarrow \exists \alpha \in [0,1] : f \sim \ell_{f^\circ,\alpha,f^\dagger}$$

The value α must be indicated by the decision-maker

There is a unique α_f for each f: by contradiction, assume two

$$f \sim \ell_{f^\circ,\alpha,f^\dagger} \sim \ell_{f^\circ,\beta,f^\dagger}$$

By monotony

$$\begin{cases} \ell_{f^\circ,\alpha,f^\dagger} \preceq \ell_{f^\circ,\beta,f^\dagger} \text{ and } f^\dagger \preceq f^\circ \Rightarrow \alpha \geq \beta \\ \ell_{f^\circ,\beta,f^\dagger} \preceq \ell_{f^\circ,\alpha,f^\dagger} \text{ and } f^\dagger \preceq f^\circ \Rightarrow \beta \geq \alpha \end{cases} \Rightarrow \beta = \alpha$$

Now, set the utility $u(f,1) = \alpha_f$

Example

Consider the following decision problem in conditions of risk

$f(x,\omega)$	$\omega^{(1)}$	$\omega^{(2)}$	$\omega^{(3)}$
$x^{(1)}$	10	20	50
$x^{(2)}$	50	10	30

$\pi(\omega)$	0.25	0.50	0.25

The second step amounts to asking the decision-maker for each impact f

- finding a binary lottery $\ell_{f^\circ,\alpha_f,f^\dagger}$ equivalent to f
  - $20 \sim (10, \alpha_{20}) \oplus (50, 1-\alpha_{20})$
  - $30 \sim (10, \alpha_{30}) \oplus (50, 1-\alpha_{30})$
- imposing the lottery probability as the utility of the impact
  - $u(20) = \alpha_{20} = 0.40$
  - $u(30) = \alpha_{30} = 0.60$

The Von Neumann-Morgenstern theorem

3. set the utility of general lotteries:
   To keep it simple, consider a binary one: $\ell = (f^{(1)}, \pi_1) \oplus (f^{(2)}, \pi_2)$

   Each final impact is equivalent to a binary lottery

   $$f^{(r)} \sim (f^{\dagger}, 1 - \alpha_{f^{(r)}}) \oplus (f^{\circ}, \alpha_{f^{(r)}}) \qquad (\text{in short } \ell^{(r)})$$

   By substitution, we obtain from $f^{(1)}$ a compound lottery

   $$f^{(1)} \sim \ell^{(1)} \Rightarrow (f^{(1)}, \pi_1) \oplus (f^{(2)}, \pi_2) \sim (\ell^{(1)}, \pi_1) \oplus (f^{(2)}, \pi_2)$$

   Again by substitution we replace also $f^{(2)}$

   $$f^{(2)} \sim \ell^{(2)} \Rightarrow (f^{(2)}, \pi_2) \oplus (\ell^{(1)}, \pi_1) \sim (\ell^{(2)}, \pi_2) \oplus (\ell^{(1)}, \pi_1)$$

   By transitivity, the original lottery is equivalent to the final one

   $$\ell = (f^{(1)}, \pi_1) \oplus (f^{(2)}, \pi_2) \sim (\ell^{(1)}, \pi_1) \oplus (\ell^{(2)}, \pi_2)$$

   The order in the combinations is irrelevant

   The result is a compound two-level lottery that returns only $f^{\dagger}$ and f°

   $$\ell \sim ((f^{\dagger}, 1 - \alpha_{f^{(2)}}) \oplus (f^{\circ}, \alpha_{f^{(2)}}), \pi_2) \oplus ((f^{\dagger}, 1 - \alpha_{f^{(1)}}) \oplus (f^{\circ}, \alpha_{f^{(1)}}), \pi_1)$$

The Von Neumann-Morgenstern theorem

By reduction the overall lottery can be replaced by a simple lottery

$$\ell \sim (f^\dagger, \pi_1(1-\alpha_{f(1)}) + \pi_2(1-\alpha_{f(2)})) \oplus (f^\circ, \pi_1\alpha_{f(1)} + \pi_2\alpha_{f(2)})$$

Now, set the utility of ℓ to the probability of f° in the final lottery:

$$u(\ell(f,\pi)) = \sum_{\omega \in \Omega} \pi(\omega)\; u(f(\omega))$$

since

- every scenario ω corresponds to a path reaching f°
- the paths consist of two arcs
- $\pi(\omega)$ is the probability on the first arc of the path
- $\alpha_{f(\omega)} = u(f(\omega))$ is the probability on the second arc of the path

The Von Neumann-Morgenstern theorem

The utility of a lottery is the expected value of the utility in the scenarios

$$u(\ell) = E[u(f)]$$

As Bernouilli proposed, combine perceived utilities, not impacts!

In summary, the Von Neumann-Morgenstern theorem

- receives from the decision-maker $u(f)$ for all $f \in F$
- returns $u(\ell)$ for all $\ell \in L$

Since in a decision problem an alternative is a lottery

$$\max_{x \in X} E[u(f(x)]$$

Example

Since we know the utilities of the deterministic impacts

f	10	20	30	50
$u(f)$	0	0.40	0.60	1

we can determine the utility of any lottery, such as

$f(x,\omega)$	$\omega^{(1)}$	$\omega^{(2)}$	$\omega^{(3)}$
$x^{(1)}$	10	20	50
$x^{(2)}$	50	10	30
π	0.25	0.50	0.25

The utilities are

- $u\left(x^{(1)}\right) = 0.25 \cdot 10 + 0.50 \cdot 20 + 0.25 \cdot 50 = 0.45$
- $u\left(x^{(2)}\right) = 0.25 \cdot 50 + 0.50 \cdot 10 + 0.25 \cdot 30 = 0.40$

and the former is better than the latter

Stochastic utility and multi-attribute utility

Stochastic utility and multi-attribute utility exhibit strong similarities

- they are convex combinations of utilities with coefficients

$$\tilde{u}(f) = \sum_{l \in P} w_l \tilde{u}_l(f_l) \qquad u(\ell(f,\pi)) = \sum_{\omega \in \Omega} \pi(\omega)\, \alpha_{f(\omega)}$$

- the utilities are normalised between a worst and a best impact
- the coefficients are normalised between 0 and 1 with unitary sum

but there are also strong differences

- additivity is intrinsically satisfied (no special condition required), because the different scenarios do not interact, while indicators do
- the attribute set P is replaced by scenario set Ω, possibly infinite
- the weights w_l are replaced by probabilities $\pi(\omega)$ that do not necessarily require the decision-maker (they might be frequencies)

Attitudes towards risk

Risk profile is the shape of $u(f)$ for all $f \in F$:

- it determines $u(\ell(f,\pi))$ for all $\ell \in L$, combining $u(f)$ and $\pi(\omega)$
- therefore, it shows the attitude of the decision-maker towards risk

For the sake of simplicity, we will assume that

- f is a benefit
- the impact set is an interval $F = \left[f^{\dagger}, f^{\circ}\right]$
- consequently, the risk profile increases from $\left(f^{\dagger}, 0\right)$ to $(f^{\circ}, 0)$

The extension to more general cases is possible

Now consider

- an intermediate deterministic impact $f_\alpha = (1-\alpha) f^\dagger + \alpha f^\circ$
- a binary lottery $\ell_{f^\circ,\alpha,f^\dagger}$ combining f° and $f^\dagger$ with coefficient α

They are equivalent for the expected value criterium

$$\phi_{EV}\left(f_\alpha\right) = f_\alpha = \phi_{EV}\left(\ell_{f^\circ,\alpha,f^\dagger}\right)$$

What about their stochastic utilities $u(f_\alpha)$ and $u(\ell_{f^\circ,\alpha,f^\dagger})$?

The utility of any binary lottery $\ell_{f^{(1)},\alpha,f^{(2)}}$ can be computed from the segment between points $\left(f^{(1)}, u(f^{(1)})\right)$ and $\left(f^{(2)}, u(f^{(2)})\right)$

$$u\left(\ell_{f^\circ,\alpha,f^\dagger}\right) = \alpha u(f^\circ) + (1-\alpha) u(f^\dagger)$$

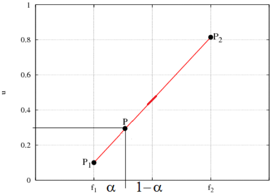

Attitudes towards risk

Three relevant cases (not exhaustive) exist:

1. convex case: the risk profile is above the segment from $(f^\dagger, 0))$ to $(f^\circ, 1)$ and the lottery is preferred to the deterministic impact

$$u(f_\alpha) \leq u\left(\ell_{f^\circ,\alpha,f^\dagger}\right) \quad \text{for all } \alpha \in [0,1]$$

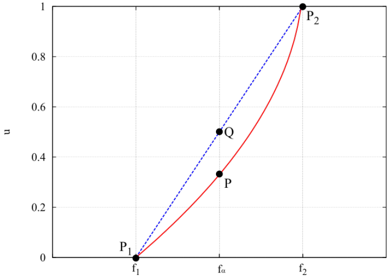

The decision-maker is risk-prone

Attitudes towards risk

Three relevant cases (not exhaustive) exist:

2. linear case: the risk profile is lying on the segment from $(f^\dagger, 0))$ to $(f^\circ, 1)$ and the lottery and the deterministic impact are equivalent

$$u(f_\alpha) = u\left(\ell_{f^\circ,\alpha,f^\dagger}\right) \quad \text{for all } \alpha \in [0,1]$$

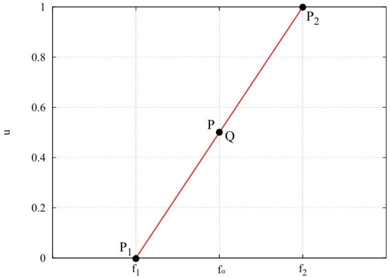

The decision-maker is risk-neutral

The expected value criterium is confirmed

Attitudes towards risk

Three relevant cases (not exhaustive) exist:

3. concave case: the risk profile is below the segment from $(f^\dagger, 0))$ to $(f^\circ, 1)$ and the deterministic impact is preferred to the lottery

$$u(f_\alpha) \geq u\left(\ell_{f^\circ,\alpha,f^\dagger}\right) \quad \text{for all } \alpha \in [0,1]$$

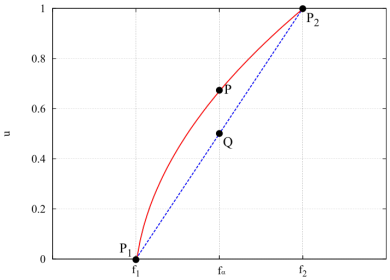

The decision-maker is risk-averse

Example

Suppose that the impact is a benefit and its set is $F = [0, 1\,000]$

The decision-maker has a concave risk profile

$$u(f) = \sqrt{\frac{f}{1\,000}}$$

and is, therefore, risk-averse

The feasible region consists of two alternatives:

- $x = (250, 1)$ *(deterministic impact)*
- $x' = (810, 0.1) \oplus (360, 0.5) \oplus (160, 0.4)$ *(lottery)*

The corresponding utilities are

- $u(x) = \sqrt{\frac{250}{1\,000}} = 0.5$
- $u(x') = 0.1 \cdot u(810) + 0.5 \cdot u(360) + 0.4 \cdot u(160) =$
  $0.1 \cdot 0.9 + 0.5 \cdot 0.6 + 0.4 \cdot 0.4 = 0.55$

The lottery is better

Certainty equivalent

Assume that the risk profile $u(f)$ is invertible (strictly increasing): the inverse function $f(u)$

- provides an equivalent information
- can be useful to present a situation in more intuitive terms

Taking part to the lottery is like obtaining a certain gain

Given lottery ℓ, its certainty equivalent $CE(\ell)$ is the deterministic impact that is equivalent to the lottery:

$$CE(\ell) = f(u(\ell)) \quad \Leftrightarrow \quad f(u(\ell)) \sim \ell$$

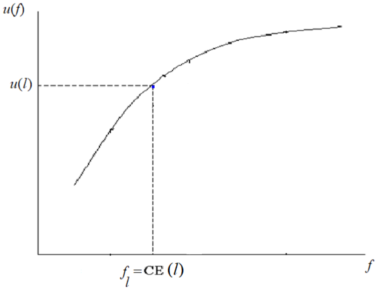

Example

The risk profile is invertible and its inverse is

$$u(f) = \sqrt{\frac{f}{1\,000}} \quad \Leftrightarrow \quad f(u) = 1\,000\,u^2$$

Its certainty equivalent is $CE(x') = f(0.55) = 1\,000 \cdot 0.55^2 = 302.5$

The decision-maker would consider the alternative x'
equivalent to a deterministic gain of 302.5, that is larger than $f(x)$

Risk premium

The certainty equivalent can differ from the expected value of the impact

Given a lottery ℓ of utility u

- the expected value of the impact is $\phi_{EV}(\ell) = E[f(\ell, \omega)]$
- the certainty equivalent is $CE(\ell) = f(u(\ell))$

The risk premium $RP(\ell)$ is the difference of the two terms

$$RP(\ell) = \phi_{EV}(\ell) - CE(\ell) = E[f(\ell, \omega)] - f(u(\ell))$$

It measures the variation on all impacts required to make the lottery equivalent to deterministically gaining its original expected value

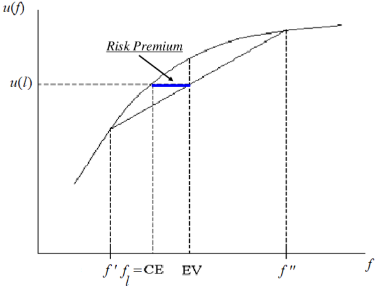

Example

As we have seen, lottery $x' = (810, 0.1) \oplus (360, 0.5) \oplus (160, 0.4)$
has a certainty equivalent $CE(x') = 302.5$

Its impact, however, has an expected value equal to

$$\phi_{EV}(x') = E[f(x', \omega)] = 0.1{\cdot}810 + 0.5{\cdot}360 + 0.4{\cdot}160 = 81 + 180 + 64 = 325$$

that is larger (confirming that the decision-maker is risk-averse)

In order to make the lottery x' equivalent to its expected impact,
we should modify all impacts by a fixed amount equal to the risk premium

$$RP(x') = \phi_{EV}(x') - CE(x') = 325 - 302.5 = 22.5$$

increasing them (risk aversion, once again) to

$$(832.5, 0.1) \oplus (382.5, 0.5) \oplus (182.5, 0.4)$$

These numbers aim to be more intuitive than abstract normalised utilities