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

Structured models of preference

We assume

- a preference relation Π with a consistent utility function $u(f)$
- a certain environment: $|\Omega| = 1 \Rightarrow f(x, \bar{\omega})$ reduces to $f(x)$
- a single decision-maker: $|D| = 1 \Rightarrow \Pi_d$ reduces to Π

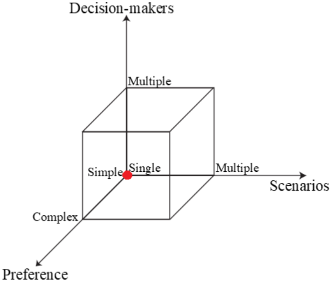

but we known the preference Π, not the utility function $u(f)$

How to build it?

From a preference relation to a consistent value function

Specific models of preference have specific applications

Case by case, they might work or not

We want a general way to derive $u(f)$ from Π

1. introduce a graphical tool (indifference map)
2. turn the graph into a function, with a complex error-prone process
3. define a special case with a simpler process (additive value functions)
4. characterise the preference relations falling within the special case

This could be useful also outside decision-making

- videogames: rank characters in Role-Play Games (RPG) starting from vectors of attributes (skill, strength, wizardry, ...)
- image processing: measure contrast in images, in order to increase it quantitatively with filters
- ...

(*any case in which a ranking should be turned into a measure*)

Indifference curves

Let the preference relation Π be a weak order

An indifference curve is a set $I \subseteq F$ of reciprocally indifferent impacts

Indifference curves always enjoy the following properties

- the curves cover F: every impact belongs to a curve
- any two indifference curves I and I' have an empty intersection
  (*transitivity would merge them*)
- the weak order on impacts maps onto a total order on curves
  (*indifference curves enjoy antisymmetry*)

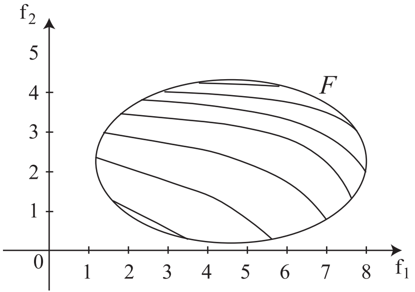

Indifference curves

Usually additional technical assumptions are made

- continuity implies that the curves are regular mathematical objects and not completely general sets of points
- the utility function $u(f)$ has a continuous infinity of values $c \in \mathbb{R}$ (*excluding lexicographic preference, complete indifference*)
- each indifference curve is expressed in implicit form

$$u(f) = c$$

- when the implicit form $u(f) = c$ can be turned into an explicit form

$$f_l = f_l(c, f_1, \ldots, f_{l-1}, f_{l+1}, \ldots, f_p)$$

an indifference curve is a hypersurface of $p - 1$ dimensions in $\mathbb{R}^p$

When $p = 2$, it is a line in the plane

Given a weak order preference relation Π on F, its indifference map $\mathcal{I}_\Pi$ is the ordered family of indifference curves covering F

The correspondence between Π and $\mathcal{I}_\Pi$ is one to one

- Π identifies all groups of indifferent impacts (curves) and their order
- $\mathcal{I}_\Pi$ identifies the preference between all pairs of impacts

As preference Π admits infinite utility functions $u(f)$, correspondingly the indifference map $\mathcal{I}_\Pi$ corresponds to infinite utility functions $u(f)$

Example: both $u(f) = f_1^2 f_2^3 = c$ and $u'(f) = 2\log_2 f_1 + 3\log_2 f_2 = c$ describe exactly the same preference relation and indifference map

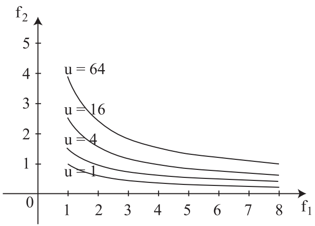

Determining a utility function (the long way)

Given a preference relation Π on F:

1. extract a sample $\tilde{F}$ from F (*dense enough, but not too much*)
2. ask the decision-maker to
   - sort the sampled impacts
   - identify their equivalence classes
3. draw an interpolating curve for each equivalence class
4. guess a parametric utility function family from their shape

$$u = u_\alpha(f) \quad \text{with} \quad \alpha = [\alpha_1 \dots \alpha_p]^T$$

5. each pair of indifferent impacts implies an equation on α

$$f \sim f' \Leftrightarrow u(f(\alpha)) = u(f'(\alpha))$$

6. add a normalisation condition to select one of the equivalent utilities
7. make consistency checks by comparing other pairs of impacts;
   if they fail, go back to 4 and change the parametric utility

Example

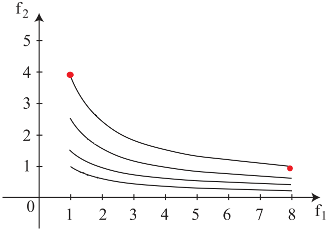

- they look like Cobb-Douglas curves that model basic human features

$$u_\alpha(f) = \prod_{l=1}^{p} f_l^{\alpha_l} = c$$

- find $p-1$ pairs of indifferent impacts (in this case, just one)

$$(8,1) \sim (1,4) \quad \Rightarrow \quad 8^{\alpha_1} 1^{\alpha_2} = 1^{\alpha_1} 4^{\alpha_2} \quad \Rightarrow \quad \alpha_2 = 3/2\alpha_1$$

- since $(u_\alpha(f))^\beta = u_{\beta \cdot \alpha}(f)$ is equivalent to $u_\alpha(f)$ for any $\beta > 0$, normalisation condition $\sum_l \alpha_l = 1$ removes a false freedom degree

  ($p-1$ *pairs of indifferent impacts are necessary and sufficient*)

Problems

The process is in general very complex and error-prone because

- large samples are costly (*time for sorting*)
- the sample must include at least $p - 1$ pairs of indifferent impacts found by trial and error
- small samples ($\approx p$) are imprecise (*false indifferent pairs*)
- numerical errors over many equations combine in cascade
- mutually dependent pairs are useless (*redundant equations on* α)
- high dimensional spaces make it hard to draw curves and guess u_α (*complicated case studies have hundreds of indicators*)

Indeed, economics textbooks usually assume $p = 2$

The process can be helped if additional assumptions hold that make it easier to find pairs of indifferent impacts

Useful properties

Some properties (not guaranteed!) help to draw curves and guess $u(f)$:

- invertibility: $u(f) = c$ can be solved with respect to each f_l
  - always verified when the indicators are costs or benefits
  - not verified for the thermostat: each humidity/temperature corresponds to two temperatures/humidities

  Example: $u(f) = f_1^2 f_2^3 = c \Rightarrow f_1 = \sqrt{c/f_2^3}$ or $f_2 = \sqrt[3]{c/f_1^2}$
- monotony: strictly decreasing or increasing indifference curves
  (*an increase in one is balanced by a decrease in another*)
  - always verified when the indicators are all costs or all benefits
- convexity or concavity ("law of diminishing marginal utility"):
  - example for benefits (convex curves):
    - high salary makes free time more precious,
    - high free time makes salary more precious
  - example for costs (concave curves):
    - low pollution makes taxes more annoying,
    - low taxes make pollution more annoying

Additivity (a game changer)

Additive utility function is a sum of terms depending on single indicators

$$u(f_1, \ldots, f_p) = \sum_{l=1}^{p} u_l(f_l)$$

It is a specific case, but brings many simplifications, as we can

- ask different decision-makers for each indicator f_l (*split the work*)
- ask decision-makers with experience in the sector (*more reliable*)
- compare scalar values f_l instead of vectors f (*easier and better*)
- build functions $u_l(f_l)$ with one argument (*easier and better*)

Since a utility function has infinitely many forms,
a nonadditive function can have an additive equivalent form

This is not guaranteed! How can we know?

Example (Cobb-Douglas functions):

$$u(f) = \prod_{l=1}^{p} f_l^{\alpha_l} \quad \text{is equivalent to} \quad u'(f) = \log u(f) = \sum_{l=1}^{p} \alpha_l f_l$$

Preferential indipendence

How to know that Π admits an additive consistent utility function?

Given the set of attribute indices $P = \{1, \dots, p\}$, focus on subset $L \subset P$

$$f = \begin{bmatrix} f_L \\ f_{P \setminus L} \end{bmatrix} \qquad \text{Example: } \begin{bmatrix} \text{environmental attributes} \\ \text{other attributes} \end{bmatrix}$$

L is preferentially indipendent from $L \subset P$ when

$$\begin{bmatrix} f_L \\ \phi \end{bmatrix} \preceq \begin{bmatrix} f'_L \\ \phi \end{bmatrix} \Leftrightarrow \begin{bmatrix} f_L \\ \psi \end{bmatrix} \preceq \begin{bmatrix} f'_L \\ \psi \end{bmatrix}$$

for all subvectors ϕ, ψ, f_L, f'_L such that the four impacts are in F

Preferences between values in L do not depend on the values out of L

Let us see several (positive and negative) examples

Examples

A cost ($L = \{\bar{I}\}$) is preferentially independent from all other indicators

$$f_{\bar{I}} \leq f'_{\bar{I}} \Leftrightarrow \begin{bmatrix} f_{\bar{I}} \\ \phi \end{bmatrix} \preceq \begin{bmatrix} f'_{\bar{I}} \\ \phi \end{bmatrix} \quad \text{for all } \phi$$

(*the lower is $f_{\bar{I}}$, the better*)

A benefit ($L = \{\bar{I}\}$) is preferentially independent from all other indicators

$$f_{\bar{I}} \geq f'_{\bar{I}} \Leftrightarrow \begin{bmatrix} f_{\bar{I}} \\ \phi \end{bmatrix} \preceq \begin{bmatrix} f'_{\bar{I}} \\ \phi \end{bmatrix} \quad \text{for all } \phi$$

(*the higher is $f_{\bar{I}}$, the better*)

But this is not true for the thermostat example ($L = \{T\}$, $P \setminus L = \{h\}$):

- at certain humidities, a lower temperature is better
- at other humidities, a higher temperature is better

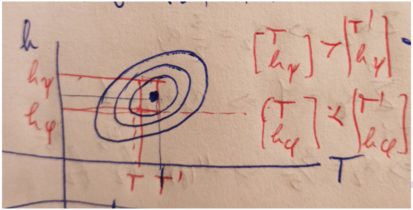

Example: the rule-abiding menu

Decide a menu, combining two elements of alternative (X = Wines $\times$ Main courses)

- a wine out of {Barolo,Nebbiolo,Erbaluce,Arneis}
- a main course out of {Stew,Roast,Meatballs,Salmon,Swordfish}

The impact function f simply projects X onto a rough set F as follows

- all red wines onto impact Red
- all white wines onto impact White
- all meat courses onto Meat
- all fish courses onto Fish

Π follows the classical rule: red wine and meat, white wine and fish

Preferential independence is violated (for both indicators):

$$(\text{Red,Meat}) \prec (\text{White,Meat}), \text{ but } (\text{Red,Fish}) \succ (\text{White,Fish})!$$

$$(\text{Red,Meat}) \prec (\text{Red,Fish}), \text{ but } (\text{White,Meat}) \succ (\text{White,Fish})!$$

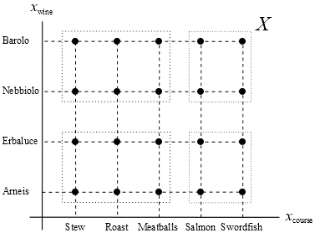

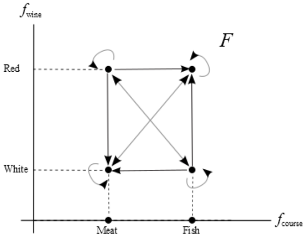

Example

Preferential independence is not symmetric

$$L \text{ independent from } P \setminus L \not\Rightarrow P \setminus L \text{ independent from } L$$

Example: let $F = \{f_1 \geq 0, f_2 \geq 1\}$ with $u(f) = (f_1 - 5)f_2$

- f_1 is independent from f_2: higher values of f_1 are always better
- f_2 is not independent from f_1:
  - lower values of f_2 are better when $f_1 < 5$
  - f_2 is irrelevant when $f_1 = 5$
  - higher values of f_2 are better when $f_1 > 5$

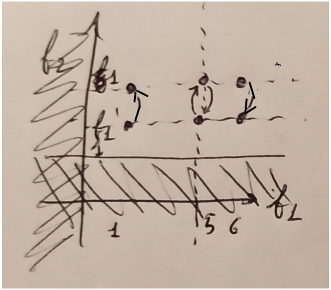

Example

Preferential independence on single indicators
does not imply independence on larger subsets

$\{I\}$ independent from $P \setminus \{I\}$, $\forall I \in P \not\Rightarrow L$ independent from $P \setminus L$, $\forall L \subseteq P$

Example: let $F = \{f_1 \geq 0, f_2 \geq 0, f_3 \geq 1\}$ with $u(f) = \dfrac{1}{(f_1 + f_3)(f_2 + f_3)}$

- each indicator is a cost: higher values reduce the utility
- (f_1, f_2) depends on f_3

$$\begin{bmatrix} 1 \\ 3 \\ 1 \end{bmatrix} \succ \begin{bmatrix} 4 \\ 1/2 \\ 1 \end{bmatrix} \text{ because } u(1,3,1) = 1/8 < u(4,1/2,1) = 2/15$$

but

$$\begin{bmatrix} 1 \\ 3 \\ 3 \end{bmatrix} \prec \begin{bmatrix} 4 \\ 1/2 \\ 3 \end{bmatrix} \text{ because } u(1,3,3) = 1/24 > u(4,1/2,3) = 2/49$$

Mutual preferential independence

Definition
Mutual preferential independence holds when
every proper subset L is independent from its complement $P \setminus L$

- how do we check it?
- why is it useful?

The definition requires to check every nonempty proper subset $P \subset L$

- $2^p - 2$ subsets
- infinite 4-tuples of subimpacts (to sample) for each subset

Theorem
Mutual preferential independence holds if and only if, given $\bar{l} \in P$, every pair $L = \{l, \bar{l}\}$ is preferentially independent from $P \setminus L$ for all $l \in P \setminus \{\bar{l}\}$

We need to check only $p - 1$ pairs (*but single indicators are not enough*)

Mutual preferential independence and additivity

Theorem

If Π admits a consistent additive utility function $u(f)$,
then Π enjoys mutual preferential independence

The problem is that we need the converse (we want to prove additivity)

Theorem

When $p \geq 3$, Π admits a consistent additive utility function $u(f)$
if and only if Π enjoys mutual preferential independence

Unfortunately, when $p = 2$, mutual preferential independence is

- necessary for additivity
- not sufficient for additivity

as it reduces to checking each indicator with respect to the other one

Luckily, two indicators can be at least visualised easily

Counterexample

The two indicators are costs, therefore mutually independent.

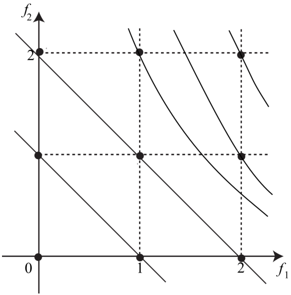

Assume additivity: $u(f_1, f_2) = u_1(f_1) + u_2(f_2)$

$$\begin{cases} (2,0) \sim (0,2) \\ (1,0) \sim (0,1) \end{cases} \Rightarrow \begin{cases} u(2,0) = u(0,2) \\ u(1,0) = u(0,1) \end{cases} \Rightarrow \begin{cases} u_1(2) + u_2(0) = u_1(0) + u_2(2) \\ u_1(1) + u_2(0) = u_1(0) + u_2(1) \end{cases}$$

Subtract the two equalities

$$u_1(2) - u_1(1) = u_2(2) - u_2(1) \Rightarrow u_1(2) + u_2(1) = u_1(1) + u_2(2) \Rightarrow$$
$$\Rightarrow u(2,1) = u(1,2) \Rightarrow (2,1) \sim (1,2)$$

which is false

The compromise between indicators should respect suitable conditions throughout F