GSOE9210 Engineering Decisions

Victor Jauregui

vicj@cse.unsw.edu.au
www.cse.unsw.edu.au/~gs9210

Maximin and minimax regret

Outline

The *Maximin* principle

Definition (The *Maximin* principle)

Assume that only the minimally preferred outcomes will occur and choose those actions that lead to the most preferred among these.

- *Maximin* and *miniMax Regret* are instances of the *Maximin* principle: original values vs regrets
- The *Maximin* principle is the main decision principle used under complete uncertainty
- We've seen *Maximin* and *miniMax Regret* on decision tables, but what about more complex decision problems (*e.g.*, multiple decision points)?

Multi-stage decisions

Example (Product development)

You head the R&D department of a small manufacturing company which is considering developing a new product. The company must decide whether to proceed with prototype development and, if development is successful, subsequently determine the production scale (*i.e.*, the size of the factory) based on projected demand for the product.

Questions

- What does *Maximin* or *miniMax Regret* mean in this problem?
- Is there a decision-table representation?

Multi-stage decisions

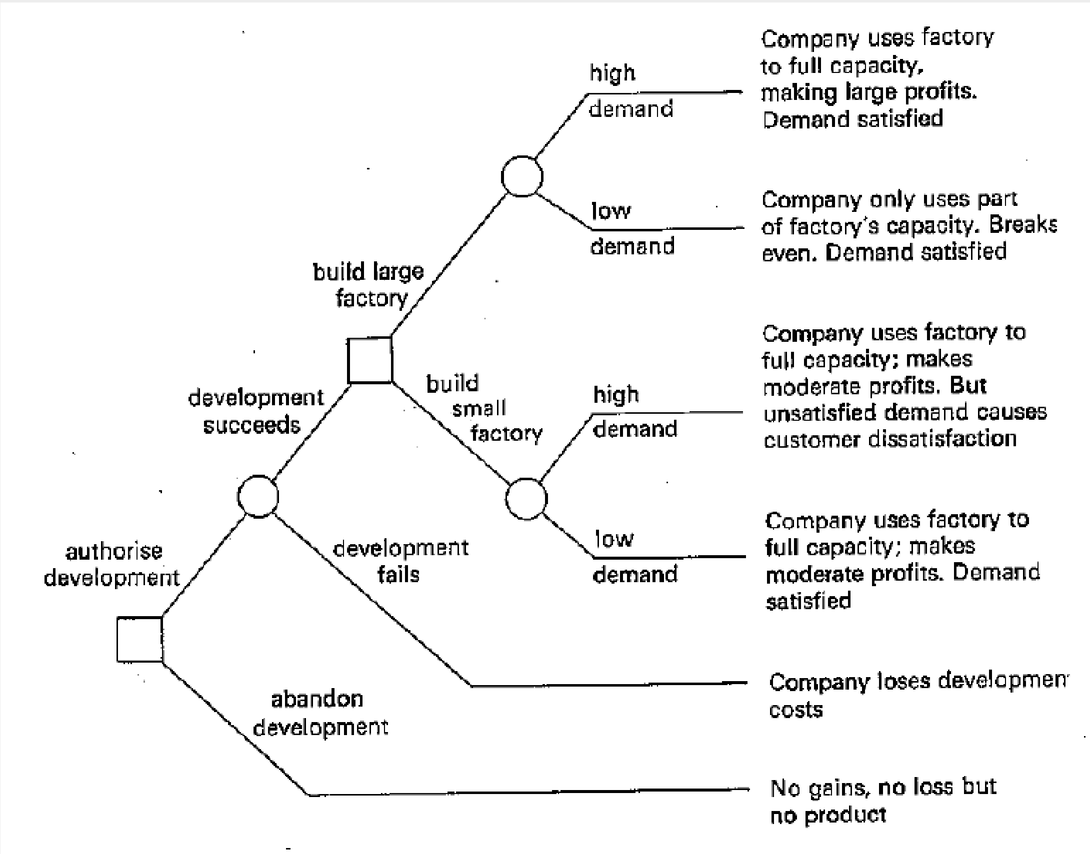

Node evaluation

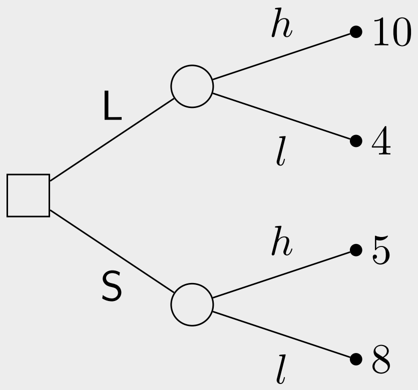

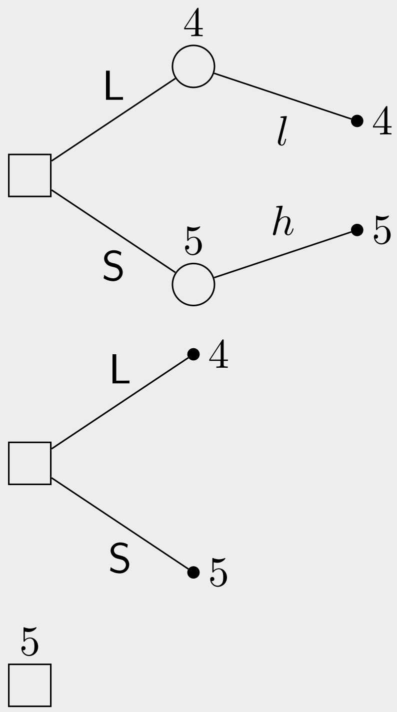

- What does *Maximin* mean in a tree?
- *Maximin* eliminates branches in chance nodes (*i.e.*, prunes the tree)
- Reduces problem to that of certainty

Node evaluation

- Each decision problem is assigned a ‘value’ by a decision rule
- The *Maximin* algorithm for decision trees:
  1. Begin at the leaves of the tree
  2. At each parent:
     1. if a chance node, *Maximin* prunes all children except the minimally preferred
     2. if a decision node, the *elimination principle*, eliminates all children except the maximally preferred
     3. propagate the remaining value up to the parent node
  3. Repeat the previous step until the root is reached
- The value propagated to the root is the *value* of the problem (under *Maximin*); *i.e.*, the value which an agent using *Maximin* assigns to the problem

Outline

Problem representation: decision tables

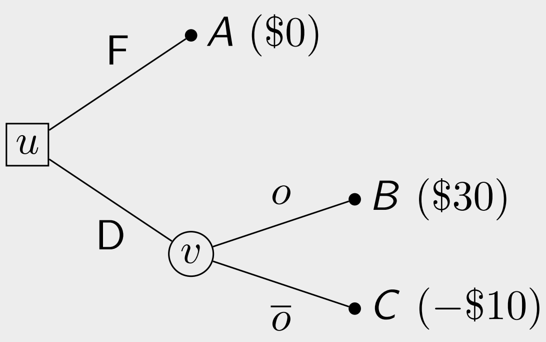

- Observation: Each combination of an action and a state uniquely determine an outcome
- Model as a binary function: $\omega : \mathcal{A} \times \mathcal{S} \to \Omega$

Represented as a table:

		$\mathcal{S}$	
	ω	o	$\overline{o}$
$\mathcal{A}$	F	A	A
	D	B	C

Decision tables:

- row = action
  column = state
- Interpretation: $\text{B} = \omega(\text{D}, o)$ means "B is the outcome of action D in state o";

Trees and tables

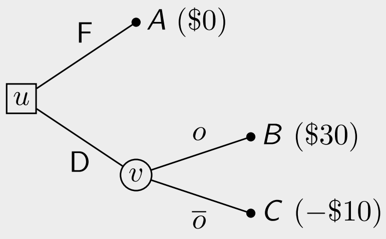

	o	$\overline{o}$
F	\$0	\$0
D	\$30	−\$10

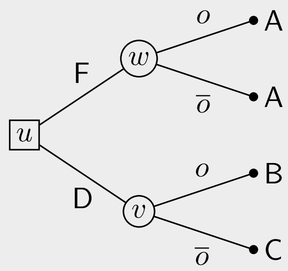

- Multiple trees may correspond to the same table
- Going from tables (*normal form*) to trees (*extensive form*) is straight forward, but the converse can be tricky
- Which representation is better: trees or tables?
- Which representation facilitates decision analysis most?

Multi-stage decisions

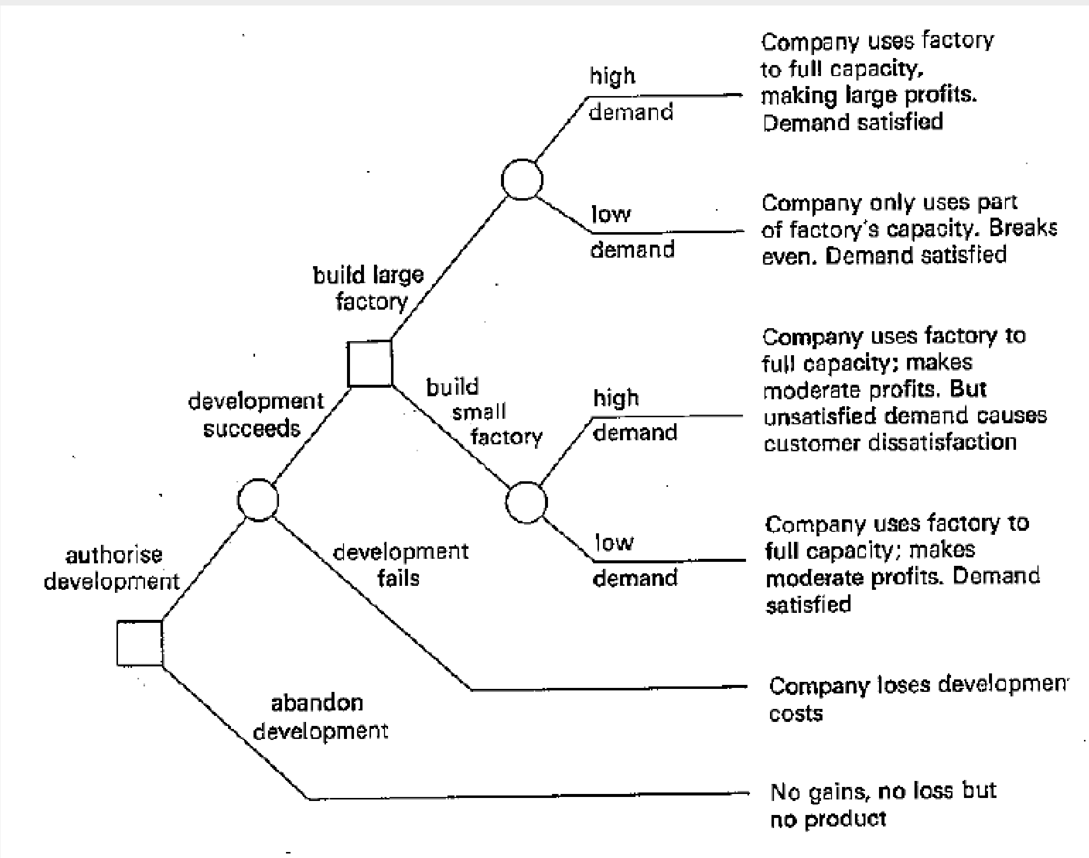

Multi-stage decisions

Example (Product development)

You head the R&D department of a small manufacturing company which is considering developing a new product. The company must decide whether to proceed with prototype development and, if development is successful, subsequently determine the production scale (*i.e.*, the size of the factory) based on projected demand for the product.

Questions

- What does *Maximin* or *miniMax Regret* mean in this problem?
- Is there a decision-table representation?

Actions to strategies

In a decision tree:

- Recall that a decision table is a representation of the outcome mapping $\omega : \mathcal{A} \times \mathcal{S} \to \Omega$
- Observation: following a path from the root to a leaf leads to a unique outcome
- Generalising:
  - A 'state' must specify *conditions* in chance nodes
  - An 'action' must specify *actions* at decision nodes

Definition (Strategy)

A *strategy* (or *policy* or *plan*) is a procedure that specifies the selection of an action at every *reachable* decision point.

Normalisation

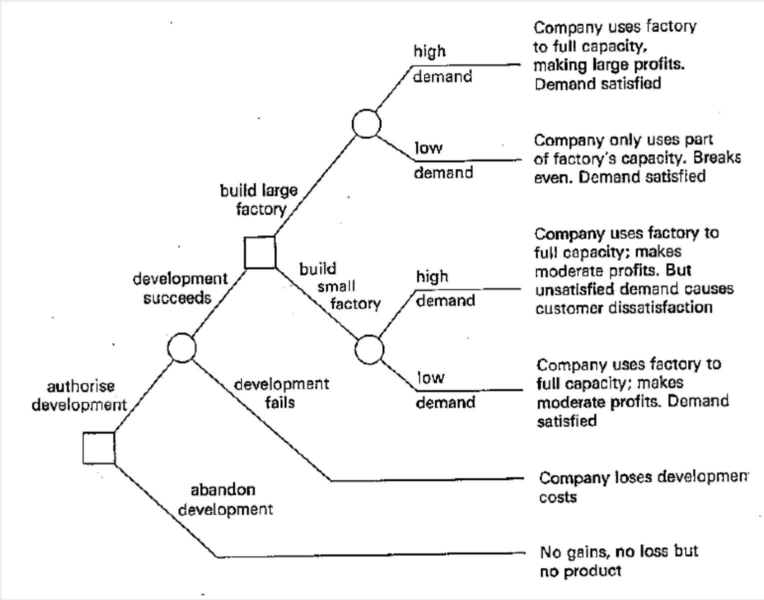

- States: $\begin{array}{ccc} s_1 & s_2 & s_3 \\ \hline s,h & s,l & f \end{array}$
- A strategy must specify an action at each *reachable* decision point; *e.g.*, "Authorise development (Au), if development succeeds (s), then build large factory (L)"

Normalisation

Encoding:

- α/A says:

  At the decision node reached via path α choose action A.

- Example: Au;s/S:

  If development has been authorised (Au) and has succeeded (s), choose to build a small factory (S).

- Strategies for this problem:

  A$_1$ Au;s/L

  A$_2$ Au;s/S

  A$_3$ Ab

Normalisation

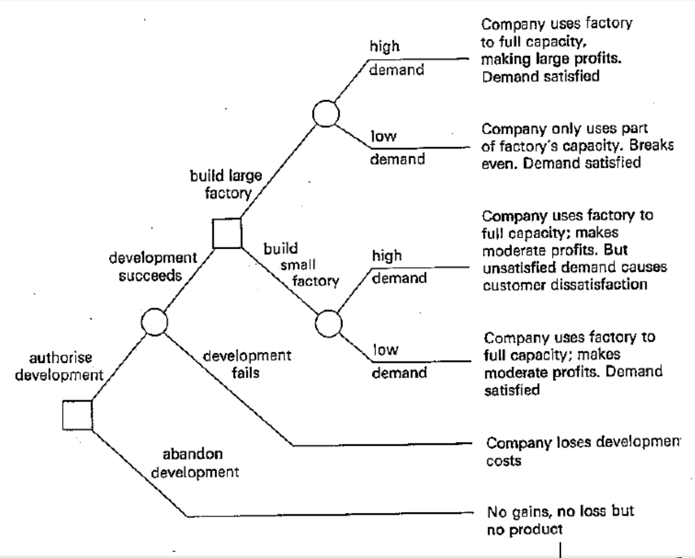

Code	*Description*
fc	full capacity
pc	partial capacity
lp	large profits
mp	moderate profits
be	break even
ldc	lose dev. costs
sat	demand satisfied
dis	dissatisfaction
sq	status quo

	s, h	s, l	f
Au;s/L	fc,lp,sat	pc,be,sat	ldc
Au;s/S	fc,mp,dis	fc,mp,sat	ldc
Ab	sq	sq	sq

Normalisation

Outcome values:

ω	v
fc,lp,sat	10
pc,be,sat	4
ldc	-1
fc,mp,dis	5
fc,mp,sat	8
sq	0

Decision table:

	s, h	s, l	f
Au;s/L	10	4	-1
Au;s/S	5	8	-1
Ab	0	0	0

Exercises

- Find the *Maximin* and *miniMax Regret* strategies for this problem.
- Evaluate this problem under *MaxiMax*, *Maximin*, *miniMax Regret* using both normal and extensive forms.

Representing information

Consider the fund-raiser example.

- Decision before weather known:

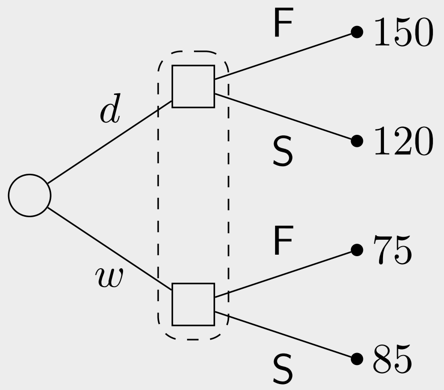

- Decision nodes part of the same *information set*
- Available strategies: F, S only

- Decision after weather known:

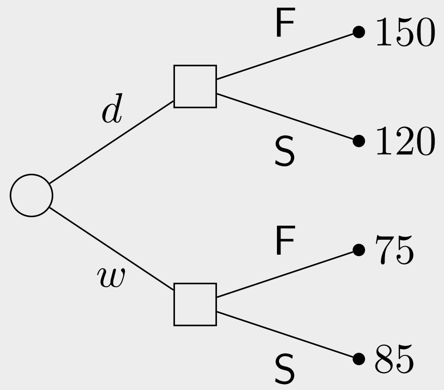

- Decision nodes distinguishable
- Possible strategy: *e.g.*, d/F

Outline

Indifference: equal preference

- Which action below is preferred above under *Maximin*?

	s_1	s_2
A	1	0
B	0	1

Definition (Indifference)

If two actions A and B are *equally preferred* then the agent is said to be *indifferent* between A and B.

- Indifference means an agent prefers two alternatives equally, not that it doesn't *know* which it prefers

Indifference classes

Definition (Indifference class)

An *indifference class* is a non-empty set of all actions/outcomes between which an agent is indifferent.

- For a given action A $\in \mathcal{A}$, the indifference class of A is given by

$$I(\mathsf{A}) = \{a \in \mathcal{A} \mid V(a) = V(\mathsf{A})\}$$

- Indifference classes partition set of all actions
- Different agents have different preferences over outcomes/actions, hence different indifference classes
- Different decision rules evaluate actions differently; *i.e.*, produce different indifference classes

Outline

Graphical representation

	s_1	s_2
A	2	3
B	4	0
C	3	3
D	5	2
E	3	5

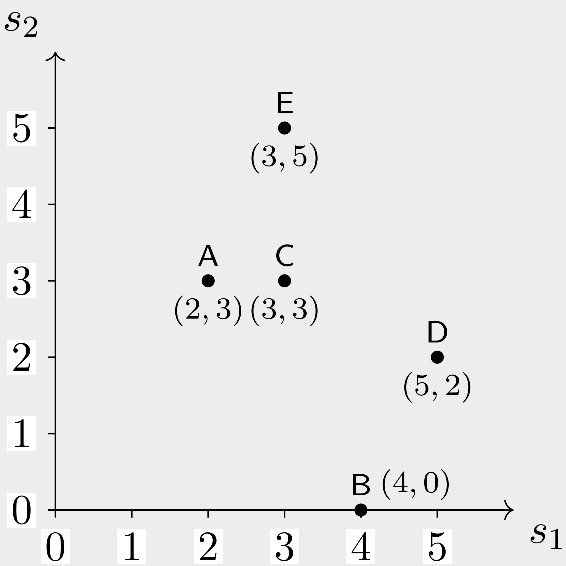

Let $v_i(a) = v(a, s_i)$ be the value of action a in state s_i. Each action a corresponds to a point (v_1, v_2), where $v_i = v(a, s_i)$.

Indifference curves: *Maximin*

For the pure actions below:

	s_1	s_2
A	2	3
B	4	0
C	3	3
D	5	2
E	3	5

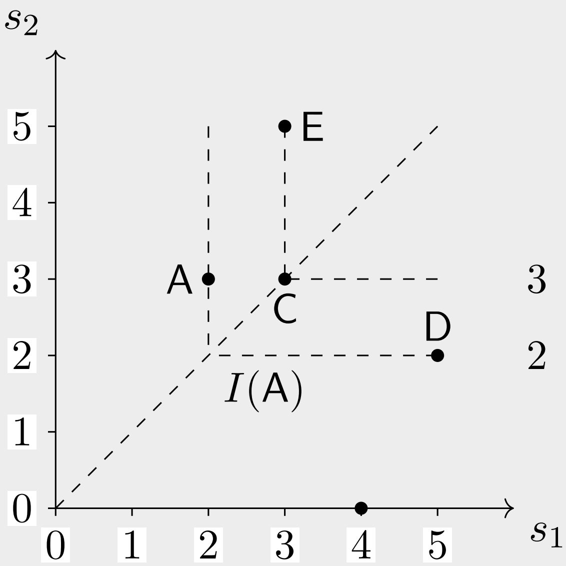

Consider curves of all points representing strategies with same *Maximin* value; *i.e.*, *Maximin indifference curves*.

Graphing regret

- Consider three actions:

	s_1	s_2
A	2	4
B	4	1
C	5	3

	s_1	s_2
A	3	0
B	1	3
C	0	1

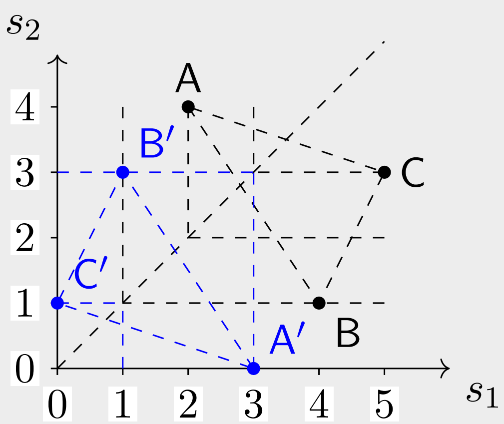

- Regrets and indifference curves for *miniMax Regret* in blue

Exercises

In regard to preference over actions, what is the relation between *Maximin* and *miniMax Regret*?

Outline

River example

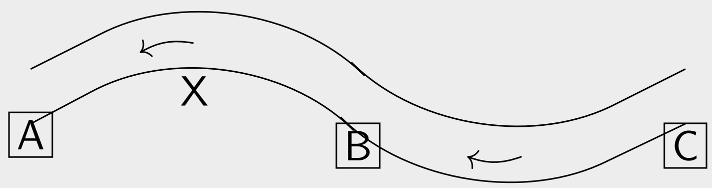

Example (River logistics)

Alice's warehouse is located at X on a river that flows down-stream from C to A. She delivers goods to a client at C via motor boats. On some days a (free) goods ferry travels up the river, stopping at A then B and C, but not at X.

The fuel required (litres) to reach C from each starting point:

	A	X	B	C
To C from:	4	3	2	0

Alice wants to minimise fuel consumption (in litres).

River example

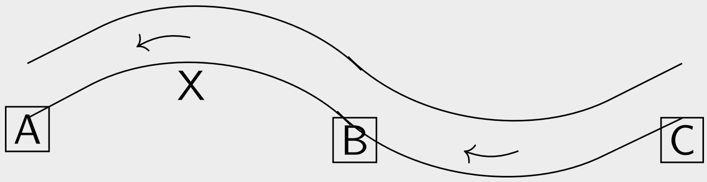

	f	$\overline{f}$
A	4	0
B	3	1
C	1	1

Alice considers three possible ways to get to C (from starting point X):

A : via A, by floating down the river

B : via B, by travelling up-stream to B

C : by travelling all the way to C

Outcomes are measured in *litres left* in a four-litre tank.

Exercise

Let $w : \Omega \to \mathbb{R}$ denote fuel consumption in litres. What transformation $f : \mathbb{R} \to \mathbb{R}$ is responsible for the values $v : \Omega \to \mathbb{R}$ in the decision table?

River example

- Axes correspond to payoffs in each of the two states; *i.e.*, payoff v_1 in state $s_1 = f$ and v_2 in $s_2 = \overline{f}$
- Actions graphed below:

	f	$\overline{f}$
A	4	0
B	3	1
C	1	1

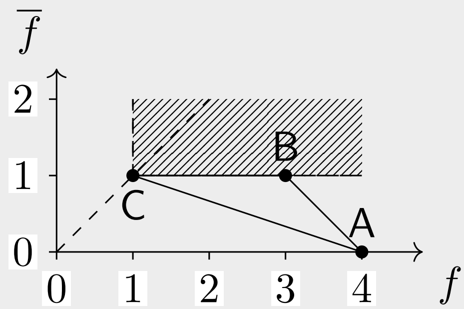

- Option C not a better response than B under any circumstances (*i.e.*, in any state)
- C worse than B in some cases and no better in all others; C can be *discarded*

Generalised dominance

Definition (Strict dominance)

Strategy A *strictly dominates* B iff every outcome of A is more preferred than the corresponding outcome of B.

Definition (Weak dominance)

Strategy A *weakly dominates* B iff every outcome of A is no less preferred than the corresponding outcome of B, and some outcome is more preferred.

	s_1	s_2	s_3
A	3	4	2
B	4	4	3
C	5	6	3

Exercise

Which strategies in the decision table shown are dominated?

Dominance and best response

Corollary

Strategy A strictly dominates B iff A is a better response than B in each possible state.

Corollary

Strategy A weakly dominates B iff A is a better response than B in some possible state and B is not a better response than A in any state.

Dominance principle

A rational agent should never choose a dominated strategy.

Admissible actions

	s_1	s_2
A	0	4
B	3	1
C	2	3
D	1	2

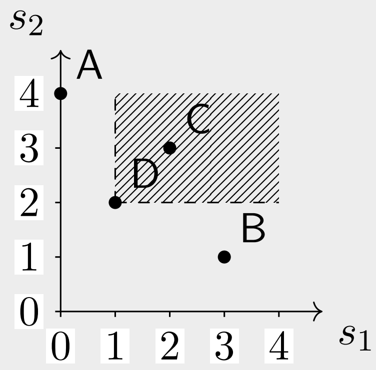

Definition (Admissible)

An action is *admissible* iff it is not dominated by any other action. An action which is not admissible is said to be *inadmissible*. The set of all admissible actions is called the *admissible frontier*.

Exercises

Which actions above are admissible?

Dominance: *MaxiMax* and *Maximin*

	s_1	s_2	M	m
A	2	2	2	2
B	2	1	2	1
C	1	1	1	1

Definition (Dominance elimination)

A decision rule is said to satisfy (strict/weak) *dominance elimination* if it never chooses actions that are (strictly/weakly) dominated.

- Dominated actions can be discarded under any rule that satisfies dominance elimination

Dominance summary

Rules that satisfy strict/weak dominance elimination.

Rule	Strict	Weak
MaxiMax	$\surd$	$\times$
Maximin	$\surd$	$\times$
Hurwicz's	$\surd$	$\times$
miniMax Regret	$\surd$	$\times$
Laplace's	$\surd$	$\surd$

Exercise

Verify the properties above.

Rule axioms

The following criteria can be used to assess the suitability of decision rules:

Axiom of dominance

A decision rule should never choose a dominated action.

Axiom of representational invariance

A decision rule's choices should be independent of representation.

Axiom of solubility

A decision rule should always select at least one action.

Axiom of state duplication independence

Adding a duplicate state should not affect a rule's decision.

Summary: decisions under complete uncertainty

- *Maximin* in extensive form
- Multi-stage decisions
- Extensive to normal form translation
- Graphical visualisation
- Indifference classes
- Dominance and admissibility