Independence, Mixture Space Theorem, and von Neumann & Morgenstern Expected Utility Theorem

Econ 3030 Fall 2025

Lecture 9

Outline

1. Convex Consumption Set and Independence
2. Mixture Space Theorem
3. Preferences Over Lotteries
4. von-Neuman & Morgenstern Expected Utility

Expected Utility

- There are n fixed mutually exclusive outcomes denoted $x_1, \ldots, x_n$.
- A probability distribution over these outcomes is $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_n)$.
- Since the outcomes are fixed, the consumer evaluates different probability distributions over outcomes.

- The expected utility function is
$$U(\boldsymbol{\pi}) = \sum_{i=1}^{n} \pi_i v(x_i).$$
where each outcome x_i has utility $v(x_i)$.
- If we let $\mathbf{x} = (x_1, \ldots, x_n)$ and $\mathbf{v}(\mathbf{x}) = (v(x_1), \ldots, v(x_n))$, we can rewrite this as
$$U(\boldsymbol{\pi}) = \boldsymbol{\pi} \cdot \mathbf{v}(\mathbf{x})$$

Remark

The representation theorem has to tell us that there exist a function v such that the utility function has the particular functional form given by U.

Convex Consumption Space

- To obtain the representation in the previous slide, we need more than completeness, transitivity, and continuity.
  - This is intuitive since the utility function $U(\cdot)$ has a very special functional form.
- We start by modifying the space over which consumption is defined.

Consumption Set is Convex

The consumption set is a convex subset of $\mathbb{R}^n$ denoted Π.

- Convexity means that

$$\text{if } \mathbf{x}, \mathbf{y} \in \Pi, \quad \text{then} \qquad \alpha\mathbf{x} + (1-\alpha)\mathbf{y} \in \Pi \quad \text{for all } \alpha \in (0,1).$$

The Plan

- Start with preferences over an abstract convex space.
- Then, add more structure to get more specific results.

Standard Assumptions

- As always, preferences must be complete and transitive

Definition

A binary relation $\succsim$ on Π is:

- complete if, for all $\mathbf{x}, \mathbf{y} \in \Pi$, $\mathbf{x} \succsim \mathbf{y}$ or $\mathbf{y} \succsim \mathbf{x}$, or both;
- transitive if, for all $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \Pi$, $\mathbf{x} \succsim \mathbf{y}$ and $\mathbf{y} \succsim \mathbf{z}$ imply $\mathbf{x} \succsim \mathbf{z}$;

Archimedean Axiom

- Since the consumption space is convex, one can use a weaker version of continuity.

Definition

A binary relation $\succsim$ on Π is Archimedean if, for all $\boldsymbol{\pi}, \boldsymbol{\rho}, \boldsymbol{\sigma} \in \Pi$,

$$\boldsymbol{\pi} \succ \boldsymbol{\rho} \succ \boldsymbol{\sigma} \Longrightarrow \begin{cases} \exists\, \alpha \in (0,1) \text{ such that } \alpha\boldsymbol{\pi} + (1-\alpha)\boldsymbol{\sigma} \succ \boldsymbol{\rho} \\ \text{and} \\ \exists\, \beta \in (0,1) \text{ such that } \boldsymbol{\rho} \succ \beta\boldsymbol{\pi} + (1-\beta)\boldsymbol{\sigma} \end{cases}$$

Exercise

Show that if $\succsim$ is continuous then it is Archimedean.

Exercise

Let $\Pi = \mathbb{R}$ and let $\succsim$ on $\mathbb{R}$ defined by the utility function

$$U(\pi) = \begin{cases} 1 & \text{if } \pi > 0 \\ 0 & \text{if } \pi = 0 \\ -1 & \text{if } \pi < 0 \end{cases}$$

. Verify that $\succsim$ is Archimedean but not continuous.

Independence

- A crucial new assumption yields additive separability of the representation.

Definition

A binary relation $\succsim$ on Π satisfies independence if, for all $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \Pi$ and $\alpha \in (0,1)$,

$$\mathbf{x} \succsim \mathbf{y} \Leftrightarrow \alpha \mathbf{x} + (1-\alpha)\mathbf{z} \succsim \alpha \mathbf{y} + (1-\alpha)\mathbf{z}.$$

Example

Suppose $\Pi = \mathbb{R}^2$ and $\succsim$ defined by

$$\mathbf{x} \succsim \mathbf{y} \text{ if and only if } x_1^2 + x_2^2 \geq y_1^2 + y_2^2$$

$\mathbf{x}$ is weakly preferred to $\mathbf{y}$ whenever the norm of $\mathbf{x}$ is weakly larger than the norm of $\mathbf{y}$

Then $(4,0) \sim (0,4)$, but

$$\frac{1}{2}(4,0) + \frac{1}{2}(2,0) = (3,0) \succ (1,2) = \frac{1}{2}(0,4) + \frac{1}{2}(2,0)$$

So $\succsim$ is not independent.

- What does independence imply geometrically?

Consequences of Independence

Proposition

If a preference relation satisfies independence its indifference classes are convex.

Proof.

Suppose $\succsim$ is independent. We need to show that

$$\boldsymbol{\pi} \sim \boldsymbol{\sigma} \Rightarrow \boldsymbol{\pi} \sim \alpha\boldsymbol{\pi} + (1-\alpha)\boldsymbol{\sigma} \sim \boldsymbol{\sigma}, \ \forall \alpha \in [0,1]$$

- If $\boldsymbol{\pi} \sim \boldsymbol{\sigma}$, clearly $\boldsymbol{\pi} \succsim \boldsymbol{\sigma}$. Thus, by independence, for all $\alpha \in [0,1]$,
$$\boldsymbol{\pi} \succsim \boldsymbol{\sigma} \Rightarrow \alpha\boldsymbol{\pi} + (1-\alpha)\boldsymbol{\pi} \succsim \alpha\boldsymbol{\sigma} + (1-\alpha)\boldsymbol{\pi} \Rightarrow \boldsymbol{\pi} \succsim \alpha\boldsymbol{\sigma} + (1-\alpha)\boldsymbol{\pi}.$$
- If $\boldsymbol{\pi} \sim \boldsymbol{\sigma}$, clearly $\boldsymbol{\sigma} \succsim \boldsymbol{\pi}$. Thus, by independence, for all $\alpha \in [0,1]$,
$$\boldsymbol{\sigma} \succsim \boldsymbol{\pi} \ \Rightarrow \alpha\boldsymbol{\sigma} + (1-\alpha)\boldsymbol{\pi} \succsim \alpha\boldsymbol{\pi} + (1-\alpha)\boldsymbol{\pi} \ \Rightarrow \ \alpha\boldsymbol{\sigma} + (1-\alpha)\boldsymbol{\pi} \succsim \boldsymbol{\pi}.$$
- Therefore we get $\boldsymbol{\pi} \sim \alpha\boldsymbol{\sigma} + (1-\alpha)\boldsymbol{\pi}$ for all $\alpha \in [0,1]$.
- The same logic shows that $\boldsymbol{\sigma} \sim \alpha\boldsymbol{\sigma} + (1-\alpha)\boldsymbol{\pi}$.
- Therefore the indifference classes are convex. □

- The following provides an alternate characterization of independence, which is sometimes useful in proofs.

Question 1, Problem Set 5.

Prove that a binary relation on Π is independent if and only if, for all $\boldsymbol{\pi}, \boldsymbol{\rho}, \boldsymbol{\sigma} \in \Pi$, and $\alpha \in (0, 1)$,

$$\boldsymbol{\pi} \succ \boldsymbol{\rho} \Leftrightarrow \alpha\boldsymbol{\pi} + (1-\alpha)\boldsymbol{\sigma} \succ \alpha\boldsymbol{\rho} + (1-\alpha)\boldsymbol{\sigma}$$

and

$$\boldsymbol{\pi} \sim \boldsymbol{\rho} \Leftrightarrow \alpha\boldsymbol{\pi} + (1-\alpha)\boldsymbol{\sigma} \sim \alpha\boldsymbol{\rho} + (1-\alpha)\boldsymbol{\sigma}$$

Linear and Affine Functions

Definition

A function $f : \Pi \to \mathbf{R}$ is affine if, for all $\boldsymbol{\pi}, \boldsymbol{\rho} \in \Pi$ and $\alpha \in [0, 1]$

$$f(\alpha\boldsymbol{\pi} + (1 - \alpha)\boldsymbol{\rho}) = \alpha f(\boldsymbol{\pi}) + (1 - \alpha) f(\boldsymbol{\rho}).$$

Definition

A function $g : \mathbf{R}^n \to \mathbf{R}$ is linear if, for all $\boldsymbol{\pi}, \boldsymbol{\rho} \in \mathbf{R}^n$ and $\alpha, \beta \in \mathbf{R}$,

$$g(\alpha\boldsymbol{\pi} + \beta\boldsymbol{\rho}) = \alpha g(\boldsymbol{\pi}) + \beta g(\boldsymbol{\rho}).$$

Exercise

Prove that a function $g : \mathbf{R}^n \to \mathbf{R}$ is affine if and only if $f(\boldsymbol{\pi}) = g(\boldsymbol{\pi}) - g(\mathbf{0}_n)$ is linear.

Mixture Space Theorem

Theorem (Mixture Space Theorem, Herstein and Milnor)

A binary relation $\succsim$ on Π (a convex subset of $\mathbb{R}^n$) is complete, transitive, independent and Archimedean if and only if there exists an affine function $U : \Pi \to \mathbb{R}$ such that

$$\boldsymbol{\pi} \succsim \boldsymbol{\rho} \Leftrightarrow U(\boldsymbol{\pi}) \geq U(\boldsymbol{\rho})$$

Moreover, if $U : \Pi \to \mathbb{R}$ represents $\succsim$, then $U' : \Pi \to \mathbb{R}$ also represents $\succsim$ if and only if there exist real numbers $a > 0$ and b such that $U'(\boldsymbol{\pi}) = aU(\boldsymbol{\pi}) + b$ for all $\boldsymbol{\pi} \in \Pi$.

- The first part of the statement states that preferences are represented by an affine utility function, while the second says that this representation is unique up to linear transformations.

Remarks

- This holds for any convex subset of an arbitrary vector space.
- The utility function $U(\cdot)$ is cardinal and not just ordinal as before (Why?).
- The Mixture Space Theorem asserts that there exists some affine representation, not that all representations are affine.

Mixture Space Theorem

Theorem (Mixture Space Theorem, Herstein and Milnor)

A binary relation $\succsim$ on Π (a convex subset of $\mathbb{R}^n$) is complete, transitive, independent and Archimedean if and only if there exists an affine function $U : \Pi \to \mathbb{R}$ representing $\succsim$:

$$\boldsymbol{\pi} \succsim \boldsymbol{\rho} \Leftrightarrow U(\boldsymbol{\pi}) \geq U(\boldsymbol{\rho})$$

Moreover, this representation is unique up to affine transformations.

- As with Debreu's utility representation proof, the main step is to find the unique number that is indifferent to a given $\boldsymbol{\pi} \in \Pi$.
  - The main difference is that the assumptions now imply an affine utility function (rather than continuous one as in Debreu's theorem).
- Proof: Math class.
- Next we will see how this theorem, when used on special convex consumption sets, yields an expected utility representation.

Preferences Over Lotteries (Von Neumann and Morgenstern (1947))

- Let $X = \{x_1, x_2, \ldots, x_n\}$ be a finite set of size n (each x_i is a 'prize').
- Let $\pi(x_i) = \pi_i$ the probability of receiving 'prize' x_i, and define

$$\Delta X = \left\{ \boldsymbol{\pi} \in \mathbb{R}^n : \forall i \; \pi_i \geq 0, \text{ and } \sum_{i=1}^{n} \pi_i = 1 \right\}$$

- ΔX is a convex subset of the vector space $\mathbb{R}^n$.
  - An element of ΔX identifies a lottery over the elements of X (the 'prizes').
  - The degenerate lottery that yields x with certainty is called a Dirac lottery on x and denoted δ_x;
    - $\boldsymbol{\delta}_{x_k}$ is the unit vector in the direction k: $\boldsymbol{\delta}_{x_k} = (0_1, \ldots, 0_{k-1}, 1_k, 0_{k+1}, \ldots, 0_n)$.
    - The set of Dirac lotteries is $\{\boldsymbol{\delta}_x : x \in X\} \subset \Delta X$, and it constitutes the extreme points of ΔX (what are the extreme points?).

$\succsim$ is defined over ΔX

Remark

A preference relation ranks probability distributions over a *fixed* finite set of objects. Since the set of prizes is fixed, the decision maker's preference order is over lotteries.

Lotteries

- If $X = \{x_1, x_2, x_3\}$, a typical lottery $\boldsymbol{\pi}$ is described using an event tree:

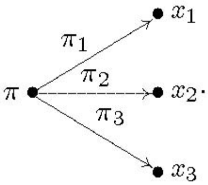

- Then $\boldsymbol{\delta}_{x_2}$, the degenerate lottery which yields x_2 with certainty, is:

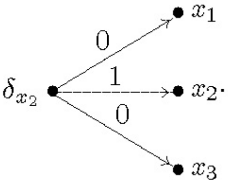

- The space ΔX assumes all uncertainty is resolved at one point in time; it does not allow for compound lotteries (lotteries over lotteries).
  - This domain restriction can be justified by introducing a 'reduction of compound lotteries' assumption as to reduce every compound lottery to a single lottery in ΔX.

Compound Lotteries

- The convex combination $\alpha\boldsymbol{\pi} + (1-\alpha)\boldsymbol{\sigma}$ might be interpreted as the compound lottery which yields $\boldsymbol{\pi}$ with probability α and $\boldsymbol{\sigma}$ with probability $1-\alpha$.
- If compounded correctly, this yields the same probabilities on consequences as $\alpha\boldsymbol{\pi} + (1-\alpha)\boldsymbol{\sigma}$:

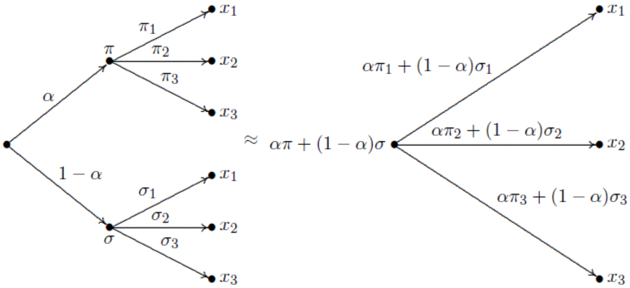

- This assumes the decision maker compounds correctly (can we state this assumption?).
  - Suppose Z is a finite subset of ΔX. A lottery $\pi \in \Delta Z$ is a compound lottery, because it is a lottery over lotteries. If one takes a $\boldsymbol{\pi} \in Z$ then $\boldsymbol{\delta}_\pi$ is an element of ΔX, a lottery over X.

Independence and Lotteries

Independence: *for all* $\boldsymbol{\pi}, \boldsymbol{\pi}', \boldsymbol{\sigma} \in \Delta X$ *and* $\alpha \in (0,1)$,

$$\boldsymbol{\pi} \succsim \boldsymbol{\pi}' \Leftrightarrow \alpha\boldsymbol{\pi} + (1-\alpha)\boldsymbol{\sigma} \succsim \alpha\boldsymbol{\pi}' + (1-\alpha)\boldsymbol{\sigma}$$

Hence:

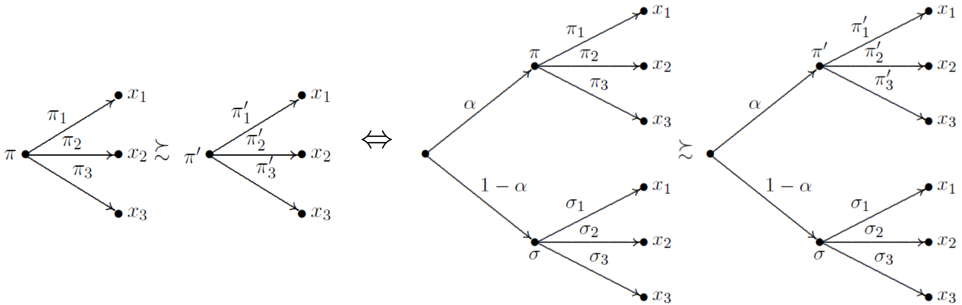

- The decision maker cares only about paths which differ.
- This is a 'normative' justification for the independence axiom on ΔX.

Expected Utility

Things we already know

- Under completeness, transitivity and continuity: there exists a continuous utility function representing the preferences.
- If X is convex, replace continuity with the Archimedean axiom and add independence: that utility function is affine.

- Under the extra structure given by ΔX, the representation theorem identifies a function $v : X \to \mathbb{R}$ such that the preference $\succsim$ is represented by the function

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

  - DM weights the utility of each outcome by the probability of receiving that outcome.
- The probability distribution over prizes $\boldsymbol{\pi}$ is given; hence the theorem identifies, via preferences, the functional form of $U(\cdot)$ and the function $v(\cdot)$.

Remark

- By letting $v_i = v(x_i)$ for $i = 1, .., n$, the function v yields a vector $\mathbf{v} \in \mathbb{R}^n$. The expected utility formula is the dot product of two vectors ($\mathbf{v}$ and $\boldsymbol{\pi}$) in $\mathbb{R}^n$.

Expected Utility Theorem

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

Given a finite set X, the preference relation $\succsim$ on ΔX (the set of all probability distributions on X) is complete, transitive, independent, and Archimedean if and only if there exists a function $v : X \to \mathbb{R}$ such that $\succsim$ is represented by

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

This representation is unique up to affine transformations.

- U represents $\succsim$ means: $\boldsymbol{\pi} \succsim \boldsymbol{\rho} \Leftrightarrow U(\boldsymbol{\pi}) \geq U(\boldsymbol{\rho})$.
- Uniqueness means: $U'(\boldsymbol{\pi}) = \sum_x v'(x)\pi(x)$ also represents $\succsim$ if and only if there exist $a > 0$ and $b \in \mathbf{R}$ such that $v'(x) = av(x) + b$ for all $x \in X$.

Remark

- The function v is called *von Neumann & Morgenstern utility index* or *Bernoulli utility*.
  - $v : X \to \mathbb{R}$ is not the representation of $\succsim$; its domain is X, which is not equal to ΔX.
  - The utility index v is a component of the utility representation given by U, which is defined on ΔX (the correct domain).

Necessity and Uniqueness in vNM's Expected Utility Theorem

Question 3, Problem Set 5

1. Necessity ($\Leftarrow$) part of vNM'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 $b \in \mathbb{R}$ such that $v'(x) = av(x) + b$ for all $x \in X$.

- We will see a proof of sufficiency next. It uses the mixture space theorem, so all we need to prove is that the affine function is the expected utility function.

Sufficiency of vNM's Expected Utility Theorem

Proof.

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

- Let $X = \{x_1, x_2, \ldots, x_n\}$; observe that ΔX is a convex subset of $\mathbb{R}^n$.
- By the Mixture Space Theorem, there is an affine utility representation $U : \Delta X \to \mathbb{R}$.
- For each i, let $v(x_i) = U(\boldsymbol{\delta}_{x_i})$ (the utility of the Dirac lottery on x_i).
  - This yields a function $v : X \to \mathbf{R}$ and pins down the utility value of each prize.
- Pick some $\boldsymbol{\pi} \in \Delta X$ and denote $\pi_i = \pi(x_i)$.
  - Verify that $\boldsymbol{\pi} = \sum_{i=1}^{n} \pi(x_i)\boldsymbol{\delta}_{x_i} = \sum_{i=1}^{n} \pi_i \delta_{x_i}$
    - this follows because $\boldsymbol{\delta}_{x_i}$ 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

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

- By construction, this implies

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

□

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 $\Delta^*\mathbb{R}$.
  - $\Delta^*\mathbb{R}$ the space of (Borel) probability measures measures on $\mathbb{R}$, or
  - $\Delta^*\mathbb{R}$ the set of density functions on $\mathbb{R}$ 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" $v : \mathbb{R} \to \mathbb{R}$ such that

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

is a representation of $\succsim$.

Next Class

- Subjective vs. Objective Probability
- Anscombe and Aumann Acts
- State Independence
- Subjective Expected Utility