Game-Theoretic View to Multiagent Systems

(Based on Shoham and Leyton-Brown (2008). *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge.)

Leen-Kiat Soh

Introduction

- In single-agent decision theory the key notion is that of an *optimal strategy.*
- In **noncooperative game theory**, the basic modeling unit is the ***individual*** (including its beliefs, preferences, and possible actions)
  - in **coalitional game theory**, it is the ***group***
- Agents are ***self-interested***
  - Each agent has its own description of which states of the world it likes—and that it acts in an attempt to bring about these states of the world.
- The dominant approach to modeling an agent's interests is ***utility theory***
- The idea of **(*expected*) utility** can be grounded in a more basic concept of preferences (and *rationality*): *von Neumann and Morgenstern*

Why is "self-interestedness" crucial?

Utility

- How valuable is $10 to you?
- $100?
- $10,000?
- $100,000?
- $101,000?
- $102,000?
- …

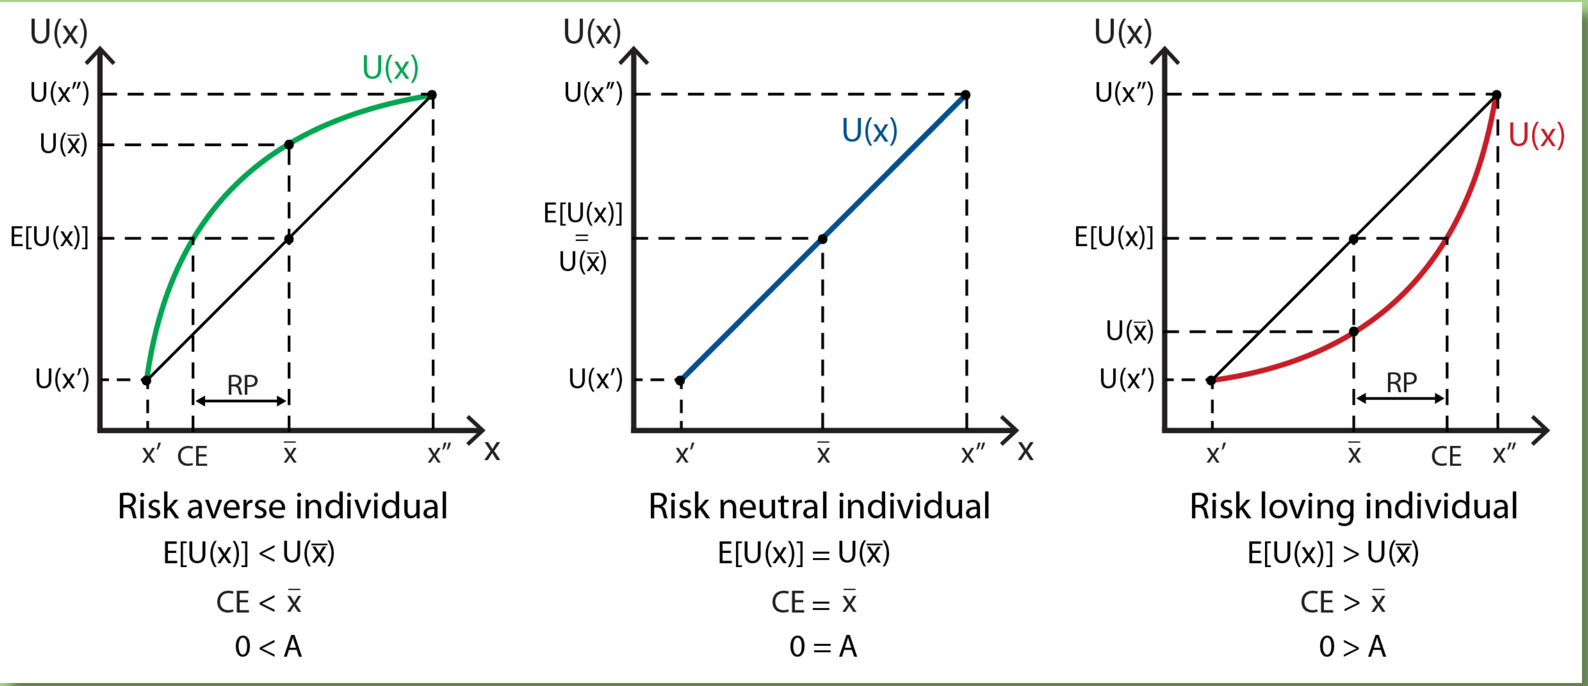

http://riskwizards.com/topic-under-discussion-risk-and-risk-aversion/

Expected Utility

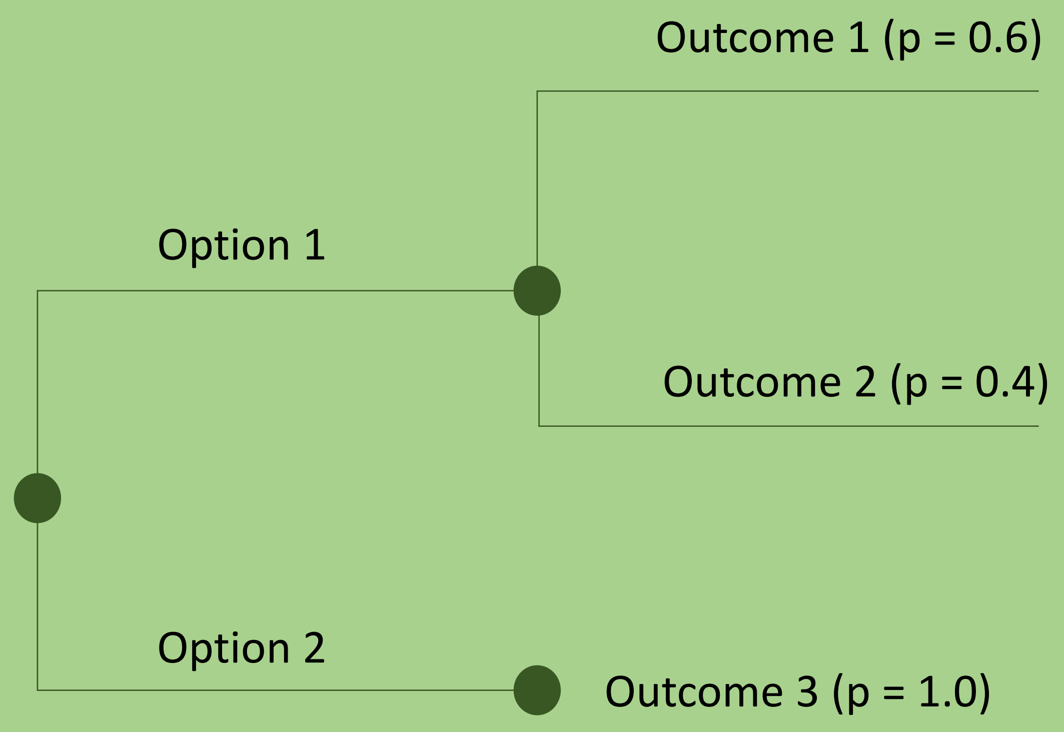

Outcome 1 = $1,000
Outcome 2 = $100
Outcome 3 = $640

Which option would you choose?

- What are $U_{Name}(\$1{,}000)$, $U_{Name}(\$640)$, $U_{Name}(\$100)$?
- What's the expected utility of Option 1? Option 2?

$$EU_{Name}(\text{Option1}) = 0.6 * U_{Name}(\$1{,}000) + 0.4 * U_{Name}(\$100)$$

$$EU_{Name}(\text{Option2}) = 1.0 * U_{Name}(\$640)$$

Rationality, risk attitudes

Utility Theory

- Let O denote a finite set of outcomes. For any pair $o_1, o_2 \in O$,
  - $o_1 \succcurlyeq o_2$: the agent weakly prefers o_1 to o_2
  - $o_1 \sim o_2$: the agent is indifferent between o_1 and o_2
  - $o_1 \succ o_2$: the agent strictly prefers o_1 to o_2.
- Now, a ***lottery*** is the random selection of one of a set of outcomes according to specified probabilities.
- Formally, a lottery is a ***probability distribution over outcomes*** written $[p_1 : o_1, \ldots, p_k : o_k]$, where each $o_i \in O$, each $p_i \geq 0$ and $\sum_{i=1}^{k} p_i = 1$.
- Let L denote the set of all lotteries
  - Extend the $\succcurlyeq$ relation to apply to the elements of L as well as to the elements of O, **effectively considering lotteries over outcomes to be outcomes themselves**

Utility Theory | Axioms

- **Axiom 3.1.1. Completeness.** $\forall o_1, o_2 \;\; o_1 \succ o_2 \; or \; o_2 \succ o_1 \; or \; o_1 \sim o_2.$
- **Axiom 3.1.2. Transitivity.** $If \; o_1 \succcurlyeq o_2 \; and \; o_2 \succcurlyeq o_3, then \; o_1 \succcurlyeq o_3.$
- **Axiom 3.1.3. Substitutability.** *If* $o_1 \sim o_2$, *then for all sequences of one or more outcomes* $o_3, \ldots, o_k$ *and sets of probabilities* $p, p_3, \ldots, p_k$ *for which* $p + \sum_{i=3}^{k} p_i = 1$,

$$[p : o_1, p_3 : o_3, \ldots, p_k : o_k] \sim [p : o_2, p_3 : o_3, \ldots, p_k : o_k].$$

Completeness: no ambiguity, preference must be there, can be partially ordered
Transitivity: consistency, inference
Substitutability states that if an agent is indifferent between two outcomes, it is also indifferent between two lotteries that differ only in which of these outcomes is offered.

Utility Theory | Axioms 2

- Let $P_\ell(o_i)$ denote the probability that outcome o_i is selected by lottery ℓ
- **Axiom 3.1.4 Decomposability.** *If* $\forall o_i \in O, P_{\ell_1}(o_i) = P_{\ell_2}(o_i)$, *then* $\ell_1 \sim \ell_2$.
  - **Decomposability states that an agent is always indifferent between lotteries that induce the same probabilities over outcomes**, no matter whether these probabilities are expressed through a single lottery or nested in a lottery over lotteries
- **Axiom 3.1.5 Monotonicity.** *If* $o_1 \succ o_2$ *and* $p > q$, *then* $[p: o_1, 1-p: o_2] \succ [q: o_1, 1-q: o_2]$.

Decomposability: Easier to compare lotteries
Monotonicity: consistent, dependable

Utility Theory | Axioms 3

- **Lemma 3.1.6** *If a preference relation $\succcurlyeq$ satisfies the axioms completeness, transitivity, decomposability, and monotonicity, and if* $o_1 \succ o_2$ *and* $o_2 \succ o_3$ *, then there exists some probability* p *such that*

$$\text{for all } p' < p,\ o_2 \succ [p' : o_1, (1 - p') : o_3], \text{ and}$$

$$\text{for all } p'' > p, [p'' : o_1, (1 - p'') : o_3] \succ o_2$$

- **Axiom 3.1.7 Continuity.** *If* $o_1 \succ o_2$ *and* $o_2 \succ o_3$*, then* $\exists p \in [0, 1]$ *such that* $o_2 \sim [p : o_1, 1 - p : o_3]$.

Utility Theory | Axioms 4

- If we accept Axioms 3.1.1, 3.1.2, 3.1.4, 3.1.5, and 3.1.7, it turns out that we have no choice but to accept the **existence of single-dimensional utility functions** whose expected values agents want to maximize
  - And if we do *not* want to reach this conclusion, we must then give up at least one of the axioms
- **Theorem 3.1.8 (von Neumann and Morgenstern, 1944)** *If a preference relation* $\succcurlyeq$ *satisfies the axioms completeness, transitivity, substitutability, decomposability, monotonicity, and continuity, then there exists a function* $u: O \rightarrow [0, 1]$ *with the properties that*

$$1.\ u(o_1) \geq u(o_2) \text{ iff } o_1 \succcurlyeq o_2, \text{ and}$$

$$2.\ u([p_1: o_1, \ldots, p_k: o_k]) = \sum_{i=1}^{k} p_i u(o_i)$$

Normal Form Games

- **a.k.a. Strategic Form Games**
- **Definition 3.2.1 (Normal-form game)** *A (finite, n-person) normal-form game is a tuple* (N, A, u)*, where:*
  - N *is a finite set of* n *players, indexed by* i
  - $A = A_1 \times \cdots \times A_n$*, where* A_i *is a finite set of* actions *available to player* i
    - *Each vector* $a = (a_1, \ldots, a_n) \in A$ *is called an* action profile
  - *u =* $(u_1, \ldots, u_n)$ *where* $u_i: A_i \to \mathbb{R}$ *is a real-valued* utility *(or* payoff*)* function *for player* i
- Note that we previously argued that utility functions should map from the set of *outcomes*, not the set of *actions*
  - Here we make the implicit assumption that $O = A$.
- A natural way to represent games is via an n-dimensional matrix.

How should an agent play a normal form game?

- We have seen that under reasonable assumptions about preferences, agents will *always* have utility function whose expected values they want to maximize
- *This suggests that acting optimally in an uncertain environment is conceptually straightforward:*
- **Agents simply need to choose the course of action that maximizes expected utility!**
- *But in real life, that's often too good to be true*

Normal Form Games | Prisoner’s Dilemma

	Player 2 No Betray	Player 2 Betray
Player 1 No Betray	**1,1**	-4,3
Player 1 Betray	3,-4	**-3,-3**

Best joint outcomes

Worst joint outcomes

Dominant strategy, Nash equilibrium

- Any rational user, when presented with this scenario once, will adopt Betray—regardless of what the other user does
  - Allowing the users to communicate beforehand will not change the outcome
- Perfectly rational agents will also adopt Betray even if they play multiple times
  - **However, if the number of times that is infinite, or uncertain, agents may adopt No Betray**

Normal Form Games | Common Payoff Game

- **Definition 3.2.2 (Common-payoff game)** *A* common-payoff game *is a game in which for all action profiles* $a \in A_1 \times \cdots \times A_n$ *and any pair of agents* i*, j, it is the case that* $u_i(a) = u_j(a)$.
- a.k.a. ***pure coordination games*** or ***team games***
  - no conflicting interests among agents; only need to coordinate to maximize benefits

E.g., two people walking towards each other in a hallway: If they choose the same side (left or right) they have some high utility; otherwise, a low utility.

	Agent 2 Left	Agent 2 Right
Agent 1 Left	1,1	0,0
Agent 1 Right	0,0	1,1

Why considering common payoff game? Agents need to understand on their own that they need to coordinate

Normal Form Games | Constant-Sum Game

- **Definition 3.2.3 (Constant-sum game)** *A two-player normal-form game is* constant sum *if there exists a constant c such that for each strategy profile* $a \in A_1 \times A_2$*, it is the case that* $u_1(a) + u_2(a) = c$.
- E.g., zero-sum game when c = 0
- Opposite of pure coordination games

	P2 Heads	P2 Tails
P1 Heads	1,-1	-1,1
P2 Tails	-1,1	1,-1

Paper-scissors-rock, constrained resource allocation (or assignment)

E.g., *Matching Pennies*. In this game, each of the two players has a penny and independently chooses to display either heads or tails. P1 pockets both if the pennies match; otherwise P2 pockets both

Connection to MAS?

Utility, expected utility, rationality?

Action selection, maximize utility, single agent decision making vs. multiple agents making decisions: may not be able to secure “optimal” strategy

Stupid Question: What if parking meters in a city allows for members of “car-pooling groups” to use left-over amount of money from parking meters, that would expire at the end of the day? What would your strategy be if you were a member of such a group?