Expected utility theory

Tai-Wei Hu

University of Bristol

N-Person Normal Form Games

A N-person *normal form game* is given as a triple:

$$G = (N, \{S_i\}_{i \in N}, \{h_i\}_{i \in N}),$$

where
(1): $N = \{1, 2, ..., N\}$—the set of players;

(2): $S_i = \{\mathbf{s}_{i1}, ..., \mathbf{s}_{i\ell_i}\}$—the set of (pure) strategies for player $i = 1, 2, ..., N$;

(3): $h_i : S_1 \times S_2 \to \mathbb{R}$—the payoff function of player $i = 1, 2, ..., N$.

Nash equilibrium

A profile of strategies $s^* = (s_1^*, ..., s_N^*)$ a *Nash equilibrium* if for all $i \in N$,

$$h_i(s_i, s_{-i}^*) \leq h_i(s_i^*, s_{-i}^*) \text{ for all } s_i \in S_i. \quad (1)$$

- unfortunately, not all games have an equilibrium

von Neumann introduced mixed strategies to obtain existence

Mixed strategies

The mixed extension of G is to replace S_i by $M_i = \Delta(S_i)$

- this also requires a new interpretation of h_1 and h_2

Three interpretations of mixed strategies

- as implemented with randomized devices
- as beliefs over other's strategies
- as unpredictable strategies

Expected Utility Theory

Somehow begins wit St. Petersburg paradox (1713)

- need to relax linear utility function
- but other solutions exist....

Developed by von Neumann and Morgenstern (1944)

- foundations to scientific studies of social behaviour

Cornerstone to modern economic analysis

- meaningful cardinal utility
- important for defining efficiency

Outcome space

For our purpose, we consider levels of consumption

- a fixed set of levels, $C = \{c_1 < c_2 < .... < c_n\}$

Expected utility theory considers *lotteries* over the outcomes

- a lottery is a probability distribution μ over C:

$$\mu : C \to [0, 1], \sum_{i=1}^{n} \mu(c_i) = 1$$

- set of lotteries over C denoted by $\Delta(C)$
- we also use c_i to denote degenerate lottery that concentrates on c_i

Preference over lotteries

The primitive is a preference relation over $\Delta(C)$, denoted $\prec$

- $\mu_1 \prec \mu_2$ means the agent will choose μ_2 over μ_1 if given the choice
- more convenient to work with weak preference $\precsim$:
  $\mu_1 \precsim \mu_2$ if $\mu_1 \prec \mu_2$ or indifferent

Relation $\precsim \subset \Delta(C) \times \Delta(C)$ is a *preference relation* if

- it is complete: $\mu_1 \precsim \mu_2$ or $\mu_2 \precsim \mu_1$
- it is transitive: $\mu_1 \precsim \mu_2$ and $\mu_2 \precsim \mu_3$ imply $\mu_1 \precsim \mu_3$

Given $\precsim$, we can define $\prec$ and $\sim$ as

$$\mu_1 \prec \mu_2 \text{ iff } \mu_2 \precsim \mu_1 \text{ does not hold; } \mu_1 \sim \mu_2 \text{ iff } \mu_1 \precsim \mu_2 \text{ and } \mu_2 \precsim \mu_1$$

Compound lotteries

Given two lotteries μ_1 and μ_2 and a number α, define

$$\mu_3 = \alpha\mu_1 + (1-\alpha)\mu_2$$

by

$$\mu_3(c_i) = \alpha\mu_1(c_i) + (1-\alpha)\mu_2(c_i) \text{ for all } c_i \in C$$

- $\mu_3 \in \Delta(C)$
- it can be interpreted as a two-stage lottery:
  - first lottery has outcome μ_1 or μ_2
  - then, execute the resulting lottery μ_1 or μ_2

Expected utility assumes identification of this compound lottery with μ_3

Expected utility representation

Preference relation $\precsim$ over $\Delta(C)$ has an *expected utility representation* if

- there is a function $u : C \to \mathbb{R}$ such that

$$\mu_1 \precsim \mu_2 \text{ if and only if } \sum_{i=1}^{n} \mu_1(c_i)u(c_i) \leq \sum_{i=1}^{n} \mu_2(c_i)u(c_i)$$

In this case, we say that the utility function u represents $\precsim$

- if u represents $\precsim$ and if $v = au + b$ for some $a > 0$, then v also represents $\precsim$
- the utility function u is called *cardinal*

Representation theorem allows us to infer u from observed choices

Axioms for expected utility

EU1 $\precsim$ is a preference relation

EU2 for any $\alpha \in (0, 1)$, and any $\mu_1, \ \mu_2, \ \mu_3 \in \Delta(C)$,

$$\alpha\mu_1 + (1-\alpha)\mu_3 \prec \alpha\mu_2 + (1-\alpha)\mu_3 \text{ if and only if } \mu_1 \prec \mu_2$$

EU3 if $\mu_1 \prec \mu_2 \prec \mu_3$, then there exist $\alpha \in (0, 1)$ such that

$$\alpha\mu_1 + (1-\alpha)\mu_3 \sim \mu_2$$

- these axioms can be verified based on actual choices
- hence the theory is *refutable*

Expected Utility Representation Theorem

Theorem

Let $\precsim$ be a relation over $\Delta(C)$.

1. *$\precsim$ has an expected utility representation iff it satisfies (EU1)-(EU3).*
2. *If u and v both represent $\precsim$, then $u = av + b$ for some $a > 0$.*

- (EU1)-(EU3) are necessary and sufficient for EU representation
- the representation is essentially unique

Proof

Necessity is rather routine

- given u, $\precsim$ is fully determined

Sufficiency requires construction of u

- begin with $u(c_1) = 0$ and $u(c_n) = 1$
- for each other c_i, EU3 implies existence of α such that

$$c_i \sim \alpha c_n + (1 - \alpha) c_1$$

  take $u(c_i) = \alpha$

- finally, show that u represents $\precsim$

Extension to all consumption levels

Up to now we assumed a finite set of outcomes

- without any restrictions on u

Can be extended to $C = \mathbb{R}_+$, and more structures to $\prec$, such as

MC monotonicity: if $c < c'$, then $c \prec c'$

C continuity: for any c, $\{c' : c \precsim c'\}$ and $\{c' : c' \precsim c\}$ are both closed

To do so, consider the set of *simple lotteries* over $\mathbb{R}_+$:

- $\mu : \mathbb{R}_+ \to [0, 1]$ is a simple lottery if

$$\mu(c) = 0 \text{ for all but finitely many } c's \text{ and } \sum_c \mu(c) = 1$$

- the set of simple lotteries is closed under compound lotteries

Risk aversion

Casual observation shows that people do not like risk

- but they make trade-off b/w risk and return
- that is, holding expected value at constant, more risk is less preferred

Formally, we can formulate risk aversion as

RA for all $\mu \in \Delta(\mathbb{R}_+)$, $\mu \precsim \mathbb{E}_\mu(c)$

- $\Delta(\mathbb{R}_+)$ is the set of all simple lotteries
- $\mathbb{E}_\mu(c) = \sum_c \mu(c)c$, expected value of c according to μ

Strict risk aversion requires strict preference whenever μ is not degenerate

Representation theorems for other properties

Theorem

Let $\precsim$ be a relation over $\Delta(\mathbb{R}_+)$ satisfying (EU1)-(EU3) represented by u.

1. *$\precsim$ satisfies MC iff u is strictly increasing.*
2. *$\precsim$ satisfies C iff u is continuous.*
3. *$\precsim$ satisfies (strict) risk aversion iff u is (strictly) concave.*

u is concave if for any c_1, c_2 and any $\alpha \in (0,1)$,

$$\alpha u(c_1) + (1-\alpha)u(c_2) \leq u[\alpha c_1 + (1-\alpha)c_2]$$

u is strictly concave if the inequality is strict

Jensen's inequality

Theorem

A function $f : \mathbb{R}_+ \to \mathbb{R}$ is (strictly) concave iff for any $\mu \in \Delta(\mathbb{R}_+)$,

$$\mathbb{E}_\mu(f) \leq (<) f[\mathbb{E}_\mu(c)].$$

- $\mathbb{E}_\mu(f) = \sum_c \mu(c) f(c)$
- characterization of risk aversion follows immediately

Can be extended to all distributions over $\mathbb{R}_+$ (not just simple ones)

Expected utility maximization

Representation theorem transforms choice problem to utility maximization

- allows for the use of calculus
- typical economic problem is to choose from a *feasible set*

Does the maximum always exit?

Theorem (Extreme Value Theorem)

Suppose that $A \subset \mathbb{R}^n$ is a compact set and that $f : \mathbb{R}^n \to \mathbb{R}$ is a continuous function. Then, there exists $x \in A$ such that

$$f(x) \geq f(x') \text{ for all } x' \in A. \tag{2}$$

- A is compact if it is closed and bounded
- we call the x that satisfies (2) a *maximum* point

First-order conditions

Extreme value theorem only gives existence; but how to find the optimum?

Theorem

Suppose that $A \subset \mathbb{R}^n$ is a compact set and that $f : \mathbb{R}^n \to \mathbb{R}$ is a differentiable function. If x^* is an interior maximum point, then

$$\frac{\partial}{\partial x_i} f(x^*) = 0 \text{ for all } i = 1, ..., n. \tag{3}$$

- $x = (x_1, ..., x_n) \in \mathbb{R}^n$

This condition is only a necessary condition for local optimum

- it only works for interior solutions
- it may not be maximum: it can be minimum as well!
- no guarantee of global optimum either

Concave functions

Theorem

Suppose that $f : \mathbb{R}_+ \to \mathbb{R}$ is a concave function.

1. *f is continuous.*
2. *f is differentiable except for countably many points.*

- that is, axiom RA implies axiom C

Theorem

A differentiable function $f : \mathbb{R}_+ \to \mathbb{R}$ is (strictly) concave if and only if $f'(x)$ is (strictly) decreasing.

- as a result, local maximum is also global maximum

Constrained optimization

Consider the following problem:

$$\max_{x \in [a,b]} f(x),$$

where $f(x)$ is a differentiable concave function

- if $f'(x^*) = 0$ for some $x^* \in (a,b)$, then x^* is a maximum point
- if $f'(a) \leq 0$, then $x^* = a$ is a maximum point
- if $f'(b) \geq 0$, then $x^* = b$ is a maximum point

Example: insurance problem

Two states of the world: high (h) and low (ℓ)

- probability of ℓ is μ
- w/o insurance, consumption at h is w_h and at ℓ is w_ℓ, $w_\ell < w_h$

One unit of insurance pays 1 at ℓ but charges premium p

- with x units of insurance, consumption levels are

$$c_h = w_h - px \text{ and } c_\ell = w_\ell + (1 - p)x$$

- a "fair" premium would be $p = \mu$

Risk aversion implies full insurance under fair premium