von Neumann & Morgenstern Expected Utility
Anscombe & Aumann Setting

Econ 2100 Fall 2018

Lecture 10, October 1

Outline

1. Von Neumann & Morgentstern Expected UtilityTheorem
2. Objective Probabilities?
3. Anscombe & Aumann Acts: Horse Races vs. Roulette Lotteries
4. State-Dependent Expected Utility
5. State Independence

From last class: Expected Utility Theorem

- $X = \{x_1, x_2, \ldots, x_n\}$ is a finite set.
- ΔX is the space of all probabilities on X:

$$\Delta X = \{\pi \in \mathbf{R}^n : \sum_{i=1}^{n} \pi_i = 1 \text{ and } \pi_i \geq 0, \forall i\},$$

- A preference relation is a binary order on ΔX.

Theorem (Expected Utility Theorem, von Neumann and Morgenstern 1947)

Let ΔX be the set of all probabilities on a finite set X. The preference relation $\succsim$ on ΔX is complete, transitive, independent and Archimedean if and only if there exists a function $v : X \to \mathbf{R}$ such that

$$U(\pi) = \sum_{x \in X} v(x)\pi(x)$$

is a representation of $\succsim$. This representation is unique up to affine transformations.

- U represents $\succsim$ means

$$\pi \succsim \rho \Leftrightarrow \sum_{x \in X} v(x)\pi(x) \geq \sum_{x \in X} v(x)\rho(x).$$

- Since the xs are fixed, this compares probability distributions.

Necessity and Uniqueness in vN&M's Expected Utility Theorem

Question 3, Problem Set 5

1. Necessity ($\Leftarrow$) part of vN&M's Expected Utility Theorem
   If there exists a vNM index $v : X \to \mathbf{R}$ such that $u(\pi) = \sum_{x \in X} v(x)\pi(x)$ is a utility representation of $\succsim$, then $\succsim$ is independent and Archimedean.
2. Uniqueness part of vNM's Expected Utility Theorem
   Let $U(\pi) = \sum_x v(x)\pi(x)$ be a utility representation of $\succsim$. Then, $U'(\pi) = \sum_x v'(x)\pi(x)$ is also representation of $\succsim$ if and only if there exist $a > 0$ and $b \in \mathbf{R}$ such that $v'(x) = av(x) + b$ for all $x \in X$.

- We will see two proofs of sufficiency. Both use the mixture space theorem, so all we need to prove is that the affine function is the expected utility function.
- The resason for the second proof is to connect the theorem to properties of the space of linear functions.
  - This is important to better understand the geometry of this result.

Sufficiency of vNM's Expected Utility Theorem

Proof.

Sufficiency (⇒) of vN&M's Expected Utility Theorem

- Let $X = \{x_1, x_2, \ldots, x_n\}$; observe that ΔX is a convex subset of $\mathbf{R}^n$.
- The Mixture Space Theorem implies existence of an affine utility representation $U : \Delta X \to \mathbf{R}$.
- For each i, let $v(x_i) = U(\delta_{x_i})$ (the utility of the Dirac lottery on x_i).
  - This pins down the utility value for each prize.
- Pick some $\pi \in \Delta X$ and denote $\pi_i = \pi(x_i)$.
  - Verify that $\pi = \sum_{i=1}^{n} \pi(x_i)\delta_{x_i} = \sum_{i=1}^{n} \pi_i \delta_{x_i}$
    - this follows because δ_{x_i} is the unit vector pointing in the i-th dimension.
- Since U is affine, each $\pi_i \geq 0$ and $\sum_{i=1}^{n} \pi_i = 1$, we know (Q4, PS 5) that

$$U(\pi) = U(\sum_{i=1}^{n} \pi_i \delta_{x_i}) = \sum_{i=1}^{n} \pi_i U(\delta_{x_i})$$

- By construction, this implies

$$U(\pi) = \sum_{i=1}^{n} \pi_i v(x_i) = \sum_{x \in X} \pi(x) v(x).$$

□

Riesz Representation Theorem

Duality between linear functions and vectors

Lemma

A function $f : \mathbf{R}^n \to \mathbf{R}$ is linear if and only if there exists a unique vector $v \in \mathbf{R}^n$ such that

$$f(x) = \sum_{i=1}^{n} v_i x_i = v \cdot x$$

- This says there is a duality between the linear functions on $\mathbf{R}^n$ and the set $\mathbf{R}^n$ itself:
  - every vector defines a linear function via the dot product, and
  - every linear function is identified with some vector.
- You saw a version of this result in math camp.
- This can be extended to more general vector spaces.

Another Proof of Sufficiency of vNM's Theorem

Proof.

Sufficiency ($\Rightarrow$) of vN&M's Expected Utility Theorem (Again)

- Let $X = \{x_1, x_2, \ldots, x_n\}$. Observe that ΔX is a convex subset of $\mathbf{R}^n$.
- By the Mixture Space Theorem, there exists $U : \Delta X \to \mathbf{R}$ that is an affine representation of $\succsim$.
- Any affine function from Π to $\mathbf{R}$ can be extended to another affine function $\overline{U} : \mathbf{R}^n \to \mathbf{R}$ such that $\overline{U}(\pi) = U(\pi)$ for all $\pi \in \Pi$. (prove as an exercise)
- WLOG, assume $\overline{U}$ is linear by subtracting the constant $\overline{U}(\mathbf{0}_n)$ (see last class).
- By Riesz, there exists a unique vector $v \in \mathbf{R}^n$ such that for all $z \in \mathbf{R}^n$

$$\overline{U}(z) = v \cdot z = \sum_{i=1}^{n} v_i z_i.$$

- Define $v(x_i) = v_i$ and observe that $\pi(x_i) = \pi_i$ for all $\pi \in \Delta X$.
- Then, for any $\pi \in \Delta X$,

$$U(\pi) = \overline{U}(\pi) = \sum_{i=1}^{n} v(x_i)\pi(x_i) = \sum_{x \in X} \pi(x) v(x).$$

□

Pessimistic Expected Utility

Example

Consider this alternative representation of $\succsim$:

$$U(\pi) = \min\{v(x) : \pi(x) > 0\}.$$

- DM evaluates each lottery by the worst element on the support. If there is a non zero chance, no matter how small, to get a terrible outcome the decision maker will evaluate the lottery as if she was going to receive that outcome for sure.
  - Extreme pessimism: "if something bad is possible, it will happen".
  - Sometimes called "infinite risk-aversion" (we will talk about risk-aversion soon).

Question 5, Problem Set 5

Suppose there exists a vNM utility index $v : X \to \mathbf{R}$ such that the following is a utility representation for $\succsim$:

$$U(\pi) = \min\{v(x) : \pi(x) > 0\}.$$

Prove or disprove the following: (a) $\succsim$ is independent; (b) $\succsim$ is Archimedean; and (c) $\min\{v'(x) : \pi(x) > 0\}$ is a representation of $\succsim$ if and only if $v' = av + b$ for some $a > 0, b \in \mathbf{R}$.

Expected Utility with Infinitely Many Prizes

- Allowing infinitely many prizes requires some more advanced functional analysis, and introduces some tricky issues.
- Suppose the real line is the space of consequences, what is the equivalent of ΔX? Call it $\Delta^*\mathbf{R}$.
  - $\Delta^*\mathbf{R}$ the space of (Borel) probability measures measures on $\mathbf{R}$, or
  - $\Delta^*\mathbf{R}$ the set of density functions on $\mathbf{R}$ with finite variance, or...
- For each choice one needs the appropriate version of "continuity".

Theorem

The preference relation $\succsim$ on $\Delta^*\mathbf{R}$ is complete, transitive, independent, and "continuous" if and only if there exists a "particular" $v : \mathbf{R} \to \mathbf{R}$ such that

$$U(f) = \int v(x)f(x)dx$$

is a representation of $\succsim$.

- The meaning of "continuous" and "particular" depend on details which are beyond the scope of this course (ask Roee).

Objective Probabilities

- **von Neumann and Morgenstern**: $\succsim$ on ΔX is complete, transitive, independent, and Archimedean if and only if there exists $v : X \to \mathbf{R}$ such that
$$U(\pi) = \sum_{x \in X} v(x)\pi(x)$$
is a representation of $\succsim$.
- Preferences rank lotteries over a given set of prizes.
- Information about the preferences implies existence of a representation with a particular functional form, and pins down the vNM utility function v.
- Probabilities (as well as consequences) are known primitives of the model:
  - The decision maker ranks pairs probability distributions.
- Where do these probabilities come from?

- In many (most) cases, however, probabilities are not given; they reflect what each decision maker thinks, and could differ across individuals.
- To account for this, we want a model that lets the likelihood of events be in the mind of the decision maker.
- In that model, from DM's preference relation one figures out her probability distribution over events as well as her utility for consequences.
- To achieve this one needs a more general consumption space and a new axiom.

Anscombe and Aumann Structure

- $\Omega = \{1, 2, \ldots, S\}$ is a finite set of states, with generic element $s \in \Omega$.
- $X = \{x_1, \ldots, x_n\}$ is a finite set of outcomes, with a generic element $x \in X$.
- $H = (\Delta X)^\Omega$ is the space of all functions from Ω to ΔX.
  - this is a convex subset of the space of functions from Ω to $\mathbf{R}^n$.

Anscombe–Aumann acts

- An Anscombe–Aumann act $h \in H$ is a function $h : \Omega \to \Delta X$;
  - A–A acts assign a lottery (an element of ΔX) to each state (an element of Ω).

- Let $h_s = h(s) \in \Delta X$, and denote $h_s(x) = [h(s)](x) \in [0, 1]$.
  - This is the probability of x conditional on s, given the act h: $h_s(x) = \mathrm{Pr}(x|s, h)$.

Notation

- A $\pi \in \Delta X$ denotes the "constant" act $f : \Omega \to \Delta X$ s.t. $f(s) = \pi$ for all $s \in \Omega$.
- Given $H \subset (\mathbf{R}^n)^\Omega$, if $f, g : \Omega \to \mathbf{R}^n$, then the function $\alpha f + \beta g : \Omega \to \mathbf{R}^n$ is defined by $[\alpha f + \beta g](s) = \alpha f(s) + \beta g(s)$.
  - This definition is crucial: Archimedean and Indepdence axioms sum acts.
    - Summing is all about the objective lotteries (not about the horse race).

Horse Lotteries and Roulette Lotteries

- There are two kinds of sources of randomness: states of the world (Ω) and lotteries over consequences (ΔX).
- How can we think about them?

Interpretation of H

- Lotteries over consequences are bets on an objective "roulette" spin:
  - outcomes' probabilities are objectively determined (everyone agrees on them).
- A state of the world represent the event that a specific "horse" named s wins a race among the field of horses Ω:
  - the decision maker subjectively assesses each horse's strength (different DMs can evaluate each horse differently).
- The theory's aim is to identify the decision maker's personal assessment of the probability that horse s will win the race using her preferences.
- To perform this identification, we set the payout on horse s equal to a lottery that depends on the outcome of a roulette spin.
- So, first the horses run the race, and afterwards the roulette is spun.
- The roulette's payoff can depend on which horse wins.

How can we describe elements of H (functions from Ω to ΔX)?

First description of H

First. *The original mathematical interpretation where* $H = (\Delta X)^{\Omega}$*.*

- Suppose $\Omega = \{s_1, s_2, s_3\}$ and $X = \{x_1, x_2, x_3\}$.
- Then a particular $h : \Omega \to \Delta X$ would be the following:

$$\begin{aligned} h(s_1) &= (0.3, 0.2, 0.5) \\ h(s_2) &= (0.4, 0.6, 0) \\ h(s_3) &= (0, 1, 0) \end{aligned}$$

Second description of H

Second. *H as a set of compound lotteries.*

- The subjective first stage lottery is over which state $s \in \Omega$ obtains, and the objective second stage lottery (conditional on s) is over which $x \in X$ obtains.
- These compound lotteries can be written as probability trees.
- For example: h is

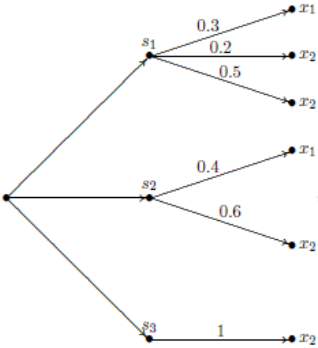

,

Third description of H

Third. *H is the set of weakly positive $m \times n$ matrices where*

$$\sum_{j=1}^{n} h_{s,j} = 1 \text{ for each } s = 1, \ldots S.$$

- Then the h on the previous slides can be represented as

$$h = \begin{bmatrix} & \Pr(x = x_1) & \Pr(x = x_2) & \Pr(x = x_3) \\ s_1 & 0.3 & 0.2 & 0.5 \\ s_2 & 0.4 & 0.6 & 0 \\ s_3 & 0 & 1 & 0 \\ \end{bmatrix}.$$

- We can then write $h_s(x) = h_{s,x}$; for example $h_{s_1}(x_2) = h_{1,2} = 0.2$.

Objective Lotteries Are Anscombe-Aumann Acts

Notation

- One can identify ΔX as a subset of H.
  - The probability distributions on X are acts that, regardless of the state, give the same lottery over outcomes.
  - These are usually called constant acts, and the set of all constant acts is denoted H_c (this subset of H is isomorphic to ΔX)

$$\Delta X \approx H_c = \{f \in H : f(s) = f(s') \text{ for all } s \in \Omega\}$$

- One can also identify X as a subset of H.
  - Elements of X are Dirac lotteries (degenerate probability distributions) in ΔX, denoted δ_x.
  - Thus, X is a subset of H_c defined as follows

$$X = \{f \in H : f(s) = f(s')\ \forall\ s \in \Omega, \text{ and } f(s) = \delta_x \text{ for some } x \in X\}$$

- Anscombe-Aumann acts generalize von Neumann and Morgenstern's setting.

- Without extra assumptions, we can use the Mixture Space Theorem on H.

State Dependent Expected Utility

Theorem

The preference relation $\succsim$ on H is complete, transitive, independent, and Archimedean if and only if there exists a set of vNM indices $v_1, \ldots, v_S : X \to \mathbf{R}$ such that

$$U(h) = \sum_{s \in \Omega} \sum_{x \in X} v_s(x) h_s(x)$$

is a utility representation.

- This follows from the Mixture Space Theorem applied to H (a convex set).
- The sufficiency proof is similar to vNM: find a function $U : \mathbf{R}^{S \times n} \to \mathbf{R}$ representing $\succsim$;
  - this linear function is uniquely determined by a vector in $\mathbf{R}^{S \times n}$.

Remark

The utility of each consequence depends on the state in which it obtains (formally, $v_s(x)$ depends on s); this is a state-dependent additive representation.

State Dependent Expected Utility

Theorem

The preference relation $\succsim$ on H is complete, transitive, independent, and Archimedean if and only if there exists a set of vNM indices $v_1, \ldots, v_S : X \to \mathbf{R}$ such that

$$U(h) = \sum_{s \in \Omega} \sum_{x \in X} v_s(x) h_s(x)$$

is a utility representation.

Can we identify a unique probability distribution over Ω?

- If we could, DM would take the 'subjective' expectation with respect to some probability distribution μ over Ω of a state-dependent vNM index v_s:

$$U(h) = \mathbf{E}_{\mu}(\mathbf{E}_{h_s}(v_s)).$$

- In other words, one would like the utility function to be

$$U(h) = \sum_{s \in \Omega} \mu(s) \left[\sum_{x \in X} v_s(x) h_s(x) \right],$$

- In the last exercise of Problem Set 5, you will show that such a μ cannot be uniquely identified.

State Dependent Expected Utility: Discussion

- More precisely, suppose we have
$$\sum_{s\in\Omega}\mu(s)\left[\sum_{x\in X} v_s(x)h_s(x)\right]$$
where μ is a probability distribution over Ω.
- Then for any $\mu' \in \Delta\Omega$ such that $\mu'(s) > 0$ for all $s \in \Omega$, there exist indices $v_1', \ldots v_S' : X \to \mathbf{R}$ such that
$$\sum_{s\in\Omega}\mu'(s)\left[\sum_{x\in X} v_s'(x)h_s(x)\right].$$

Remark

- One cannot pin down probabilities using state dependent expected utility.
  - Any representation that has $\mu, v_s(\cdot)$ is equivalent to a representation that uses $\mu', v_s'(\cdot)$.

- Because of this, under the assumptions of the state-dependent expected utility theorem, we cannot think of DM's preference identifying a unique probability over the state space.
- For this identification to be possible, one needs a "state-independent" representation in which the function v does not change across states.

Null States

Notation

Given an act $h \in (\Delta X)^\Omega$, a state $s \in \Omega$, and a lottery $\pi \in \Delta X$, define the new act $(h_{-s}, \pi) : \Omega \to \Delta X$ by $(h_{-s}, \pi) = (h_1, \ldots, h_{s-1}, \pi, h_{s+1}, \ldots, h_m)$. So

$$[(h_{-s}, \pi)](t) = \begin{cases} \pi & \text{if } t = s \\ h(t) & \text{if } t \neq s \end{cases};$$

- (h_{-s}, π) replaces h_s (the lottery that act h assigns to state s) with the lottery π while the remainder of h stays the same.

- With this notation, one can describe states the decision maker never cares about since they should be irrelevant to the representation.

Definitions

- A state $s \in \Omega$ is null if, for all $h \in (\Delta X)^\Omega$ and $\pi, \rho \in \Delta X$, $(h_{-s}, \pi) \sim (h_{-s}, \rho)$.
- A state $s \in \Omega$ is non-null if it is not null, i.e. if there exist $h \in (\Delta X)^\Omega$ and $\pi, \rho \in \Delta X$ such that $(h_{-s}, \pi) \succ (h_{-s}, \rho)$.

- A state is null if it never affects rankings.

State Independence

Definition

The binary relation $\succsim$ on H is state-independent if, for all non-null states $s, t \in \Omega$, for all acts $h, g \in H$, and for all lotteries $\pi, \rho \in \Delta X$,

$$(h_{-s}, \pi) \succsim (h_{-s}, \rho) \quad \Rightarrow \quad (g_{-t}, \pi) \succsim (g_{-t}, \rho).$$

- In words, the ranking of roulette lotteries does not depend on the state in which they obtain.
  - This only needs to hold in states the decision maker cares about.
- Implies the utility index over consequences is the same across states.

Exercise

Let $\Omega = \{rainy\ day, sunny\ day\}$ and $X = \{Umbrella, Hat\}$. The decision maker's preferences satisfy:

- $(\delta_{\mathbf{U}}, \delta_{\mathbf{H}}) \succ (\delta_{\mathbf{U}}, \delta_{\mathbf{U}})$ (she strictly prefers an umbrella in the rain and a hat in the sun to an umbrella for sure) and
- $(\delta_{\mathbf{U}}, \delta_{\mathbf{H}}) \succ (\delta_{\mathbf{H}}, \delta_{\mathbf{H}})$ (she strictly prefers an umbrella in the rain and a hat in the sun to getting sunglasses for sure).

Verify that these preferences violate state-independence.

Next Class

- Subjective Expected Utility
- Expected Utility Over Money