

Introduction to Decision and Game Theory

Mehdi Dastani
BBL-521
M.M.Dastani@uu.nl

Decision and Game Theory

- ▶ What are the subject matters of decision and game theory and which phenomena do they help us understand?
- ▶ What are the problems of decision and game theory?
- ▶ What are the elementary concepts of decision and game theory?
- ▶ What is the relevance of decision and game theory to autonomous agents research?
- ▶ How can decision- and game-theoretic concepts be put to use so as to design better systems?

Example: Defence-Attack

Situation: Attacker (Red, column player) can attack either target A or target B, but not both. Defender (Blue, row player) can defend either of two targets but not both. Target A is three times as valuable as Target B.




	A	B
A	4, 0	3, 1
B	1, 3	4, 0

Question: Which target is Red to attack and which target is Blue to defend?

Example: Defence-Attack

Situation: Attacker (Red, column player) can attack either target A or target B, but not both. Defender (Blue, row player) can defend either of two targets but not both. Target A is three times as valuable as Target B.



	A	B
A	4, 0	3, 1
B	1, 3	4, 0

Question: Which target is Red to attack and which target is Blue to defend?

Battle of the Sexes

John and Mary agreed to go out. They can attend a ballet performance or a box match. Mary would like to go to the ballet performance while John would most of all like to go to the box match. Both prefer to go to the same place rather than different ones. If they cannot communicate, where should they go?

	<i>Fight</i>	<i>Ballet</i>
<i>Fight</i>	2, 1	0, 0
<i>Ballet</i>	0, 0	1, 2


Game of Chicken

Two drivers headed each other from opposite directions. The one to turn aside loses. If neither player turn aside, the result is a deadly collision. The best outcome for each driver is to stay straight while the other turns aside and the worst outcome for both driver is to have a deadly collision. In this situation each player wants to secure his/her best outcome, risking the worst scenario.

	<i>Aside</i>	<i>Straight</i>
<i>Aside</i>	0, 0	-5, 5
<i>Straight</i>	5, -5	-10, -10


Decision Theory: An Agent Plays Against Environment

- ▶ An agent is autonomous if it is capable of *deciding* actions in order to achieve its objectives.
- ▶ Classical Decision Theory (Savage 1954)
 - ▶ **probability** and **utility** functions



Decision Theory: An Agent Plays Against Environment

- ▶ An agent is autonomous if it is capable of *deciding* actions in order to achieve its objectives.
- ▶ Classical Decision Theory (Savage 1954)
 - ▶ **probability** and **utility** functions




- ▶ Decision rule = *maximum expected utility* for each action “a” given the set of outcomes O

$$EU(a) = \sum_{o \in O} U(o) * P(o | a)$$

Decision Theory: An Agent Plays Against Environment

- ▶ An agent is autonomous if it is capable of *deciding* actions in order to achieve its objectives.
- ▶ Classical Decision Theory (Savage 1954)
 - ▶ **probability** and **utility** functions



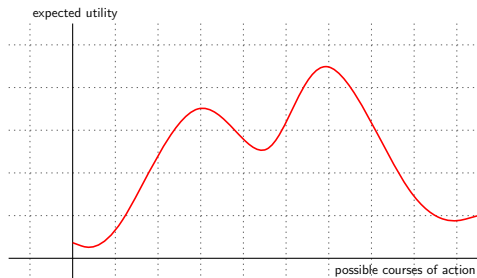
- ▶ Decision rule = *maximum expected utility* for each action “a” given the set of outcomes O

$$EU(a) = \sum_{o \in O} U(o) * P(o | a)$$

$$EU(\text{Work}) = (0.7 * 10) + (0.3 * 1) = 7.3$$

$$EU(\text{Robbery}) = (0.1 * 10) + (0.9 * 1) = 1.9$$

Decision Theory



Issue: Find the course of action that maximizes expected utility given particular environmental parameters.

Utilities and Preferences (1)

- ▶ An agent's **Utility** quantifies its degree of **preferences** over a set $O = \{o_1, \dots, o_n\}$ of outcomes.
- ▶ “The agent prefers weakly o_1 to o_2 ” is denoted by $o_1 \succsim o_2$.
 - ▶ $o_1 \succ o_2$ iff $o_1 \succsim o_2$ and not $o_2 \succsim o_1$.
 - ▶ $o_1 \sim o_2$ iff $o_1 \succsim o_2$ and $o_2 \succsim o_1$.
- ▶ An agent's *Preference*, denoted by \succsim , over a set of outcomes O is a reflexive, transitive, and complete relation on O .
 - ▶ Reflexivity: $\forall o \in O : o \succsim o$.
 - ▶ Transitivity: $\forall o_1, o_2, o_3 \in O : \text{if } o_1 \succsim o_2 \text{ and } o_2 \succsim o_3, \text{ then } o_1 \succsim o_3$.
 - ▶ Completeness: $\forall o_1, o_2 \in O : o_1 \succsim o_2 \text{ or } o_2 \succsim o_1 \text{ or } o_1 \sim o_2$.

Utilities and Preferences (2)

- ▶ Substitutability (indifference in outcomes implies indifference in actions):

If $o_1 \sim o_2$, **then** $[p : o_1, p_3 : o_3, \dots, p_k : o_k] \sim [p : o_2, p_3 : o_3, \dots, p_k : o_k]$

for all outcomes o_3, \dots, o_k and probabilities p, p_3, \dots, p_k ($p + \sum_{i=3}^k p_i = 1$).

- ▶ Decomposability (indifference in actions with similar expected outcomes):

if $\forall o_i \in O : P(o_i | a_1) = P(o_i | a_2)$, **then** $a_1 \sim a_2$

- ▶ Monotonicity:

if $o_1 > o_2$ **and** $p > q$, **then** $[p : o_1, 1 - p : o_2] > [q : o_1, 1 - q : o_2]$

- ▶ Continuity:

if $o_1 > o_2 > o_3$, **then** $\exists p \in [0, 1]$ **such that** $o_2 \sim [p : o_1, 1 - p : o_3]$

Utilities and Preferences (3)

Lemma: If a preference relation \succsim satisfies Completeness, Transitivity, Decomposability, and Monotonicity, and if $o_1 \succ o_2 \succ o_3$, then $\exists p \in [0, 1]$ such that

- ▶ $\forall p' : p' < p : o_2 \succ [p' : o_1 ; 1 - p' : o_3]$, and
- ▶ $\forall p'' : p'' > p : [p'' : o_1 ; 1 - p'' : o_3] \succ o_2$

Theorem: (Von Neumann and Morgenstern, 1944) If a preference relation \succsim satisfies Reflexivity, Transitivity, Completeness, Substitutability, Decomposability, Monotonicity, and Continuity, then there exists a utility function $u : O \rightarrow [0, 1]$ with the properties that:

- ▶ $u(o_1) \geq u(o_2)$ iff $o_1 \succsim o_2$, and
- ▶ $u([p_1 : o_1, \dots, p_k : o_k]) = \sum_{i=1}^k p_i u(o_i)$.

Exercise: Give examples of some outcomes and actions that satisfy the abovementioned properties and examine the lemma/theorem.

Utilities and Preferences (4)

Fact: All preference relations over a countable set O are representable by a utility function. These utility functions are invariant under monotonically increasing functions.

Fact: Let $O = \mathbb{R} \times \mathbb{R}$ and \succsim be the *lexicographic order* on O :

$$(o_1, o'_1) \succsim (o_2, o'_2) \text{ iff } o_1 > o_2 \text{ or both } o_1 = o_2 \text{ and } o'_1 \geq o'_2$$

Then, \succsim *cannot* be represented by a utility function.

Lexicographical Preference Order

Fact: Let $O = \mathbb{R} \times \mathbb{R}$ and \succsim be the *lexicographic order* on O :

$$(o_1, o'_1) \succsim (o_2, o'_2) \text{ iff } o_1 > o_2 \text{ or both } o_1 = o_2 \text{ and } o'_1 \geq o'_2$$

Then, \succsim *cannot* be represented by a utility function.

Proof: Assume such a utility function “ u ” exists. Then, for all positive $r \in \mathbb{R}$, it holds $(r, 2) \succsim (r, 1)$ iff $u(r, 2) > u(r, 1)$, and there exists a rational number $q \in \mathbb{Q}$ such that

$$u(r, 2) > q_r > u(r, 1)$$

Note that there exists always a rational number between any two real numbers. Now take two real numbers r and r' such that $r > r'$. We have $(r, 1) \succsim (r', 2)$ iff $u(r, 1) > u(r', 2)$ and therefore

$$u(r, 2) > q_r > u(r, 1) > q_{r'} > u(r', 2)$$

This means that if $r \neq r'$ then $q_r \neq q_{r'}$. Moreover, it is always the case that if $q_r \neq q_{r'}$ then $r \neq r'$. Together these two facts imply the existence of a one to one mapping between \mathbb{R} and \mathbb{Q} (a bijection between \mathbb{R} and \mathbb{Q}). However, such a bijection does not exist.

What is Game Theory Trying to Accomplish?

Point of Departure: Game theory as *interactive decision theory*.

Issue: *Assume many agents operating in the same environment each faced with a different optimization problem.*



What is Game Theory Trying to Accomplish?

Point of Departure: Game theory as *interactive decision theory*.

Issue: *Assume many agents operating in the same environment each faced with a different optimization problem.*



Observation I: Dependence of the outcome on *all* the players' actions

Hence, the optimality of an action depends on the optimality of the other players' actions. As this holds for all players, a circularity threatens.

What is Game Theory Trying to Accomplish?

Point of Departure: Game theory as *interactive* decision theory.

Issue: *Assume many agents operating in the same environment each faced with a different optimization problem.*



Observation I: Dependence of the outcome on *all* the players' actions

Hence, the optimality of an action depends on the optimality of the other players' actions. As this holds for all players, a circularity threatens.

Observation II: Yet, the other players' decisions cannot be considered parameters of the environment in an obvious way.

What is Game Theory Trying to Accomplish?

Point of Departure: Game theory as *interactive* decision theory.

Issue: Assume many agents operating in the same environment each faced with a different optimization problem.



Observation I: Dependence of the outcome on *all* the players' actions

Hence, the optimality of an action depends on the optimality of the other players' actions. As this holds for all players, a circularity threatens.

Observation II: Yet, the other players' decisions cannot be considered parameters of the environment in an obvious way.

Conclusion: New mathematical concepts required to take over the role of the optimum, *solution concepts*.

Nobel Prizes for Game Theory



1972

Arrow

Welfare theory

1978

Simon

Decision making

1994

Nash, Harsanyi, Selten

Equilibria

1996

Vickrey

Incentives

1998

Sen

Welfare economics

2005

Aumann and Schelling

Conflict and cooperation

2007

Hurwicz, Maskin and Myerson

Mechanism design



Games and Game Forms: The Strategic Form

- ▶ Players: Who is involved?
- ▶ Rules: What can the players do? What do they know when they act?
- ▶ Outcomes: What will happen when the players act in a particular way?
- ▶ Preferences: What are the players' preferences over the possible outcomes?

	<i>Fight</i>	<i>Ballet</i>
<i>Fight</i>	2, 1	0, 0
<i>Ballet</i>	0, 0	1, 2

Exercise: Describe Defence-Attack, Battle of the Sexes, and Game of Chicken games as strategic games.

Games and Game Forms: The Strategic Form

Definition: A *game form* is a quadruple (N, A, O, g) where:

- ▶ N is a set of n players
- ▶ $A = A_1 \times \cdots \times A_n$, an n -dimensional space of *strategy profiles* (action profiles), where A_i denote the set of *strategies* (actions) of player i
- ▶ O is a set of *outcomes*
- ▶ $g: A \rightarrow O$ is an *outcome function*

Games and Game Forms: The Strategic Form

Definition: A *game form* is a quadruple (N, A, O, g) where:

- ▶ N is a set of n players
- ▶ $A = A_1 \times \cdots \times A_n$, an n -dimensional space of *strategy profiles* (action profiles), where A_i denote the set of *strategies* (actions) of player i
- ▶ O is a set of *outcomes*
- ▶ $g: A \rightarrow O$ is an *outcome function*

Definition: A *strategic game* is a quintuple (N, A, O, g, u) , where:

- ▶ (N, A, O, g) is a game form
- ▶ $u = (u_1, \dots, u_n)$, where $u_i: A \rightarrow \mathbb{R}$ is a *utility function* for player i .

Games and Game Forms: The Strategic Form

Definition: A *game form* is a quadruple (N, A, O, g) where:

- ▶ N is a set of n players
- ▶ $A = A_1 \times \cdots \times A_n$, an n -dimensional space of *strategy profiles* (action profiles), where A_i denote the set of *strategies* (actions) of player i
- ▶ O is a set of *outcomes*
- ▶ $g: A \rightarrow O$ is an *outcome function*

Definition: A *strategic game* is a quintuple (N, A, O, g, u) , where:

- ▶ (N, A, O, g) is a game form
- ▶ $u = (u_1, \dots, u_n)$, where $u_i: A \rightarrow \mathbb{R}$ is a *utility function* for player i .

We write (N, A, u) instead of (N, A, O, g, u) by assuming $A = O$.

Exercise: Describe Defence-Attack, Battle of the Sexes, and Game of Chicken games as strategic games.

Preferences and Utilities in Games

Let O be a set of outcomes.

- ▶ $\succsim_i \subseteq O \times O$, reflexive, transitive and complete
- ▶ $\succsim = (\succsim_1, \dots, \succsim_n)$

Notations:

- ▶ $o_1 \sim_i o_2$ if both $o_1 \succsim_i o_2$ and $o_2 \succsim_i o_1$
- ▶ $o_1 \succ_i o_2$ if both $o_1 \succsim_i o_2$ and not $o_2 \succsim_i o_1$
- ▶ $a \succsim_i a'$ if $g(a) \succsim_i g(a')$

Definition: A utility function $u_i: O \rightarrow \mathbb{R}$ represents preferences \succsim_i over outcomes O so that:

$$u_i(o_1) \geq u_i(o_2) \quad \text{iff} \quad o_1 \succsim_i o_2$$

Notation: For a strategy in a game, we write $u_i(a) = u_i(g(a))$.

Security Level

Notation: Let $A = A_1 \times \dots \times A_n$ be the strategy space of n players, where A_i is the set of strategies of player i . We use (s_i, \mathbf{s}_{-i}) to denote a strategy profile $(s_1, \dots, s_i, \dots, s_n) \in A$.

Definition: The *pure security level* of player i is the least payoff he can guarantee himself, no matter what strategies the other players play, i.e.:

$$\max_{s_i} \min_{\mathbf{s}_{-i}} (u_i(s_i, \mathbf{s}_{-i}))$$

0, 4	3, 1
1, 2	2, 3

1, 0	0, 1
0, 1	1, 0

Exercise: Determine the security level of Defence-Attack, Battle of the Sexes, and Game of Chicken games.