Utility and Foundations (2)

Modeling Human Strategic Behavior

Kevin Leyton-Brown

University of British Columbia
Canada CIFAR AI Chair, Amii

James R. Wright

University of Alberta
Canada CIFAR AI Chair, Amii

Lecture Overview

Recap

Proof sketch

Fun Game!

Recap: Axioms

- Completeness

$$o_1 \succeq o_2 \text{ or } o_2 \succeq o_1$$

Recap: Axioms

- Completeness

$$o_1 \succeq o_2 \text{ or } o_2 \succeq o_1$$

- Transitivity

$$(o_1 \succeq o_2) \wedge (o_2 \succeq o_3) \implies o_1 \succeq o_3$$

Recap: Axioms

- Completeness

$$o_1 \succeq o_2 \text{ or } o_2 \succeq o_1$$

- Transitivity

$$(o_1 \succeq o_2) \wedge (o_2 \succeq o_3) \implies o_1 \succeq o_3$$

- Monotonicity

$$p > q \implies [p : good,\ (1-p) : bad] \succ [q : good,\ (1-q) : bad]$$

Recap: Axioms

- Completeness

$$o_1 \succeq o_2 \text{ or } o_2 \succeq o_1$$

- Transitivity

$$(o_1 \succeq o_2) \wedge (o_2 \succeq o_3) \implies o_1 \succeq o_3$$

- Monotonicity

$$p > q \implies [p : good,\ (1-p) : bad] \succ [q : good,\ (1-q) : bad]$$

- Substitutability

$$o_1 \sim o_2 \implies \text{Can replace } o_1 \text{ with } o_2$$

Recap: Axioms

- Completeness

$$o_1 \succeq o_2 \text{ or } o_2 \succeq o_1$$

- Transitivity

$$(o_1 \succeq o_2) \wedge (o_2 \succeq o_3) \implies o_1 \succeq o_3$$

- Monotonicity

$$p > q \implies [p : good,\ (1-p) : bad] \succ [q : good,\ (1-q) : bad]$$

- Substitutability

$$o_1 \sim o_2 \implies \text{Can replace } o_1 \text{ with } o_2$$

- Decomposability

$$P_{\ell_1}(o) = P_{\ell_2}(o)) \implies \ell_1 \sim \ell_2$$

Recap: Axioms

- Completeness

$$o_1 \succeq o_2 \text{ or } o_2 \succeq o_1$$

- Transitivity

$$(o_1 \succeq o_2) \wedge (o_2 \succeq o_3) \implies o_1 \succeq o_3$$

- Monotonicity

$$p > q \implies [p : good,\ (1-p) : bad] \succ [q : good,\ (1-q) : bad]$$

- Substitutability

$$o_1 \sim o_2 \implies \text{Can replace } o_1 \text{ with } o_2$$

- Decomposability

$$P_{\ell_1}(o) = P_{\ell_2}(o)) \implies \ell_1 \sim \ell_2$$

- Continuity

$$o_1 \succ o_2 \succ o_3 \implies \exists p \in [0,1] : o_2 \sim [p : o_1,\ (1-p) : o_3]$$

Recap: Representation Theorem

Theorem [von Neumann & Morgenstern, 1944]

Suppose that a preference relation $\succeq$ satisfies the axioms Completeness, Transitivity, Monotonicity, Substitutability, Decomposability, and Continuity.

Then there exists a function $u : O \to \mathbb{R}$ such that

1. $\forall o_1, o_2 \in O : o_1 \succeq o_2 \iff u(o_1) \geq u(o_2)$, and
2. $\forall [p_1 : o_1, \ \ldots, p_k : o_k] \in O : u([p_1 : o_1, \ \ldots, p_k : o_k]) = \sum_{j=1}^{k} p_j u(o_j)$.

That is, there exists a utility function u that represents $\succeq$.

Lecture Overview

Recap

Proof sketch

Fun Game!

Proof sketch

1. Choose o^+, o^- such that $o^- \preceq o \preceq o^+$ for all o

Proof sketch

1. Choose o^+, o^- such that $o^- \preceq o \preceq o^+$ for all o
   - (this turns out to be without loss of generality)

Proof sketch

1. Choose o^+, o^- such that $o^- \preceq o \preceq o^+$ for all o
   - (this turns out to be without loss of generality)
2. Construct $u(o) = p$ such that $o \sim [p : o^+,\ (1-p) : o^-]$

Proof sketch

1. Choose o^+, o^- such that $o^- \preceq o \preceq o^+$ for all o
   - (this turns out to be without loss of generality)
2. Construct $u(o) = p$ such that $o \sim [p : o^+,\ (1-p) : o^-]$
3. Substitutability lets us replace everything with these “canonical” lotteries; Monotonicity lets us assert the ordering between them.

Caveats & Details: Uniqueness

For a given set of preferences, the utility function is **not uniquely defined**.

Caveats & Details: Uniqueness

For a given set of preferences, the utility function is **not uniquely defined**.

Comparisons of expected values are invariant to **positive affine transformations**:

$$X \succeq Y \iff \mathbb{E}[u(X)] \geq \mathbb{E}[u(Y)]$$

Caveats & Details: Uniqueness

For a given set of preferences, the utility function is **not uniquely defined**.

Comparisons of expected values are invariant to **positive affine transformations**:

$$\begin{aligned} X \succeq Y &\iff \mathbb{E}[u(X)] \geq \mathbb{E}[u(Y)] \\ &\iff c\,\mathbb{E}[u(X)] \geq c\,\mathbb{E}[u(Y)] \end{aligned}$$

Caveats & Details: Uniqueness

For a given set of preferences, the utility function is **not uniquely defined**.

Comparisons of expected values are invariant to **positive affine transformations**:

$$\begin{aligned} X \succeq Y &\iff \mathbb{E}[u(X)] \geq \mathbb{E}[u(Y)] \\ &\iff c\,\mathbb{E}[u(X)] \geq c\,\mathbb{E}[u(Y)] \\ &\iff c\,\mathbb{E}[u(X)]+b \geq c\,\mathbb{E}[u(Y)]+b \end{aligned}$$

Caveats & Details: Uniqueness

For a given set of preferences, the utility function is **not uniquely defined**.

Comparisons of expected values are invariant to **positive affine transformations**:

$$\begin{aligned} X \succeq Y \iff & \mathbb{E}[u(X)] \geq \mathbb{E}[u(Y)] \\ \iff & c\,\mathbb{E}[u(X)] \geq c\,\mathbb{E}[u(Y)] \\ \iff & c\,\mathbb{E}[u(X)]+b \geq c\,\mathbb{E}[u(Y)]+b \\ \iff & \mathbb{E}[cu(X)+b] \geq \mathbb{E}[cu(Y)+b] \end{aligned}$$

for all $b \in \mathbb{R}$ and $c > 0$

Lecture Overview

Recap

Proof sketch

Fun Game!

Fun Game: Buying Lottery Tickets

Write down the following numbers:

1. How much would you pay to play the lottery

$$[0.3 : \$5,\ 0.3 : \$7,\ 0.4 : \$9]?$$

Fun Game: Buying Lottery Tickets

Write down the following numbers:

1. How much would you pay to play the lottery

$$[0.3 : \$5,\ 0.3 : \$7,\ 0.4 : \$9]?$$

2. How much would you pay to play the lottery

$$[p : \$5,\ q : \$7,\ (1 - p - q) : \$9]?$$

Fun Game: Buying Lottery Tickets

Write down the following numbers:

1. How much would you pay to play the lottery

$$[0.3 : \$5,\ 0.3 : \$7,\ 0.4 : \$9]?$$

2. How much would you pay to play the lottery

$$[p : \$5,\ q : \$7,\ (1 - p - q) : \$9]?$$

3. How much would you pay to play the lottery

$$[p : \$5,\ q : \$7,\ (1 - p - q) : \$9]$$

If you knew that the last seven draws had been $5, 5, 7, 5, 9, 9, 5$*?*

Beyond von Neumann & Morgenstern

- The first game was a pretty good match for the utility theory that we just learned.
- **Question:** If two rational agents have different prices for $[0.3 : \$5,\ 0.3 : \$7,\ 0.4 : \$9]$, what does that suggest about their **preferences for money**?

Beyond von Neumann & Morgenstern

- The first game was a pretty good match for the utility theory that we just learned.
- **Question:** If two rational agents have different prices for $[0.3:\$5,\ 0.3:\$7,\ 0.4:\$9]$, what does that suggest about their **preferences for money**?
- The second game was not such a great match!
- **Question:** If two rational agents have different prices for $[p:\$5,\ q:\$7,\ (1-p-q):\$9]$, can we infer anything about the two agents' **preferences for money**?

Beyond von Neumann & Morgenstern

- The first game was a pretty good match for the utility theory that we just learned.
- **Question:** If two rational agents have different prices for $[0.3:\$5,\ 0.3:\$7,\ 0.4:\$9]$, what does that suggest about their **preferences for money**?
- The second game was not such a great match!
- **Question:** If two rational agents have different prices for $[p:\$5,\ q:\$7,\ (1-p-q):\$9]$, can we infer anything about the two agents' **preferences for money**?
- If the two agents agree about the price for $[p:\$5,\ q:\$7,\ (1-p-q):\$9]$ but then disagree once they hear what the last few draws were?

Beyond von Neumann & Morgenstern

- The first game was a pretty good match for the utility theory that we just learned.
- **Question:** If two rational agents have different prices for $[0.3 : \$5,\ 0.3 : \$7,\ 0.4 : \$9]$, what does that suggest about their **preferences for money**?
- The second game was not such a great match!
- **Question:** If two rational agents have different prices for $[p : \$5,\ q : \$7,\ (1-p-q) : \$9]$, can we infer anything about the two agents' **preferences for money**?
- If the two agents agree about the price for $[p : \$5,\ q : \$7,\ (1-p-q) : \$9]$ but then disagree once they hear what the last few draws were?
- von Neumann and Morgenstern's utility theory assumes **known, objective** probabilities.
- There are other representation theorems [e.g., Savage 1954] that state that rational agents must (a) have probabilistic beliefs, (b) update those beliefs as if by conditioning, (c) maximize the expected value of some utility function wrt them

Utility Summary

Utility theory proves that agents whose preferences obey certain simple axioms about preferences over lotteries must act as if they were maximizing the expected value of a scalar function.

- **“Rational”** agents are those whose behaviour satisfies the **axioms**
- ***If you don’t buy the axioms, then you shouldn’t buy that this theorem is about rational behavior.***
- Conversely, if you don’t buy that rational agents must behave in this way, then there must be at least one axiom that you disagree with.

This approach extends to “subjective” probabilities:

- Axioms about **preferences over uncertain “acts”** that do not describe how agents manipulate probabilities.

Game Representations

Kevin Leyton-Brown

University of British Columbia
Canada CIFAR AI Chair, Amii

James R. Wright

University of Alberta
Canada CIFAR AI Chair, Amii

Lecture Overview

Normal-Form

Repeated

Extensive Form

Bayesian Games

TCP Backoff Game

Should you send your packets using correctly-implemented TCP (which has a "backoff" mechanism) or using a defective implementation (which doesn't)?

- Consider this situation as a two-player game:
  - **both use a correct implementation:** both get 1 ms delay
  - **one correct, one defective:** 4 ms delay for correct, 0 ms for defective
  - **both defective:** both get a 3 ms delay.

TCP Backoff Game

- Consider this situation as a two-player game:
  - **both use a correct implementation:** both get 1 ms delay
  - **one correct, one defective:** 4 ms delay for correct, 0 ms for defective
  - **both defective:** both get a 3 ms delay.
- Go into a breakout room. Play once with each person.
- Questions:
  - What **action** should a player of the game take?
  - Would all users behave **the same** in this scenario?
  - What global **patterns of behaviour** should the system designer expect?
  - Under what **changes to the delay numbers** would behavior be the same?
  - What effect would **communication** have?
  - Does it matter if I believe that my opponent is **rational**?

Defining Games

- Finite, n-person game: $\langle N, A, u \rangle$:
  - N is a finite set of n **players**, indexed by i
  - $A = \langle A_1, \ldots, A_n \rangle$ is a tuple of **action sets** for each player i
    - $a \in A$ is an **action profile**
  - $u = \langle u_1, \ldots, u_n \rangle$, a **utility function** for each player, where $u_i : A \mapsto \mathbb{R}$

- Writing a 2-player game as a **matrix**:
  - row player is player 1, column player is player 2
  - rows are actions $a \in A_1$, columns are $a' \in A_2$
  - cells are outcomes, written as a tuple of utility values for each player

Games in Matrix Form

Here's the **TCP Backoff Game** written as a matrix ("normal form").

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

More General Form

Prisoner's dilemma is any game

	C	D
C	a, a	b, c
D	c, b	d, d

with $c > a > d > b$.

Games of Pure Competition

Players have **exactly opposed** interests

- There must be precisely two players *(otherwise they can't have exactly opposed interests)*
- For all action profiles $a \in A$, $u_1(a) + u_2(a) = c$ for some constant c
  - Special case: zero sum
- Thus, we only need to store a utility function for one player
  - in a sense, it's a one-player game

Matching Pennies

One player wants to **match**; the other wants to **mismatch**.

	Heads	Tails
Heads	$1, -1$	$-1, 1$
Tails	$-1, 1$	$1, -1$

Rock-Paper-Scissors

Generalized matching pennies.

	Rock	Paper	Scissors
Rock	$0, 0$	$-1, 1$	$1, -1$
Paper	$1, -1$	$0, 0$	$-1, 1$
Scissors	$-1, 1$	$1, -1$	$0, 0$

...Believe it or not, there's an annual international competition for this game!

Games of Cooperation

Players have **exactly the same** interests.

- no conflict: all players want the same things
- $\forall a \in A, \forall i, j,\ u_i(a) = u_j(a)$
- we often write such games with a single payoff per cell
- why are these even still games?

Coordination Game

Which **side of the road** should you drive on?

	Left	Right
Left	$1,1$	$0,0$
Right	$0,0$	$1,1$

General Games: Battle of the Sexes

The most interesting games combine elements of **cooperation *and* competition**.

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

General Games: Battle of the Sexes

The most interesting games combine elements of **cooperation *and* competition**.

	B	F
B	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

Play this game in breakout rooms. Be fast!