

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 theory

We assume

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



The framework is the same as in the previous lesson, but

- **decisions are taken in stages**
- **part of the scenario unravels before part of the decision is taken**

Part of the alternative depends on part of the scenario

Decision theory

The model is similar to a turn-based game between

- decision-maker
- external world

The decision process takes place in a **finite number $t_{\max}$ of stages**

1a) the decision-maker assigns a value to a subvector $x^{(0)}$

1b) part of the scenario unravels, revealing the value of subvector $\omega^{(0)}$

2a) the decision-maker assigns a value to another subvector $x^{(1)}$

2b) part of the scenario unravels, revealing the value of subvector $\omega^{(1)}$

...

$t_{\max}a$) the decision-maker assigns a value to the last subvector $x^{(t_{\max}-1)}$

$t_{\max}b$) the scenario fully unravels, revealing subvector $\omega^{(t_{\max}-1)}$

The situation is similar to a turn-based game between

- decision-maker
- external world

Strategies versus solutions

A solution in basic decision problems is a real vector ($x \in \mathbb{R}^n$)

In decision theory a solution can be a strategy

It must be a **time-consistent strategy**: the decision variables of each stage depend on the exogenous variables of the previous stages

$$x^{(t)} = x^{(t)}(\omega^{(0)}, \dots, \omega^{(t-1)}) \quad t \in \{0, \dots, t_{\max} - 1\}$$

These problems are more complex to represent

- an evaluation matrix is possible, replacing solutions with strategies
(*that tend to be much more numerous*)
- a tree representation is more natural and, therefore, common


Both require finite problems and pose problems in combinatorial ones

We shall discuss matrix representations later on

Decision tree

A **decision tree** is a tree with $2t_{\max} + 1$ levels chronologically ordered with respect to when information is revealed to the decision-maker:

- the **even levels** $2t$ with $t \in \{0, \dots, t_{\max} - 1\}$ represents **decisions** fixing $x^{(t)}$
- the **odd levels** $2t + 1$ with $t \in \{0, \dots, t_{\max} - 1\}$ represent **scenario elements** fixing $\omega^{(t)}$
- level** $2t_{\max}$ represents the **final configurations** (x, ω) , with the associated impacts $f(x, \omega)$ or stochastic utilities $u(f(x, \omega))$



The arcs going out of a node correspond to

- for decision nodes, the possible **values** of $x^{(t)}$
- for scenario nodes, the possible **values** of $\omega^{(t)}$ and associated information (**conditional probabilities** summing to 1)

Example

A company wants to launch a new model on the market

- three models are available: $X = \{A, B, C\}$
- three demand scenarios are predictable: $X = \{\text{High, Medium, Low}\}$
- profits $f(x, \omega)$ and scenario probabilities $\pi(\omega)$ have been generated with suitable predictive models

If only one model can be launched, which one should be chosen?

The problem can be represented in the standard way,
with an evaluation matrix and a probability vector

$f(x, \omega)$		Demand level ω		
Model x		Low	Medium	High
A		200 000	350 000	600 000
B		250 000	350 000	540 000
C		300 000	375 000	490 000

	Low	Medium	High
Probability π	0.1	0.5	0.4


Example

A company wants to launch a new model on the market

- three models are available: $X = \{A, B, C\}$
- three demand scenarios are predictable: $X = \{\text{High, Medium, Low}\}$
- profits $f(x, \omega)$ and scenario probabilities $\pi(\omega)$ have been generated with suitable predictive models

If only one model can be launched, which one should be chosen?

But it can also be represented (equivalently) with a decision tree



In this simple case, the tree looks redundant


Backward induction

Anyway, how can the problem be solved on a decision tree?

The **backward induction** algorithm visits the tree from leaves to root, assigning a value to the current node based on the values of its children

- in the odd levels $2t + 1$, apply a choice criterium ϕ
- in the even levels $2t$, select the best value and mark the corresponding arc

The marked arcs provide the optimal strategy




Let us apply the worst-case criterium

Example: the worst-case criterium

If we apply the worst-case criterium

- 1 the leaves (level 2) are associated to the impacts $f(x, \omega)$
- 2 the nodes at level 1 depend on scenarios unravelling after a decision: copy in the node the worst (minimum) value in its children, that is 200 000, 250 000 and 300 000
- 3 the root (level 0) depends on the solution choice: copy in the node the best (maximum) value in its children: 300 000 mark the arc corresponding to model C

Consequently, $x^\circ = C$ and $\phi_{\text{worst}}(x^\circ) = 300\,000$ (overall, $C \prec A \prec B$)




Example: the Laplace criterium

If we apply the Laplace criterium

- 1 the leaves (level 2) are associated to the impacts $f(x, \omega)$
- 2 the nodes at level 1 depends on scenarios unravelling after a decision: copy in the node the arithmetic mean of its children, that is 383 333, 380 000 and 388 333
- 3 the root (level 0) depends on the solution choice: copy in the node the best (maximum) value in its children: 300 000 mark the arc corresponding to model C

Consequently, $x^\circ = C$ and $\phi_{\text{worst}}(x^\circ) = 388\,333$ (overall, $C \prec A \prec B$)




Example: the expected value criterium

If we apply the expected value criterium

- 1 the leaves (level 2) are associated to the impacts $f(x, \omega)$
- 2 the nodes at level 1 depends on scenarios unravelling after a decision: copy in the node the expected value of its children, that is 435 000, 416 000 and 413 500
- 3 the root (level 0) depends on the solution choice: copy in the node the best (maximum) value in its children: 300 000 mark the arc corresponding to model C

Consequently, $x^\circ = C$ and $\phi_{\text{worst}}(x^\circ) = 435,000$ (overall, $A \prec B \prec C$)



Natural applications of decision trees

The tree representation is better than the matrix representation when

- ① scenario probabilities are conditioned by the decision
 - The decision-maker has an influence on the external world
- ② the decision takes place in more stages
 - The decision becomes a strategy
- ③ random experiments to improve the estimation of probabilities can be performed before the decision
 - The probabilities are conditioned by the outcome of the experiment
 - The decision is a strategy depending on the outcome

Scenario probabilities conditioned by the decision

Suppose that the choice of the model to launch affects the demand, modifying the probability of each scenario

Replace probability vector $\pi(\omega)$ by conditional probability matrix $\pi(\omega|x)$

$f(x, \omega)$		Demand level ω		
Model x		Low	Medium	High
A		200 000	350 000	600 000
B		250 000	350 000	540 000
C		300 000	375 000	490 000

$\pi(\omega x)$		Demand level ω		
Model x		Low	Medium	High
A		0.3	0.5	0.2
B		0.1	0.5	0.4
C		0.1	0.6	0.3

Multi-stage decisions

After launching a model and discovering the demand level, the company could try an advertising campaign and assess its effect

To make it extremely (unrealistically) simple

- the company either makes a campaign or not
(instead of choosing among several alternatives)
- the campaign has a predictable deterministic effect
(instead of several possible ones)


This yields a multi-stage model

- $t_{\max} = 2$ decision stages imply a decision tree with 5 levels
 - ① choice of the model: $x^{(1)} \in \{A, B, C\}$
 - ② demand level: $\omega^{(1)} \in \{\text{High, Medium, Low}\}$
 - ③ choice of the campaign: $x^{(2)} \in \{\text{Yes, No}\}$
 - ④ effect of the campaign: $\omega^{(2)} \in \{\text{Given}\}$
 - ⑤ impact: $f(x^{(1)}, \omega^{(1)}, x^{(2)}, \omega^{(2)})$, combining the original revenues and costs with the campaign cost and the increase in sales
- in general, all probabilities are conditioned by the previous levels

$$\pi(\omega^{(1)}|x^{(1)}) \quad \pi(\omega^{(2)}|x^{(1)}, \omega^{(1)}, x^{(2)}) \quad \dots$$

(not in this very simplified case)

Multi-stage decisions



- 1 the leaves (level 4) are associated to the impacts $f(x, \omega)$
- 2 the nodes at level 3 depend on scenarios about the campaign effect, that here are deterministic: just copy the children value
- 3 the nodes at level 2 depend on making or not the campaign: copy in the node the best value in its children and mark the associated arc
- 4 the nodes at level 1 depend on scenarios about the demand level: apply the expected value criterium with the probability estimates
- 5 the root (level 0) depends on the model chosen: copy in the node the best value in its children and mark the associated arc

$x^\circ(\omega)$ is a time-consistent strategy:

- choose model C
- if the demand is low or medium, advertise; if it is high, do not advertise

Random experiments

Random experiment is an **action** whose outcome depends on the scenario

Given the outcome ω'

- in general, the exact scenario ω remains unknown
- but **all scenarios incompatible with outcome ω' can be ruled out**
- **the scenario set reduces from Ω to $\bar{\Omega}(\omega') \subseteq \Omega$,**
that is the **subset of scenarios compatible with ω'**
- therefore, in general $\pi(\omega|\omega') \neq \pi(\omega)$

The experiment provides information on the scenario

Example: build a model of a bridge and let models of car move on it


- if the model collapses, the bridge might stand, but it is unlikely
- if the model stands, the bridge might collapse, but it is unlikely

It is uncertain, but precious, information

Random experiments and decision-making

Random experiments can be incorporated in decision trees

- 1 decide whether to **make the experiment or not** (x')
- 2 the experiment has uncertain **outcome ω' with probability $\pi(\omega')$** ; if it is not made, the outcome is deterministic
- 3 the **main decision** takes place (x)
- 4 the uncertain **scenario ω has**
 - **probability $\pi(\omega|\omega')$ conditioned by the outcome of the experiment**
 - **has an **impact $f(x, x', \omega)$ including the cost of the experiment****
(*building and running models, making interviews or polls*)



In summary, we need $f(x, x', \omega)$, $\pi(\omega')$ and $\pi(\omega|\omega')$

The problem is that usually we do not have the required probabilities

Example

We want to make a trip

- we can wear light clothes, bring also an umbrella or warm clothes

$$X = \{L, U, W\}$$

- the weather can be nice or bad (and, of course, we cannot predict it)

$$\Omega = \{N, B\}$$

- we have *a priori* information on the typical seasonal probabilities

	N	B
$\pi(\omega)$	0.4	0.6

- the impacts, expressed as costs, are estimated as

$f(x, \omega)$	N	B
L	0	5
LU	1	3
W	3	2

Example

$f(x, \omega)$	N	B
L	0	5
U	1	3
W	3	2

	N	B
$\pi(\omega)$	0.4	0.6

A risk-neutral decision-maker simply applies the expected value criterium

- $\phi_{EV}(L) = 0.4 \cdot 0 + 0.6 \cdot 5 = 3.0$
- $\phi_{EV}(U) = 0.4 \cdot 1 + 0.6 \cdot 3 = 2.2$
- $\phi_{EV}(W) = 0.4 \cdot 3 + 0.6 \cdot 2 = 2.4$

and decide to wear light clothes and bring an umbrella

But barometer gives current pressure, related to future weather

Example

A first decision, therefore, is to read or not to read

$$X' = \{Y, N\}$$

The barometric reading has three possible outcomes: fair, change, or rain

$$\Omega' = \{F, C, R\}$$

In order to build the decision tree, we need

- $\pi(\omega')$: absolute probabilities of the barometric readings
- $\pi(\omega|\omega')$: conditional probabilities of the scenarios with respect to the barometric readings

Example

The barometer manufacturer provides the precision of the instrument based on historical observations (*unreliable values, better than nothing*)

- $\pi(\omega'|\omega)$: conditional probabilities of the readings with respect to the actual future weather

$\pi(\omega' \omega)$	N	B
F	0.60	0.20
C	0.25	0.30
R	0.15	0.50


Of course, the values on each column sum up to 1

We know $\pi(\omega)$ and $\pi(\omega'|\omega)$: how to compute $\pi(\omega')$ and $\pi(\omega|\omega')$?

Example

The situation can be modelled as follows

- the overall rectangle is the scenario set Ω
- the single scenarios ω are the disjoint portions A_i
(in our example, only two: N and B)
- scenario probabilities $\pi(\omega)$ are proportional to the areas
- the green area is one of the possible experiment outcomes ω'
- we know the fraction of green area in each portion A_i , i. e. $\pi(\omega'|\omega)$



We want

- the green area $\pi(\omega')$
- the fraction $\pi(\omega|\omega')$ of each portion A_i in the green area

Bayes theorem

The law of conditional probability and the symmetry of set intersection imply that

$$\pi(A_i|B) \pi(B) = \pi(A_i \cap B) = \pi(B \cap A_i) = \pi(B|A_i) \pi(A_i)$$

In our case we ignore

- $\pi(A_i|B)$, that is $\pi(\omega|\omega')$
- $\pi(B)$, that is $\pi(\omega')$

but we know

- $\pi(B|A_i)$, that is $\pi(\omega'|\omega)$
- $\pi(A_i)$, that is $\pi(\omega)$

The process is quite simple

- 1 Compute the **conjoint probabilities** $\pi(A_i \cup B)$

$$\pi(\omega \cdot \omega') = \pi(\omega'|\omega) \pi(\omega)$$

These are the elementary shards of the partition

Example

In our example

$\pi(\omega' \omega)$	N	B
F	0.60	0.20
C	0.25	0.30
R	0.15	0.50

$\pi(\omega)$	0.40	0.60
---------------	------	------

imply by simple products

$\pi(\omega \cdot \omega')$	N	B
F	0.24	0.12
C	0.10	0.18
R	0.06	0.30

By definition, the overall matrix sums to 1

Example

- ② Compute the **total probabilities of the experiment outcomes** ($\pi(B)$)

$$\pi(\omega') = \sum_{\omega \in \Omega} \pi(\omega \cdot \omega')$$

In our example

$\pi(\omega' \omega)$	N	B
F	0.60	0.20
C	0.25	0.30
R	0.15	0.50

$\pi(\omega)$	0.40	0.60
---------------	------	------

$\pi(\omega \cdot \omega')$	N	B
F	0.24	0.12
C	0.10	0.18
R	0.06	0.30

$\pi(\omega')$
0.36
0.28
0.36

By definition, the new vector sums to 1

Example

- ③ Compute the **conditional probabilities** $(\pi(B|A_i))$

$$\pi(\omega|\omega') = \frac{\pi(\omega \cdot \omega')}{\pi(\omega')}$$

In our example

$\pi(\omega' \omega)$	N	B
F	0.60	0.20
C	0.25	0.30
R	0.15	0.50

$\pi(\omega)$	0.40	0.60
---------------	------	------


$\pi(\omega \cdot \omega')$	N	B
F	0.24	0.12
C	0.10	0.18
R	0.06	0.30

$\pi(\omega')$
0.36
0.28
0.36

$\pi(\omega' \omega)$	N	B
F	24/36	12/36
C	10/28	18/28
R	6/36	30/36

By definition, each row of the new matrix sums to 1

Example



Example

The upper branch (no experiment) has


- a single deterministic scenario ω'
- the original probabilities $\pi(\omega)$

The lower branch (experiment) has

- three scenarios ω'
- conditional probabilities $\pi(\omega|\omega')$

Backward induction provides a strategy

- make the experiment
- if $\omega' = F$, then $x \in \{L, U\}$
- if $\omega' = C$, then $x = U$
- if $\omega' = R$, then $x = W$



Value of the experiment

The expected cost is

- 2.2 with no experiment
- 2.02 with the experiment (plus the experiment cost, if any)

Value of the experiment is the improvement allowed by the experiment

In this case, $v = 2.2 - 2.02 = 0.18$

It is the maximum cost one is willing to pay for the information provided

