Utility Theory

James Wright

January 16/2014

Outline

Overview

Utility, informally

- A utility function is a real-valued function that indicates how much agents like an outcome.
- In the presence of uncertainty, rational agents act to maximize their expected utility.
- Utility is a foundational concept in game theory.

Overview

Utility, informally

- A utility function is a real-valued function that indicates how much agents like an outcome.
- In the presence of uncertainty, rational agents act to maximize their expected utility.
- Utility is a foundational concept in game theory.
- But it is a nontrivial claim:
  1. Why should we believe that an agent's preferences can be adequately represented by a single number?
  2. Why should agents maximize expectations rather than some other criterion?

Overview

Utility, informally

- A utility function is a real-valued function that indicates how much agents like an outcome.
- In the presence of uncertainty, rational agents act to maximize their expected utility.
- Utility is a foundational concept in game theory.
- But it is a nontrivial claim:
  1. Why should we believe that an agent's preferences can be adequately represented by a single number?
  2. Why should agents maximize expectations rather than some other criterion?
- Von Neumann and Morgenstern's theorem shows why (and when!) these are true.
- It is also a good example of some common elements in game theory (and economics):
  - Behaving "as-if"
  - Axiomatic characterization

Outline

Formal setting

Definition

Let O be a set of possible outcomes. A lottery is a probability distribution over outcomes. Write $[p_1 : o_1, p_2 : o_2, \ldots, p_k : o_k]$ for the lottery that assigns probability p_1 to outcome o_1, etc.

Formal setting

Definition

Let O be a set of possible outcomes. A lottery is a probability distribution over outcomes. Write $[p_1 : o_1, p_2 : o_2, \ldots, p_k : o_k]$ for the lottery that assigns probability p_1 to outcome o_1, etc.

Definition

For a specific preference relation $\succeq$, write:

1. $o_1 \succeq o_2$ if the agent weakly prefers o_1 to o_2;
2. $o_1 \succ o_2$ if the agent strictly prefers o_1 to o_2; and
3. $o_1 \sim o_2$ if the agent is indifferent between o_1 and o_2.

Formal setting

Definition

Let O be a set of possible outcomes. A lottery is a probability distribution over outcomes. Write $[p_1 : o_1, p_2 : o_2, \ldots, p_k : o_k]$ for the lottery that assigns probability p_1 to outcome o_1, etc.

Definition

For a specific preference relation $\succeq$, write:

1. $o_1 \succeq o_2$ if the agent weakly prefers o_1 to o_2;
2. $o_1 \succ o_2$ if the agent strictly prefers o_1 to o_2; and
3. $o_1 \sim o_2$ if the agent is indifferent between o_1 and o_2.

Definition

A utility function is a function $u : O \to \mathbb{R}$. A utility function represents a set of preferences if:

1. $o_1 \succeq o_2 \iff u(o_1) \geq u(o_2)$; and
2. $u([p_1 : o_1, \ldots, p_k : o_k]) = \sum_{i=1}^{k} p_i u(o_i)$.

Representation theorem

Von Neumann and Morgenstern, 1944

Theorem

Suppose a preference relation $\succeq$ satisfies the axioms Completeness, Transitivity, Monotonicity, Substitutability, Decomposability, and Continuity. Then there exists a function $u : O \to [0, 1]$ such that

1. $o_1 \succeq o_2 \iff u(o_1) \geq u(o_2)$; *and*
2. $u([p_1 : o_1, \ldots, p_k : o_k]) = \sum_{i=1}^{k} p_i u(o_i)$.

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

Outline

Completeness and transitivity

Definition (Completeness)

$$\forall o_1, o_2 : o_1 \succ o_2 \text{ or } o_2 \succ o_1 \text{ or } o_1 \sim o_2.$$

Completeness and transitivity

Definition (Completeness)

$$\forall o_1, o_2 : o_1 \succ o_2 \text{ or } o_2 \succ o_1 \text{ or } o_1 \sim o_2.$$

Definition (Transitivity)

$$o_1 \succeq o_2 \text{ and } o_2 \succeq o_3 \implies o_1 \succeq o_3.$$

Completeness and transitivity

Definition (Completeness)

$$\forall o_1, o_2 : o_1 \succ o_2 \text{ or } o_2 \succ o_1 \text{ or } o_1 \sim o_2.$$

Definition (Transitivity)

$$o_1 \succeq o_2 \text{ and } o_2 \succeq o_3 \implies o_1 \succeq o_3.$$

Money pump justification.

- Suppose that $o_1 \succ o_2$ and $o_2 \succ o_3$ and $o_3 \succ o_1$.
- Starting from o_3, you should be willing to pay 1 cent (say) to switch to o_2.
- But from o_2 you should be willing to pay 1 cent to switch to o_1.
- But from o_1 you should be willing to pay 1 cent to switch back to o_3 ...

Monotonicity

Definition (Monotonicity)

If $o_1 \succ o_2$ and $p > q$, then
$[p : o_1, (1-p) : o_2] \succ [q : o_1, (1-q) : o_2]$.

Monotonicity

Definition (Monotonicity)

If $o_1 \succ o_2$ and $p > q$, then
$[p : o_1, (1-p) : o_2] \succ [q : o_1, (1-q) : o_2]$.

You should prefer a 90% chance of getting \$1000 to a 50% chance of getting \$10.

Substitutability

Definition (Substitutability)

If $o_1 \sim o_2$, then for all sequences $o_3, \ldots, o_k$ and $p, p_3, \ldots, p_k$ with $p + \sum_{i=3}^{k} p_i = 1$,

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

Substitutability

Definition (Substitutability)

If $o_1 \sim o_2$, then for all sequences $o_3, \ldots, o_k$ and $p, p_3, \ldots, p_k$ with $p + \sum_{i=3}^{k} p_i = 1$,

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

If I like apples and bananas equally, then I should be indifferent between a 30% chance of getting a banana or a 30% chance of getting an apple.

Decomposability

Definition (Decomposability)

Let $P_\ell(o_i)$ denote the probability that lottery ℓ selects outcome o_i.
If $P_{\ell_1}(o_i) = P_{\ell_2}(o_i)\ \forall o_i \in O$, then $\ell_1 \sim \ell_2$.

Decomposability

Definition (Decomposability)

Let $P_\ell(o_i)$ denote the probability that lottery ℓ selects outcome o_i.
If $P_{\ell_1}(o_i) = P_{\ell_2}(o_i)\ \forall o_i \in O$, then $\ell_1 \sim \ell_2$.

Example.
Let $\ell_1 = [0.5 : [0.5 : o_1, 0.5 : o_2], 0.5 : o_3]$.
Let $\ell_2 = [0.25 : o_1, 0.25 : o_2, 0.5 : o_3]$.

Decomposability

Definition (Decomposability)

Let $P_\ell(o_i)$ denote the probability that lottery ℓ selects outcome o_i.
If $P_{\ell_1}(o_i) = P_{\ell_2}(o_i)\ \forall o_i \in O$, then $\ell_1 \sim \ell_2$.

Example.
Let $\ell_1 = [0.5 : [0.5 : o_1, 0.5 : o_2], 0.5 : o_3]$.
Let $\ell_2 = [0.25 : o_1, 0.25 : o_2, 0.5 : o_3]$.

Then $\ell_1 \sim \ell_2$, because:
$P_{\ell_1}(o_1) = P_{\ell_2}(o_1) = 0.25$,
$P_{\ell_1}(o_2) = P_{\ell_2}(o_2) = 0.25$,
$P_{\ell_1}(o_3) = P_{\ell_2}(o_3) = 0.5$.

Continuity

Definition (Continuity)

If $o_1 \succ o_2 \succ o_3$, then $\exists p \in [0, 1]$ such that $o_2 \sim [p : o_1, (1 - p) : o_3]$.

Outline

Proof sketch

Construct the utility function

1. For $\succeq$ satisfying Completeness, Transitivity, Monotonicity, Decomposability and $o_1 \succ o_2 \succ o_3$,
   $\exists p$ such that:
   1. $o_2 \succ [q : o_1, (1-q) : o_3] \quad \forall q < p$, and
   2. $o_2 \prec [q : o_1, (1-q) : o_3] \quad \forall q > p$.

Proof sketch

Construct the utility function

1. For $\succeq$ satisfying Completeness, Transitivity, Monotonicity, Decomposability and $o_1 \succ o_2 \succ o_3$,
   $\exists p$ such that:
   1. $o_2 \succ [q : o_1, (1-q) : o_3] \quad \forall q < p$, and
   2. $o_2 \prec [q : o_1, (1-q) : o_3] \quad \forall q > p$.
2. For $\succeq$ additionally satisfying Continuity,
   $\exists p : o_2 \sim [p : o_1, (1-p) : o_3]$.

Proof sketch

Construct the utility function

1. For $\succeq$ satisfying Completeness, Transitivity, Monotonicity, Decomposability and $o_1 \succ o_2 \succ o_3$,
   $\exists p$ such that:
   1. $o_2 \succ [q : o_1, (1-q) : o_3] \quad \forall q < p$, and
   2. $o_2 \prec [q : o_1, (1-q) : o_3] \quad \forall q > p$.
2. For $\succeq$ additionally satisfying Continuity,
   $\exists p : o_2 \sim [p : o_1, (1-p) : o_3]$.
3. Choose maximal $\overline{o} \in O$ and minimal $\underline{o} \in O$.
4. Construct $u(o) = p$ such that $o \sim [p : \overline{o}, (1-p) : \underline{o}]$.

Proof sketch

Check the properties

1. $u(o_1) > u(o_2) \implies o_1 \succ o_2$:

Proof sketch

Check the properties

1. $u(o_1) > u(o_2) \implies o_1 \succ o_2$:
   - $u(o) = p$ such that $o \sim [p : \overline{o}, (1-p) : \underline{o}]$

Proof sketch

Check the properties

1. $u(o_1) > u(o_2) \implies o_1 \succ o_2$:
   - $u(o) = p$ such that $o \sim [p : \overline{o}, (1-p) : \underline{o}]$
2. $u([p_1 : o_1, \ldots, p_k : o_k]) = \sum_{i=1}^{k} p_i u(o_i)$:
   1. Let $u^* = u([p_1 : o_1, \ldots, p_k : o_k])$.

Proof sketch

Check the properties

1. $u(o_1) > u(o_2) \implies o_1 \succ o_2$:
   - $u(o) = p$ such that $o \sim [p : \overline{o}, (1-p) : \underline{o}]$
2. $u([p_1 : o_1, \ldots, p_k : o_k]) = \sum_{i=1}^{k} p_i u(o_i)$:
   1. Let $u^* = u([p_1 : o_1, \ldots, p_k : o_k])$.
   2. Replace o_i by ℓ_i, giving:
      $u^* = u([p_1 : [u(o_1) : \overline{o}, (1-u(o_1)) : \underline{o}], \ldots])$.

Proof sketch

Check the properties

1. $u(o_1) > u(o_2) \implies o_1 \succ o_2$:
   - $u(o) = p$ such that $o \sim [p : \overline{o}, (1-p) : \underline{o}]$
2. $u([p_1 : o_1, \ldots, p_k : o_k]) = \sum_{i=1}^{k} p_i u(o_i)$:
   1. Let $u^* = u([p_1 : o_1, \ldots, p_k : o_k])$.
   2. Replace o_i by ℓ_i, giving: $u^* = u([p_1 : [u(o_1) : \overline{o}, (1 - u(o_1)) : \underline{o}], \ldots])$.
   3. Question: What is the probability of getting $\overline{o}$?

Proof sketch

Check the properties

1. $u(o_1) > u(o_2) \implies o_1 \succ o_2$:
   - $u(o) = p$ such that $o \sim [p : \overline{o}, (1-p) : \underline{o}]$
2. $u([p_1 : o_1, \ldots, p_k : o_k]) = \sum_{i=1}^{k} p_i u(o_i)$:
   1. Let $u^* = u([p_1 : o_1, \ldots, p_k : o_k])$.
   2. Replace o_i by ℓ_i, giving:
      $u^* = u([p_1 : [u(o_1) : \overline{o}, (1 - u(o_1)) : \underline{o}], \ldots])$.
   3. Question: What is the probability of getting $\overline{o}$?
   4. Answer: $\sum_{i=1}^{k} p_i u(o_i)$

Proof sketch

Check the properties

1. $u(o_1) > u(o_2) \implies o_1 \succ o_2$:
   - $u(o) = p$ such that $o \sim [p : \overline{o}, (1-p) : \underline{o}]$
2. $u([p_1 : o_1, \ldots, p_k : o_k]) = \sum_{i=1}^{k} p_i u(o_i)$:
   1. Let $u^* = u([p_1 : o_1, \ldots, p_k : o_k])$.
   2. Replace o_i by ℓ_i, giving:
      $u^* = u([p_1 : [u(o_1) : \overline{o}, (1 - u(o_1)) : \underline{o}], \ldots])$.
   3. Question: What is the probability of getting $\overline{o}$?
   4. Answer: $\sum_{i=1}^{k} p_i u(o_i)$
   5. So $u^* = u\left(\left[\left(\sum_{i=1}^{k} p_i u(o_i)\right) : \overline{o}, \left(1 - \sum_{i=1}^{k} p_i u(o_i)\right) : \underline{o}\right]\right)$.
   6. By definition of u then,
      $u([p_1 : o_1, \ldots, p_k : o_k]) = \sum_{i=1}^{k} p_i u(o_i)$.

Outline

Fun game

Buying random dollars

Write down the following numbers:

Fun game

Buying random dollars

Write down the following numbers:

1. How much would you pay for a ticket in the lottery $[\frac{1}{3} : \$5, \frac{1}{3} : \$7, \frac{1}{3} : \$9]$?

Fun game

Buying random dollars

Write down the following numbers:

1. How much would you pay for a ticket in the lottery $[\frac{1}{3} : \$5, \frac{1}{3} : \$7, \frac{1}{3} : \$9]$?
2. How much would you pay for a ticket in the lottery $[p : \$5, q : \$7, (1 - p - q) : \$9]$?

Fun game

Buying random dollars

Write down the following numbers:

1. How much would you pay for a ticket in the lottery $[\frac{1}{3} : \$5, \frac{1}{3} : \$7, \frac{1}{3} : \$9]$?
2. How much would you pay for a ticket in the lottery $[p : \$5, q : \$7, (1 - p - q) : \$9]$?
3. How much would you pay for a ticket in the lottery $[p : \$5, q : \$7, (1 - p - q) : \$9]$ if you knew the last seven draws had been 5,5,7,5,9,9,5?

Outline

Beyond von Neumann-Morgenstern

- The first step of the fun game was a good match to the utility theory we just learned.
  - If two people have different prices for step 1, what does that say about their utility functions for money?

Beyond von Neumann-Morgenstern

- The first step of the fun game was a good match to the utility theory we just learned.
    - If two people have different prices for step 1, what does that say about their utility functions for money?
- The second and third steps, not so much!
    - If two people have different prices for step 2, what does *that* say about their utility functions?

Beyond von Neumann-Morgenstern

- The first step of the fun game was a good match to the utility theory we just learned.
  - If two people have different prices for step 1, what does that say about their utility functions for money?
- The second and third steps, not so much!
  - If two people have different prices for step 2, what does *that* say about their utility functions?
  - What if two people have the same prices for step 2 but different prices for step 3?

Representation theorem

Savage 1954

Theorem

Suppose a preference relation satisfies P1–P6; then there exists a utility function U and a probability measure P such that

$$\mathbf{f} \preceq \mathbf{g} \text{ iff } \sum_i P[B_i]U[f_i] \leq \sum_i P[B_i]U[g_i].$$

Representation theorem

Savage 1954

Theorem

Suppose a preference relation satisfies P1–P6; then there exists a utility function U and a probability measure P such that

$$\mathbf{f} \preceq \mathbf{g} \text{ iff } \sum_i P[B_i]U[f_i] \leq \sum_i P[B_i]U[g_i].$$

Savage "postulates"

P1 $\succeq$ is a simple order.

P2 For every $\mathbf{f}, \mathbf{g}$, and B, either $\mathbf{f} \preceq \mathbf{g}$ given B or $\mathbf{g} \preceq \mathbf{f}$ given B.

P3 If $\mathbf{f}(s) = g, \mathbf{f}'(s) = g'$ for every $s \in B$, then $\mathbf{f} \preceq \mathbf{f}'$ given B if and only if $g \preceq g'$.

P4 For every A, B, $P[A] \leq P[B]$ or $P[B] \leq P[A]$.

P5 It is false that for every f, f', $f \preceq f'$.

P6 (Sure-thing principle)

Summary

- Using very simple axioms about preferences over uncertain outcomes, utility theory proves that rational agents ought to act as if they were maximizing the expected value of a real-valued function.

Summary

- Using very simple axioms about preferences over uncertain outcomes, utility theory proves that rational agents ought to act as if they were maximizing the expected value of a real-valued function.

Summary

- Using very simple axioms about preferences over uncertain outcomes, utility theory proves that rational agents ought to act as if they were maximizing the expected value of a real-valued function.
- Can extend beyond this to “subjective” probabilities, using axioms that do not describe how agents manipulate probabilities.

References

Savage, L. (1954).
The Foundations of Statistics.
Dover Publications.

Shoham, Y. and Leyton-Brown, K. (2008).
Multiagent Systems: Algorithmic, Game-theoretic, and Logical Foundations.
Cambridge University Press.

Von Neumann, J. and Morgenstern, O. (1944).
Theory of Games and Economic Behavior.
Princeton University Press.