Decision Making Under Uncertainty

Econ 2100 Fall 2018

Lecture 9, September 26

Outline

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

Decision Making Under Uncertainty: Prelude

- So far, consumption has been an "here and now" matter.
  - Preference orderings compare alternatives available for immediate use.
- Next, the decision maker must choose now among items that will be consumed in the future.
  - This is reasonably straightforward when the future is known with certainty, one only needs to worry about discounting (future consumption may not be as valuable as current consumption).
  - Things are more complicated (and interesting) when there is uncertainty about what will happen.
  - Then, future consumption depends on current choices through the way in which uncertainty is resolved.
- We will typically think of a future consumption vector as a random variable:
  - only one of many possibilities (states of the world) will occur, but
  - an exhaustive list of all these possibilities (the state space) describes possible future consumption.
- Preferences compare (now) alternatives that will be consumed in the future.
- We will see classic results on utility functions representing preferences. These results are more about getting a specific useful functional form than establishing existence of a utility function.

Uncertainty and Utility: Debreu's Approach

Two Periods

Decisions are made today but items are consumed tomorrow.

Uncertainty and Utility: Debreu's Approach

Two Periods

Decisions are made today but items are consumed tomorrow.

- List of all the ways in which uncertainty can be resolved tomorrow.

Uncertainty and Utility: Debreu's Approach

Two Periods

Decisions are made today but items are consumed tomorrow.

- List of all the ways in which uncertainty can be resolved tomorrow.
  - Each of this possible resolutions is a 'state of the world', denoted s, and S is the set of all states of the world.

Uncertainty and Utility: Debreu's Approach

Two Periods

Decisions are made today but items are consumed tomorrow.

- List of all the ways in which uncertainty can be resolved tomorrow.
  - Each of this possible resolutions is a 'state of the world', denoted s, and S is the set of all states of the world.
- A 'commodity' is defined by its physical description (e.g. umbrella), as well as the state in which that good is consumed (e.g. tomorrow is a rainy day).

Uncertainty and Utility: Debreu's Approach

Two Periods

Decisions are made today but items are consumed tomorrow.

- List of all the ways in which uncertainty can be resolved tomorrow.
  - Each of this possible resolutions is a 'state of the world', denoted s, and S is the set of all states of the world.
- A 'commodity' is defined by its physical description (e.g. umbrella), as well as the state in which that good is consumed (e.g. tomorrow is a rainy day).
  - Thus, "an umbrella on a rainy day" is a commodity; "an umbrella on a sunny day" is a different commodity.

Uncertainty and Utility: Debreu's Approach

Two Periods

Decisions are made today but items are consumed tomorrow.

- List of all the ways in which uncertainty can be resolved tomorrow.
  - Each of this possible resolutions is a 'state of the world', denoted s, and S is the set of all states of the world.
- A 'commodity' is defined by its physical description (e.g. umbrella), as well as the state in which that good is consumed (e.g. tomorrow is a rainy day).
  - Thus, "an umbrella on a rainy day" is a commodity; "an umbrella on a sunny day" is a different commodity.
- Do this for all L "physical goods" in all states S, and define the consumption space as $\mathbf{R}^{L \times S}$.

Uncertainty and Utility: Debreu's Approach

Two Periods

Decisions are made today but items are consumed tomorrow.

- List of all the ways in which uncertainty can be resolved tomorrow.
  - Each of this possible resolutions is a 'state of the world', denoted s, and S is the set of all states of the world.
- A 'commodity' is defined by its physical description (e.g. umbrella), as well as the state in which that good is consumed (e.g. tomorrow is a rainy day).
  - Thus, "an umbrella on a rainy day" is a commodity; "an umbrella on a sunny day" is a different commodity.
- Do this for all L "physical goods" in all states S, and define the consumption space as $\mathbf{R}^{L \times S}$.
  - A consumption vector specifies quantities of all these goods.

Uncertainty and Utility: Debreu's Approach

Two Periods

Decisions are made today but items are consumed tomorrow.

- List of all the ways in which uncertainty can be resolved tomorrow.
  - Each of this possible resolutions is a 'state of the world', denoted s, and S is the set of all states of the world.
- A 'commodity' is defined by its physical description (e.g. umbrella), as well as the state in which that good is consumed (e.g. tomorrow is a rainy day).
  - Thus, "an umbrella on a rainy day" is a commodity; "an umbrella on a sunny day" is a different commodity.
- Do this for all L "physical goods" in all states S, and define the consumption space as $\mathbf{R}^{L\times S}$.
  - A consumption vector specifies quantities of all these goods.
    - "1.2 umbrellas if it rains", "0.7 umbrellas if it is sunny", and so on.

Uncertainty and Utility: Debreu's Approach

Two Periods

Decisions are made today but items are consumed tomorrow.

- List of all the ways in which uncertainty can be resolved tomorrow.
  - Each of this possible resolutions is a 'state of the world', denoted s, and S is the set of all states of the world.
- A 'commodity' is defined by its physical description (e.g. umbrella), as well as the state in which that good is consumed (e.g. tomorrow is a rainy day).
  - Thus, "an umbrella on a rainy day" is a commodity; "an umbrella on a sunny day" is a different commodity.
- Do this for all L "physical goods" in all states S, and define the consumption space as $\mathbf{R}^{L \times S}$.
  - A consumption vector specifies quantities of all these goods.
    - "1.2 umbrellas if it rains", "0.7 umbrellas if it is sunny", and so on.
- If a preference relation over this consumption space is complete, transitive, and continuous, then there exist a utility function representing it.

Uncertainty and Utility: Debreu's Approach

Two Periods

Decisions are made today but items are consumed tomorrow.

- List of all the ways in which uncertainty can be resolved tomorrow.
  - Each of this possible resolutions is a 'state of the world', denoted s, and S is the set of all states of the world.
- A 'commodity' is defined by its physical description (e.g. umbrella), as well as the state in which that good is consumed (e.g. tomorrow is a rainy day).
  - Thus, "an umbrella on a rainy day" is a commodity; "an umbrella on a sunny day" is a different commodity.
- Do this for all L "physical goods" in all states S, and define the consumption space as $\mathbf{R}^{L\times S}$.
  - A consumption vector specifies quantities of all these goods.
    - "1.2 umbrellas if it rains", "0.7 umbrellas if it is sunny", and so on.
- If a preference relation over this consumption space is complete, transitive, and continuous, then there exist a utility function representing it.
  - Wether or not goods are deterministic does not change that theory.

Uncertainty and Utility: Debreu's Approach

Two Periods

Decisions are made today but items are consumed tomorrow.

- List of all the ways in which uncertainty can be resolved tomorrow.
  - Each of this possible resolutions is a 'state of the world', denoted s, and S is the set of all states of the world.
- A 'commodity' is defined by its physical description (e.g. umbrella), as well as the state in which that good is consumed (e.g. tomorrow is a rainy day).
  - Thus, "an umbrella on a rainy day" is a commodity; "an umbrella on a sunny day" is a different commodity.
- Do this for all L "physical goods" in all states S, and define the consumption space as $\mathbf{R}^{L \times S}$.
  - A consumption vector specifies quantities of all these goods.
    - "1.2 umbrellas if it rains", "0.7 umbrellas if it is sunny", and so on.
- If a preference relation over this consumption space is complete, transitive, and continuous, then there exist a utility function representing it.
  - Wether or not goods are deterministic does not change that theory.
- This approach is formally correct, yet hard to use in applications.

Uncertainty and Utility: Debreu's Approach

Two Periods

Decisions are made today but items are consumed tomorrow.

- List of all the ways in which uncertainty can be resolved tomorrow.
  - Each of this possible resolutions is a 'state of the world', denoted s, and S is the set of all states of the world.
- A 'commodity' is defined by its physical description (e.g. umbrella), as well as the state in which that good is consumed (e.g. tomorrow is a rainy day).
  - Thus, "an umbrella on a rainy day" is a commodity; "an umbrella on a sunny day" is a different commodity.
- Do this for all L "physical goods" in all states S, and define the consumption space as $\mathbf{R}^{L \times S}$.
  - A consumption vector specifies quantities of all these goods.
    - "1.2 umbrellas if it rains", "0.7 umbrellas if it is sunny", and so on.
- If a preference relation over this consumption space is complete, transitive, and continuous, then there exist a utility function representing it.
  - Wether or not goods are deterministic does not change that theory.
- This approach is formally correct, yet hard to use in applications.
- Furthermore, it hides how decision making may depend on states' likelihoods.

Uncertainty and Utility: Debreu's Approach

Two Periods

Decisions are made today but items are consumed tomorrow.

- List of all the ways in which uncertainty can be resolved tomorrow.
  - Each of this possible resolutions is a 'state of the world', denoted s, and S is the set of all states of the world.
- A 'commodity' is defined by its physical description (e.g. umbrella), as well as the state in which that good is consumed (e.g. tomorrow is a rainy day).
  - Thus, "an umbrella on a rainy day" is a commodity; "an umbrella on a sunny day" is a different commodity.
- Do this for all L "physical goods" in all states S, and define the consumption space as $\mathbf{R}^{L\times S}$.
  - A consumption vector specifies quantities of all these goods.
    - "1.2 umbrellas if it rains", "0.7 umbrellas if it is sunny", and so on.
- If a preference relation over this consumption space is complete, transitive, and continuous, then there exist a utility function representing it.
  - Wether or not goods are deterministic does not change that theory.
- This approach is formally correct, yet hard to use in applications.
- Furthermore, it hides how decision making may depend on states' likelihoods.
  - When rain is more likely, "umbrella when it rains" should be more attractive.

Lotteries and Probabilities

A box contains green and red balls. One ball will be drawn from this urn.

Lotteries and Probabilities

A box contains green and red balls. One ball will be drawn from this urn.

- Urn I

Green	Red
$G = 60$	$R = 40$

Lotteries and Probabilities

A box contains green and red balls. One ball will be drawn from this urn.

- Urn I

Green	Red
$G = 60$	$R = 40$

- Consider these "lotteries":

Lotteries and Probabilities

A box contains green and red balls. One ball will be drawn from this urn.

- Urn I

Green	Red
$G = 60$	$R = 40$

- Consider these "lotteries":
  - A: you win \$10 if a Green ball is drawn and zero otherwise

Lotteries and Probabilities

A box contains green and red balls. One ball will be drawn from this urn.

- Urn I

Green	Red
$G = 60$	$R = 40$

- Consider these "lotteries":
  - A: you win \$10 if a Green ball is drawn and zero otherwise
  - B: you win \$10 if a Red ball is drawn and zero otherwise

Lotteries and Probabilities

A box contains green and red balls. One ball will be drawn from this urn.

- Urn I

Green	Red
$G = 60$	$R = 40$

- Consider these "lotteries":
  - A: you win \$10 if a Green ball is drawn and zero otherwise
  - B: you win \$10 if a Red ball is drawn and zero otherwise
- Would you rather have A or B? Why?

Lotteries and Probabilities

A box contains green and red balls. One ball will be drawn from this urn.

- Urn I

Green	Red
$G = 60$	$R = 40$

- Consider these "lotteries":
  - A: you win \$10 if a Green ball is drawn and zero otherwise
  - B: you win \$10 if a Red ball is drawn and zero otherwise
- Would you rather have A or B? Why?

- Urn II

Green	Red
$G = 40$	$R = 60$

- Consider these "lotteries":
  - C: you win \$10 if a Green ball is drawn and zero otherwise
  - D: you win \$10 if a Red ball is drawn and zero otherwise
- Would you rather have C or D? Why?

Lotteries and Probabilities

A box contains green and red balls. One ball will be drawn from this urn.

- Urn I

Green	Red
$G = 60$	$R = 40$

- Consider these "lotteries":
  - A: you win \$10 if a Green ball is drawn and zero otherwise
  - B: you win \$10 if a Red ball is drawn and zero otherwise
- Would you rather have A or B? Why?

- Urn II

Green	Red
$G = 40$	$R = 60$

- Consider these "lotteries":
  - C: you win \$10 if a Green ball is drawn and zero otherwise
  - D: you win \$10 if a Red ball is drawn and zero otherwise
- Would you rather have C or D? Why?

- The probability of each outcome influences preferences.

Lotteries and Probabilities

A box contains green and red balls. One ball will be drawn from this urn.

- Urn I

Green	Red
$G = 60$	$R = 40$

- Consider these “lotteries”:
  - E: you win $10 if a Green ball is drawn and zero otherwise
  - F: you win $100 if a Red ball is drawn and zero otherwise
- Would you rather have E or F? Why?

Lotteries and Probabilities

A box contains green and red balls. One ball will be drawn from this urn.

- Urn I

Green	Red
$G = 60$	$R = 40$

- Consider these "lotteries":
  - E: you win \$10 if a Green ball is drawn and zero otherwise
  - F: you win \$100 if a Red ball is drawn and zero otherwise
- Would you rather have E or F? Why?

- Urn II

Green	Red
$G = 40$	$R = 60$

- Consider these "lotteries":
  - G: you win \$100 if a Green ball is drawn and zero otherwise
  - H: you win \$10 if a Red ball is drawn and zero otherwise
- Would you rather have G or H? Why?

Lotteries and Probabilities

A box contains green and red balls. One ball will be drawn from this urn.

- Urn I

Green	Red
$G = 60$	$R = 40$

- Consider these "lotteries":
    - E: you win $10 if a Green ball is drawn and zero otherwise
    - F: you win $100 if a Red ball is drawn and zero otherwise
- Would you rather have E or F? Why?

- Urn II

Green	Red
$G = 40$	$R = 60$

- Consider these "lotteries":
    - G: you win $100 if a Green ball is drawn and zero otherwise
    - H: you win $10 if a Red ball is drawn and zero otherwise
- Would you rather have G or H? Why?

- The "consequence" of each outcome also influences preferences.

Lotteries and Probabilities

A box contains green and red balls. One ball will be drawn from this urn.

- Urn I

Green	Red
$G = 60$	$R = 40$

- Consider these "lotteries":
  - E: you win $10 if a Green ball is drawn and zero otherwise
  - F: you win $100 if a Red ball is drawn and zero otherwise
- Would you rather have E or F? Why?

- Urn II

Green	Red
$G = 40$	$R = 60$

- Consider these "lotteries":
  - G: you win $100 if a Green ball is drawn and zero otherwise
  - H: you win $10 if a Red ball is drawn and zero otherwise
- Would you rather have G or H? Why?

- The "consequence" of each outcome also influences preferences.

Conjecture

Is the expected value of each lottery a reasonable utility function? Does it describe how much one is willing to pay for it?

Expected Values and Probabilities

St. Petersburg Paradox

Toss a fair coin until the first Head obtains.

Expected Values and Probabilities

St. Petersburg Paradox

Toss a fair coin until the first Head obtains.

- consider a gamble that pays $\$2^t$ if the first head appears on the t-th toss

Expected Values and Probabilities

St. Petersburg Paradox

Toss a fair coin until the first Head obtains.

- consider a gamble that pays \$$2^t$ if the first head appears on the t-th toss
  - \$2 if H on the 1^{st} toss, \$4 if H on the 2^{sd}, \$8 if H on the 3^{rd},...

Expected Values and Probabilities

St. Petersburg Paradox

Toss a fair coin until the first Head obtains.

- consider a gamble that pays \$$2^t$ if the first head appears on the t-th toss
  - \$2 if H on the 1^{st} toss, \$4 if H on the 2^{sd}, \$8 if H on the 3^{rd},...
- What is the value (utility) of this gamble? Would the expected gains work?

Expected Values and Probabilities

St. Petersburg Paradox

Toss a fair coin until the first Head obtains.

- consider a gamble that pays \$$2^t$ if the first head appears on the t-th toss
  - \$2 if H on the 1^{st} toss, \$4 if H on the 2^{sd}, \$8 if H on the 3^{rd},...
- What is the value (utility) of this gamble? Would the expected gains work?

$$\frac{1}{2}2 + \frac{1}{4}4 + \frac{1}{8}8 + ... = \infty$$

Expected Values and Probabilities

St. Petersburg Paradox

Toss a fair coin until the first Head obtains.

- consider a gamble that pays \$$2^t$ if the first head appears on the t-th toss
  - \$2 if H on the 1^{st} toss, \$4 if H on the 2^{sd}, \$8 if H on the 3^{rd},...
- What is the value (utility) of this gamble? Would the expected gains work?

$$\frac{1}{2}2 + \frac{1}{4}4 + \frac{1}{8}8 + \ldots = \infty$$

  - is anyone willing to pay that much for this lottery?

Expected Values and Probabilities

St. Petersburg Paradox

Toss a fair coin until the first Head obtains.

- consider a gamble that pays \$$2^t$ if the first head appears on the t-th toss
  - \$2 if H on the 1^{st} toss, \$4 if H on the 2^{sd}, \$8 if H on the 3^{rd},...
- What is the value (utility) of this gamble? Would the expected gains work?

$$\frac{1}{2}2 + \frac{1}{4}4 + \frac{1}{8}8 + ... = \infty$$

  - is anyone willing to pay that much for this lottery?
- As a solution, Bernoulli suggested using $\log 2^t$ to evaluate it:

Expected Values and Probabilities

St. Petersburg Paradox

Toss a fair coin until the first Head obtains.

- consider a gamble that pays \$$2^t$ if the first head appears on the t-th toss
  - \$2 if H on the 1^{st} toss, \$4 if H on the 2^{sd}, \$8 if H on the 3^{rd},...
- What is the value (utility) of this gamble? Would the expected gains work?

$$\frac{1}{2}2 + \frac{1}{4}4 + \frac{1}{8}8 + ... = \infty$$

  - is anyone willing to pay that much for this lottery?
- As a solution, Bernoulli suggested using $\log 2^t$ to evaluate it:

$$\frac{1}{2}\log 2 + \frac{1}{4}\log 4 + \frac{1}{8}\log 8 + ... \text{ a finite number}$$

Expected Values and Probabilities

St. Petersburg Paradox

Toss a fair coin until the first Head obtains.

- consider a gamble that pays $\$2^t$ if the first head appears on the t-th toss
  - \$2 if H on the 1^{st} toss, \$4 if H on the 2^{sd}, \$8 if H on the 3^{rd},...
- What is the value (utility) of this gamble? Would the expected gains work?

$$\frac{1}{2}2 + \frac{1}{4}4 + \frac{1}{8}8 + \ldots = \infty$$

  - is anyone willing to pay that much for this lottery?
- As a solution, Bernoulli suggested using $\log 2^t$ to evaluate it:

$$\frac{1}{2}\log 2 + \frac{1}{4}\log 4 + \frac{1}{8}\log 8 + \ldots \text{ a finite number}$$

- This solution works for this gamble... but how about a gamble that pays e^{2t} for H on the t-th toss?

Expected Values and Probabilities

St. Petersburg Paradox

Toss a fair coin until the first Head obtains.

- consider a gamble that pays $\$2^t$ if the first head appears on the t-th toss
  - \$2 if H on the 1^{st} toss, \$4 if H on the 2^{sd}, \$8 if H on the 3^{rd},...
- What is the value (utility) of this gamble? Would the expected gains work?

$$\frac{1}{2}2 + \frac{1}{4}4 + \frac{1}{8}8 + ... = \infty$$

  - is anyone willing to pay that much for this lottery?
- As a solution, Bernoulli suggested using $\log 2^t$ to evaluate it:

$$\frac{1}{2}\log 2 + \frac{1}{4}\log 4 + \frac{1}{8}\log 8 + ... \text{ a finite number}$$

- This solution works for this gamble... but how about a gamble that pays e^{2t} for H on the t-th toss?

$$\frac{1}{2}\log e^2 + \frac{1}{4}\log e^4 + \frac{1}{8}\log e^8 + ... = \infty$$

Expected Values and Probabilities

St. Petersburg Paradox

Toss a fair coin until the first Head obtains.

- consider a gamble that pays $\$2^t$ if the first head appears on the t-th toss
  - \$2 if H on the 1^{st} toss, \$4 if H on the 2^{sd}, \$8 if H on the 3^{rd},...
- What is the value (utility) of this gamble? Would the expected gains work?

$$\frac{1}{2}2 + \frac{1}{4}4 + \frac{1}{8}8 + ... = \infty$$

  - is anyone willing to pay that much for this lottery?
- As a solution, Bernoulli suggested using $\log 2^t$ to evaluate it:

$$\frac{1}{2}\log 2 + \frac{1}{4}\log 4 + \frac{1}{8}\log 8 + ... \text{ a finite number}$$

- This solution works for this gamble... but how about a gamble that pays e^{2t} for H on the t-th toss?

$$\frac{1}{2}\log e^2 + \frac{1}{4}\log e^4 + \frac{1}{8}\log e^8 + ... = \infty$$

  - we are back at the paradox.

Expected Values and Probabilities

St. Petersburg Paradox

Toss a fair coin until the first Head obtains.

- consider a gamble that pays $\$2^t$ if the first head appears on the t-th toss
  - \$2 if H on the 1^{st} toss, \$4 if H on the 2^{sd}, \$8 if H on the 3^{rd},...
- What is the value (utility) of this gamble? Would the expected gains work?

$$\frac{1}{2}2 + \frac{1}{4}4 + \frac{1}{8}8 + \ldots = \infty$$

  - is anyone willing to pay that much for this lottery?
- As a solution, Bernoulli suggested using $\log 2^t$ to evaluate it:

$$\frac{1}{2}\log 2 + \frac{1}{4}\log 4 + \frac{1}{8}\log 8 + \ldots \text{ a finite number}$$

- This solution works for this gamble... but how about a gamble that pays e^{2t} for H on the t-th toss?

$$\frac{1}{2}\log e^2 + \frac{1}{4}\log e^4 + \frac{1}{8}\log e^8 + \ldots = \infty$$

  - we are back at the paradox.
- Expected gains do not seeem to capture how we feel about it.

Expected Utility (imprecisely described)

Expected Utility (imprecisely described)

- Suppose there is only one good: call it money.

Expected Utility (imprecisely described)

- Suppose there is only one good: call it money.
- The state space consists of S mutually exclusive "states of the world";

Expected Utility (imprecisely described)

- Suppose there is only one good: call it money.
- The state space consists of S mutually exclusive "states of the world";
  - a generic element is denoted $s \in S$.

Expected Utility (imprecisely described)

- Suppose there is only one good: call it money.
- The state space consists of S mutually exclusive "states of the world";
  - a generic element is denoted $s \in S$.
- Let $\pi = (\pi_1, ..., \pi_S)$ be a probability distribution over S:

$$1 \geq \pi_s \geq 0 \qquad \text{and} \qquad \sum_{s \in S} \pi_s = 1$$

Expected Utility (imprecisely described)

- Suppose there is only one good: call it money.
- The state space consists of S mutually exclusive "states of the world";
  - a generic element is denoted $s \in S$.
- Let $\pi = (\pi_1, ..., \pi_S)$ be a probability distribution over S:

$$1 \geq \pi_s \geq 0 \qquad \text{and} \qquad \sum_{s \in S} \pi_s = 1$$

- A consumption bundle $x = x_1, ..., x_S$ is a random variable ($x \in \mathbf{R}^S$).

Expected Utility (imprecisely described)

- Suppose there is only one good: call it money.
- The state space consists of S mutually exclusive "states of the world";
  - a generic element is denoted $s \in S$.
- Let $\pi = (\pi_1, ..., \pi_S)$ be a probability distribution over S:

$$1 \geq \pi_s \geq 0 \qquad \text{and} \qquad \sum_{s \in S} \pi_s = 1$$

- A consumption bundle $x = x_1, ..., x_S$ is a random variable ($x \in \mathbf{R}^S$).
  - the agent "consumes" the amount x_s if and only if state s happens.

Expected Utility (imprecisely described)

- Suppose there is only one good: call it money.
- The state space consists of S mutually exclusive "states of the world";
  - a generic element is denoted $s \in S$.
- Let $\pi = (\pi_1, ..., \pi_S)$ be a probability distribution over S:

$$1 \geq \pi_s \geq 0 \qquad \text{and} \qquad \sum_{s \in S} \pi_s = 1$$

- A consumption bundle $x = x_1, ..., x_S$ is a random variable ($x \in \mathbf{R}^S$).
  - the agent "consumes" the amount x_s if and only if state s happens.
- Consumption bundles are evaluated with the expected value of their utility

$$U(x) = \sum_{s=1}^{S} \pi_s u(x_s).$$

Expected Utility (imprecisely described)

- Suppose there is only one good: call it money.
- The state space consists of S mutually exclusive "states of the world";
  - a generic element is denoted $s \in S$.
- Let $\pi = (\pi_1, ..., \pi_S)$ be a probability distribution over S:

$$1 \geq \pi_s \geq 0 \qquad \text{and} \qquad \sum_{s \in S} \pi_s = 1$$

- A consumption bundle $x = x_1, ..., x_S$ is a random variable ($x \in \mathbf{R}^S$).
  - the agent "consumes" the amount x_s if and only if state s happens.
- Consumption bundles are evaluated with the expected value of their utility

$$U(x) = \sum_{s=1}^{S} \pi_s u(x_s).$$

  - This function 'separates' probability from utility of consumption.

Expected Utility (imprecisely described)

- Suppose there is only one good: call it money.
- The state space consists of S mutually exclusive "states of the world";
  - a generic element is denoted $s \in S$.
- Let $\pi = (\pi_1, ..., \pi_S)$ be a probability distribution over S:

$$1 \geq \pi_s \geq 0 \qquad \text{and} \qquad \sum_{s \in S} \pi_s = 1$$

- A consumption bundle $x = x_1, ..., x_S$ is a random variable ($x \in \mathbf{R}^S$).
  - the agent "consumes" the amount x_s if and only if state s happens.
- Consumption bundles are evaluated with the expected value of their utility

$$U(x) = \sum_{s=1}^{S} \pi_s u(x_s).$$

  - This function 'separates' probability from utility of consumption.
  - The utility function $U(\cdot)$ is linear in probabilities and 'utils of consequences'.

Expected Utility (imprecisely described)

- Suppose there is only one good: call it money.
- The state space consists of S mutually exclusive "states of the world";
  - a generic element is denoted $s \in S$.
- Let $\pi = (\pi_1, ..., \pi_S)$ be a probability distribution over S:

$$1 \geq \pi_s \geq 0 \qquad \text{and} \qquad \sum_{s \in S} \pi_s = 1$$

- A consumption bundle $x = x_1, ..., x_S$ is a random variable ($x \in \mathbf{R}^S$).
  - the agent "consumes" the amount x_s if and only if state s happens.
- Consumption bundles are evaluated with the expected value of their utility

$$U(x) = \sum_{s=1}^{S} \pi_s u(x_s).$$

  - This function 'separates' probability from utility of consumption.
  - The utility function $U(\cdot)$ is linear in probabilities and 'utils of consequences'.
  - If we let $u(x) = (u(x_1), ..., u(x_S))$, we can write this as $\pi \cdot u(x)$.

Expected Utility (imprecisely described)

- Suppose there is only one good: call it money.
- The state space consists of S mutually exclusive "states of the world";
  - a generic element is denoted $s \in S$.
- Let $\pi = (\pi_1, ..., \pi_S)$ be a probability distribution over S:

$$1 \geq \pi_s \geq 0 \qquad \text{and} \qquad \sum_{s \in S} \pi_s = 1$$

- A consumption bundle $x = x_1, ..., x_S$ is a random variable ($x \in \mathbf{R}^S$).
  - the agent "consumes" the amount x_s if and only if state s happens.
- Consumption bundles are evaluated with the expected value of their utility

$$U(x) = \sum_{s=1}^{S} \pi_s u(x_s).$$

  - This function 'separates' probability from utility of consumption.
  - The utility function $U(\cdot)$ is linear in probabilities and 'utils of consequences'.
  - If we let $u(x) = (u(x_1), ..., u(x_S))$, we can write this as $\pi \cdot u(x)$.
- Decision making is connected to the theory of probability developed in mathematics and statistics.

Expected Utility (imprecisely described)

- Suppose there is only one good: call it money.
- The state space consists of S mutually exclusive "states of the world";
  - a generic element is denoted $s \in S$.
- Let $\pi = (\pi_1, ..., \pi_S)$ be a probability distribution over S:

$$1 \geq \pi_s \geq 0 \qquad \text{and} \qquad \sum_{s \in S} \pi_s = 1$$

- A consumption bundle $x = x_1, ..., x_S$ is a random variable ($x \in \mathbf{R}^S$).
  - the agent "consumes" the amount x_s if and only if state s happens.
- Consumption bundles are evaluated with the expected value of their utility

$$U(x) = \sum_{s=1}^{S} \pi_s u(x_s).$$

  - This function 'separates' probability from utility of consumption.
  - The utility function $U(\cdot)$ is linear in probabilities and 'utils of consequences'.
  - If we let $u(x) = (u(x_1), ..., u(x_S))$, we can write this as $\pi \cdot u(x)$.
- Decision making is connected to the theory of probability developed in mathematics and statistics.
  - An advantage of this approach is that we can use the rules of probability theory to evaluate how information about different events enters the decision making process (this is handy for game theory and information economics).

Convex Consumption Space

- To obtain representations similar to the one in the previous slide, we need more assumptions than completeness, transitivity, and continuity.
    - This is intuitive since the utility function has a very special functional form.
- We also neet to modify the space over which consumption is defined.

Convex Consumption Space

- To obtain representations similar to the one in the previous slide, we need more assumptions than completeness, transitivity, and continuity.
  - This is intuitive since the utility function has a very special functional form.
- We also neet to modify the space over which consumption is defined.

Consumption Set is Convex

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

Convex Consumption Space

- To obtain representations similar to the one in the previous slide, we need more assumptions than completeness, transitivity, and continuity.
  - This is intuitive since the utility function has a very special functional form.
- We also neet to modify the space over which consumption is defined.

Consumption Set is Convex

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

- Convexity means that

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

Convex Consumption Space

- To obtain representations similar to the one in the previous slide, we need more assumptions than completeness, transitivity, and continuity.
  - This is intuitive since the utility function has a very special functional form.
- We also neet to modify the space over which consumption is defined.

Consumption Set is Convex

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

- Convexity means that

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

- Start with preferences over an abstract convex space.

Convex Consumption Space

- To obtain representations similar to the one in the previous slide, we need more assumptions than completeness, transitivity, and continuity.
    - This is intuitive since the utility function has a very special functional form.
- We also neet to modify the space over which consumption is defined.

Consumption Set is Convex

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

- Convexity means that

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

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

Convex Consumption Space

- To obtain representations similar to the one in the previous slide, we need more assumptions than completeness, transitivity, and continuity.
  - This is intuitive since the utility function has a very special functional form.
- We also neet to modify the space over which consumption is defined.

Consumption Set is Convex

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

- Convexity means that

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

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

Definition

A binary relation $\succsim$ on Π is:

Convex Consumption Space

- To obtain representations similar to the one in the previous slide, we need more assumptions than completeness, transitivity, and continuity.
  - This is intuitive since the utility function has a very special functional form.
- We also neet to modify the space over which consumption is defined.

Consumption Set is Convex

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

- Convexity means that

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

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

Definition

A binary relation $\succsim$ on Π is:

- complete if, for all $x, y \in \Pi$, $x \succsim y$ or $y \succsim x$, or both;

Convex Consumption Space

- To obtain representations similar to the one in the previous slide, we need more assumptions than completeness, transitivity, and continuity.
  - This is intuitive since the utility function has a very special functional form.
- We also neet to modify the space over which consumption is defined.

Consumption Set is Convex

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

- Convexity means that

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

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

Definition

A binary relation $\succsim$ on Π is:

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

Archimedean Axiom

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

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 $\pi, \rho, \sigma \in \Pi$,

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

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 $\pi, \rho, \sigma \in \Pi$,

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

Exercise

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

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 $\pi, \rho, \sigma \in \Pi$,

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

Exercise

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

Exercise

Let $\Pi = \mathbf{R}$ and let $\succsim$ on $\mathbf{R}$ 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 $x, y, z \in \Pi$ and $\alpha \in (0,1)$,

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

Independence

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

Definition

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

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

Example

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

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

this says: x is weakly preferred to y whenever the norm of x is weakly larger than the norm of y.

Independence

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

Definition

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

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

Example

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

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

this says: x is weakly preferred to y whenever the norm of x is weakly larger than the norm of y.

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

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

Independence

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

Definition

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

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

Example

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

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

this says: x is weakly preferred to y whenever the norm of x is weakly larger than the norm of 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.

Independence

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

Definition

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

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

Example

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

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

this says: x is weakly preferred to y whenever the norm of x is weakly larger than the norm of 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 this imply geometrically?

Consequences of Independence

Proposition

If $\succsim$ satisfies independence its indifference classes are convex.

Consequences of Independence

Proposition

If $\succsim$ satisfies independence its indifference classes are convex.

Proof.

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

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

Consequences of Independence

Proposition

If $\succsim$ satisfies independence its indifference classes are convex.

Proof.

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

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

- If $\pi \sim \sigma$, clearly $\pi \succsim \sigma$. Thus, by independence, for all $\alpha \in [0,1]$,

$$\pi \succsim \sigma \Rightarrow \alpha\pi + (1-\alpha)\pi \succsim \alpha\sigma + (1-\alpha)\pi \Rightarrow \pi \succsim \alpha\sigma + (1-\alpha)\pi.$$

Consequences of Independence

Proposition

If $\succsim$ satisfies independence its indifference classes are convex.

Proof.

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

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

- If $\pi \sim \sigma$, clearly $\pi \succsim \sigma$. Thus, by independence, for all $\alpha \in [0,1]$,
$$\pi \succsim \sigma \Rightarrow \alpha\pi + (1-\alpha)\pi \succsim \alpha\sigma + (1-\alpha)\pi \Rightarrow \pi \succsim \alpha\sigma + (1-\alpha)\pi.$$
- If $\pi \sim \sigma$, clearly $\sigma \succsim \pi$. Thus, by independence, for all $\alpha \in [0,1]$,
$$\sigma \succsim \pi \ \Rightarrow \alpha\sigma + (1-\alpha)\pi \succsim \alpha\pi + (1-\alpha)\pi \ \Rightarrow \ \alpha\sigma + (1-\alpha)\pi \succsim \pi.$$

Consequences of Independence

Proposition

If $\succsim$ satisfies independence its indifference classes are convex.

Proof.

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

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

- If $\pi \sim \sigma$, clearly $\pi \succsim \sigma$. Thus, by independence, for all $\alpha \in [0,1]$,
$$\pi \succsim \sigma \Rightarrow \alpha\pi + (1-\alpha)\pi \succsim \alpha\sigma + (1-\alpha)\pi \Rightarrow \pi \succsim \alpha\sigma + (1-\alpha)\pi.$$
- If $\pi \sim \sigma$, clearly $\sigma \succsim \pi$. Thus, by independence, for all $\alpha \in [0,1]$,
$$\sigma \succsim \pi \;\Rightarrow \alpha\sigma + (1-\alpha)\pi \succsim \alpha\pi + (1-\alpha)\pi \;\Rightarrow\; \alpha\sigma + (1-\alpha)\pi \succsim \pi.$$
- Therefore we get $\pi \sim \alpha\sigma + (1-\alpha)\pi$ for all $\alpha \in [0,1]$.

Consequences of Independence

Proposition

If $\succsim$ satisfies independence its indifference classes are convex.

Proof.

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

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

- If $\pi \sim \sigma$, clearly $\pi \succsim \sigma$. Thus, by independence, for all $\alpha \in [0,1]$,
$$\pi \succsim \sigma \Rightarrow \alpha\pi + (1-\alpha)\pi \succsim \alpha\sigma + (1-\alpha)\pi \Rightarrow \pi \succsim \alpha\sigma + (1-\alpha)\pi.$$
- If $\pi \sim \sigma$, clearly $\sigma \succsim \pi$. Thus, by independence, for all $\alpha \in [0,1]$,
$$\sigma \succsim \pi \ \Rightarrow \alpha\sigma + (1-\alpha)\pi \succsim \alpha\pi + (1-\alpha)\pi \ \Rightarrow \ \alpha\sigma + (1-\alpha)\pi \succsim \pi.$$
- Therefore we get $\pi \sim \alpha\sigma + (1-\alpha)\pi$ for all $\alpha \in [0,1]$.
- The same logic shows that $\sigma \sim \alpha\sigma + (1-\alpha)\pi$.

Consequences of Independence

Proposition

If $\succsim$ satisfies independence its indifference classes are convex.

Proof.

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

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

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

Characterization of Independence

- 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 $\pi, \rho, \sigma \in \Pi$, and $\alpha \in (0,1)$,

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

and

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

Linear and Affine Functions

Definition

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

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

Linear and Affine Functions

Definition

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

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

Definition

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

$$f(\alpha\pi + \beta\rho) = \alpha f(\pi) + \beta f(\rho).$$

Linear and Affine Functions

Definition

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

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

Definition

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

$$f(\alpha\pi + \beta\rho) = \alpha f(\pi) + \beta f(\rho).$$

Exercise

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

Mixture Space Theorem

Theorem (Mixture Space Theorem, Herstein and Milnor)

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

$$\pi \quad \succsim \quad \rho \Leftrightarrow U(\pi) \geq U(\rho)$$

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

Mixture Space Theorem

Theorem (Mixture Space Theorem, Herstein and Milnor)

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

$$\pi \quad \succsim \quad \rho \Leftrightarrow U(\pi) \geq U(\rho)$$

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

- The first part of the statement says that the preferences are represented by an affine utility function.

Mixture Space Theorem

Theorem (Mixture Space Theorem, Herstein and Milnor)

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

$$\pi \quad \succsim \quad \rho \Leftrightarrow U(\pi) \geq U(\rho)$$

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

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

Mixture Space Theorem

Theorem (Mixture Space Theorem, Herstein and Milnor)

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

$$\pi \quad \succsim \quad \rho \Leftrightarrow U(\pi) \geq U(\rho)$$

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

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

Remarks

Mixture Space Theorem

Theorem (Mixture Space Theorem, Herstein and Milnor)

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

$$\pi \quad \succsim \quad \rho \Leftrightarrow U(\pi) \geq U(\rho)$$

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

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

Remarks

- This holds for any convex subset of an arbitrary vector space.

Mixture Space Theorem

Theorem (Mixture Space Theorem, Herstein and Milnor)

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

$$\pi \quad \succsim \quad \rho \Leftrightarrow U(\pi) \geq U(\rho)$$

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

- The first part of the statement says that the preferences are represented by an affine utility function.
- The second that the 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?).

Mixture Space Theorem

Theorem (Mixture Space Theorem, Herstein and Milnor)

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

$$\pi \succsim \rho \Leftrightarrow U(\pi) \geq U(\rho)$$

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

- The first part of the statement says that the preferences are represented by an affine utility function.
- The second that the 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* $\mathbf{R}^n$*) is complete, transitive, independent and Archimedean if and only if there exists an affine function* $U : \Pi \to \mathbf{R}$ *representating* $\succsim$*:*

$$\pi \succsim \rho \Leftrightarrow U(\pi) \geq U(\rho)$$

Moreover, if $U : \Pi \to \mathbf{R}$ *is an affine representation of* $\succsim$*, then* $U' : \Pi \to \mathbf{R}$ *is an affine representation of* $\succsim$ *if and only if there exist real numbers* $a > 0$ *and* b *such that* $U'(\pi) = aU(\pi) + b$ *for all* $\pi \in \Pi$*.*

Mixture Space Theorem

Theorem (Mixture Space Theorem, Herstein and Milnor)

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

$$\pi \succsim \rho \Leftrightarrow U(\pi) \geq U(\rho)$$

Moreover, if $U : \Pi \to \mathbf{R}$ is an affine representation of $\succsim$, then $U' : \Pi \to \mathbf{R}$ is an affine representation of $\succsim$ if and only if there exist real numbers $a > 0$ and b such that $U'(\pi) = aU(\pi) + b$ for all $\pi \in \Pi$.

- As with Debreu's utility representation proof, the main step is to find the unique number that is indifferent to a given $\pi \in \Pi$.

Mixture Space Theorem

Theorem (Mixture Space Theorem, Herstein and Milnor)

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

$$\pi \succsim \rho \Leftrightarrow U(\pi) \geq U(\rho)$$

Moreover, if $U : \Pi \to \mathbf{R}$ is an affine representation of $\succsim$, then $U' : \Pi \to \mathbf{R}$ is an affine representation of $\succsim$ if and only if there exist real numbers $a > 0$ and b such that $U'(\pi) = aU(\pi) + b$ for all $\pi \in \Pi$.

- As with Debreu's utility representation proof, the main step is to find the unique number that is indifferent to a given $\pi \in \Pi$.
  - The main difference is that the assumptions now imply an affine utility function (rather than continuous one as in Debreu's theorem).

Mixture Space Theorem

Theorem (Mixture Space Theorem, Herstein and Milnor)

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

$$\pi \succsim \rho \Leftrightarrow U(\pi) \geq U(\rho)$$

Moreover, if $U : \Pi \to \mathbf{R}$ is an affine representation of $\succsim$, then $U' : \Pi \to \mathbf{R}$ is an affine representation of $\succsim$ if and only if there exist real numbers $a > 0$ and b such that $U'(\pi) = aU(\pi) + b$ for all $\pi \in \Pi$.

- As with Debreu's utility representation proof, the main step is to find the unique number that is indifferent to a given $\pi \in \Pi$.
  - The main difference is that the assumptions now imply an affine utility function (rather than continuous one as in Debreu's theorem).
- Proof: Roee's class.

Mixture Space Theorem

Theorem (Mixture Space Theorem, Herstein and Milnor)

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

$$\pi \succsim \rho \Leftrightarrow U(\pi) \geq U(\rho)$$

Moreover, if $U : \Pi \to \mathbf{R}$ is an affine representation of $\succsim$, then $U' : \Pi \to \mathbf{R}$ is an affine representation of $\succsim$ if and only if there exist real numbers $a > 0$ and b such that $U'(\pi) = aU(\pi) + b$ for all $\pi \in \Pi$.

- As with Debreu's utility representation proof, the main step is to find the unique number that is indifferent to a given $\pi \in \Pi$.
    - The main difference is that the assumptions now imply an affine utility function (rather than continuous one as in Debreu's theorem).
- Proof: Roee's 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 be a finite set of size n and ΔX be the space of probabilities on X.

Preferences Over Lotteries

Von Neumann and Morgenstern (1947)

Let X be a finite set of size n and ΔX be the space of probabilities on X.

- Enumerate $X = \{x_1, x_2, \ldots, x_n\}$, and let $\pi(x_i) = \pi_i$ so that

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

Preferences Over Lotteries

Von Neumann and Morgenstern (1947)

Let X be a finite set of size n and ΔX be the space of probabilities on X.

- Enumerate $X = \{x_1, x_2, \ldots, x_n\}$, and let $\pi(x_i) = \pi_i$ so that
$$\Delta X = \left\{ \pi \in \mathbf{R}^n : \sum_{i=1}^{n} \pi_i = 1 \text{ and } \pi_i \geq 0, \forall i \right\}$$
- ΔX is a convex subset of the vector space $\mathbf{R}^n$.

Preferences Over Lotteries

Von Neumann and Morgenstern (1947)

Let X be a finite set of size n and ΔX be the space of probabilities on X.

- Enumerate $X = \{x_1, x_2, \ldots, x_n\}$, and let $\pi(x_i) = \pi_i$ so that
$$\Delta X = \left\{ \pi \in \mathbf{R}^n : \sum_{i=1}^{n} \pi_i = 1 \text{ and } \pi_i \geq 0, \forall i \right\}$$
- ΔX is a convex subset of the vector space $\mathbf{R}^n$.
    - An element of ΔX identifies a lottery over the elements of X (the 'prizes').

Preferences Over Lotteries

Von Neumann and Morgenstern (1947)

Let X be a finite set of size n and ΔX be the space of probabilities on X.

- Enumerate $X = \{x_1, x_2, \ldots, x_n\}$, and let $\pi(x_i) = \pi_i$ so that
$$\Delta X = \left\{ \pi \in \mathbf{R}^n : \sum_{i=1}^{n} \pi_i = 1 \text{ and } \pi_i \geq 0, \forall i \right\}$$
- ΔX is a convex subset of the vector space $\mathbf{R}^n$.
  - An element of ΔX identifies a lottery over the elements of X (the 'prizes').
    - One can write a lottery π also as $(\pi_1, x_1; \pi_2, x_2; \ldots; \pi_n, x_n)$.

Preferences Over Lotteries

Von Neumann and Morgenstern (1947)

Let X be a finite set of size n and ΔX be the space of probabilities on X.

- Enumerate $X = \{x_1, x_2, \ldots, x_n\}$, and let $\pi(x_i) = \pi_i$ so that
$$\Delta X = \left\{ \pi \in \mathbf{R}^n : \sum_{i=1}^{n} \pi_i = 1 \text{ and } \pi_i \geq 0, \forall i \right\}$$
- ΔX is a convex subset of the vector space $\mathbf{R}^n$.
  - An element of ΔX identifies a lottery over the elements of X (the 'prizes').
    - One can write a lottery π also as $(\pi_1, x_1; \pi_2, x_2; \ldots; \pi_n, x_n)$.
  - The degenerate lottery that yields some outcome x with certainty is called a Dirac lottery on x and denoted δ_x;

Preferences Over Lotteries

Von Neumann and Morgenstern (1947)

Let X be a finite set of size n and ΔX be the space of probabilities on X.

- Enumerate $X = \{x_1, x_2, \ldots, x_n\}$, and let $\pi(x_i) = \pi_i$ so that
$$\Delta X = \left\{ \pi \in \mathbf{R}^n : \sum_{i=1}^{n} \pi_i = 1 \text{ and } \pi_i \geq 0, \forall i \right\}$$
- ΔX is a convex subset of the vector space $\mathbf{R}^n$.
  - An element of ΔX identifies a lottery over the elements of X (the 'prizes').
    - One can write a lottery π also as $(\pi_1, x_1; \pi_2, x_2; \ldots; \pi_n, x_n)$.
  - The degenerate lottery that yields some outcome x with certainty is called a Dirac lottery on x and denoted δ_x;
    - δ_{x_k} is the unit vector in the direction k: $\delta_{x_k} = (0_1, \ldots, 0_{k-1}, 1_k, 0_{k+1}, \ldots, 0_n)$.

Preferences Over Lotteries

Von Neumann and Morgenstern (1947)

Let X be a finite set of size n and ΔX be the space of probabilities on X.

- Enumerate $X = \{x_1, x_2, \dots, x_n\}$, and let $\pi(x_i) = \pi_i$ so that
$$\Delta X = \left\{ \pi \in \mathbf{R}^n : \sum_{i=1}^{n} \pi_i = 1 \text{ and } \pi_i \geq 0, \forall i \right\}$$
- ΔX is a convex subset of the vector space $\mathbf{R}^n$.
  - An element of ΔX identifies a lottery over the elements of X (the 'prizes').
    - One can write a lottery π also as $(\pi_1, x_1; \pi_2, x_2; \dots; \pi_n, x_n)$.
  - The degenerate lottery that yields some outcome x with certainty is called a Dirac lottery on x and denoted δ_x;
    - δ_{x_k} is the unit vector in the direction k: $\delta_{x_k} = (0_1, \dots, 0_{k-1}, 1_k, 0_{k+1}, \dots, 0_n)$.
    - The set of Dirac lotteries is $\{\delta_x : x \in X\} \subset \Delta X$, and it constitutes the extreme points of ΔX (what are the extreme points?).

Preferences Over Lotteries

Von Neumann and Morgenstern (1947)

Let X be a finite set of size n and ΔX be the space of probabilities on X.

- Enumerate $X = \{x_1, x_2, \ldots, x_n\}$, and let $\pi(x_i) = \pi_i$ so that

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

- ΔX is a convex subset of the vector space $\mathbf{R}^n$.
  - An element of ΔX identifies a lottery over the elements of X (the 'prizes').
    - One can write a lottery π also as $(\pi_1, x_1; \pi_2, x_2; \ldots; \pi_n, x_n)$.
  - The degenerate lottery that yields some outcome x with certainty is called a Dirac lottery on x and denoted δ_x;
    - δ_{x_k} is the unit vector in the direction k: $\delta_{x_k} = (0_1, \ldots, 0_{k-1}, 1_k, 0_{k+1}, \ldots, 0_n)$.
    - The set of Dirac lotteries is $\{\delta_x : x \in X\} \subset \Delta X$, 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 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 $\pi = (\pi_1, x_1; \pi_2, x_2; \pi_3, x_3)$ is described using an event tree:

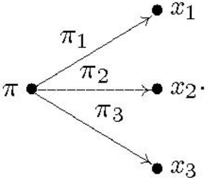

- Then δ_{x_2}, the degenerate lottery which yields x_2 with certainty, is:

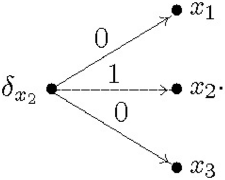

Lotteries

- If $X = \{x_1, x_2, x_3\}$, a typical lottery $\pi = (\pi_1, x_1; \pi_2, x_2; \pi_3, x_3)$ is described using an event tree:

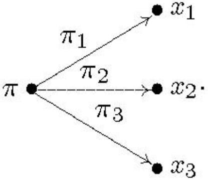

- Then δ_{x_2}, the degenerate lottery which yields x_2 with certainty, is:

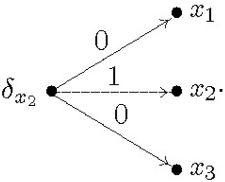

- The space ΔX assumes all uncertainty is resolved at one point in time; it does not allow for compound lotteries (lotteries over lotteries).

Lotteries

- If $X = \{x_1, x_2, x_3\}$, a typical lottery $\pi = (\pi_1, x_1; \pi_2, x_2; \pi_3, x_3)$ is described using an event tree:

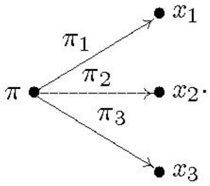

- Then δ_{x_2}, the degenerate lottery which yields x_2 with certainty, is:

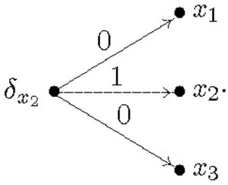

- The space ΔX assumes all uncertainty is resolved at one point in time; it does not allow for compound lotteries (lotteries over lotteries).
  - This domain restriction can be justified by introducing a 'reduction of compound lotteries' assumption as to reduce every compound lottery to a single lottery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

Compound Lotteries

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

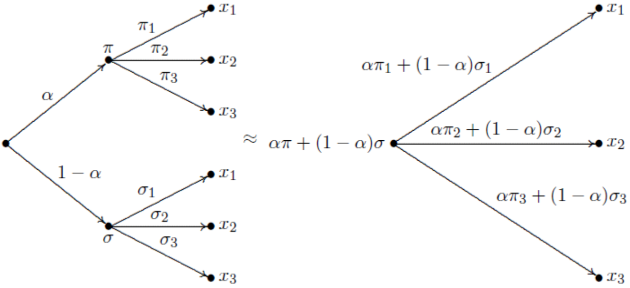

Compound Lotteries

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

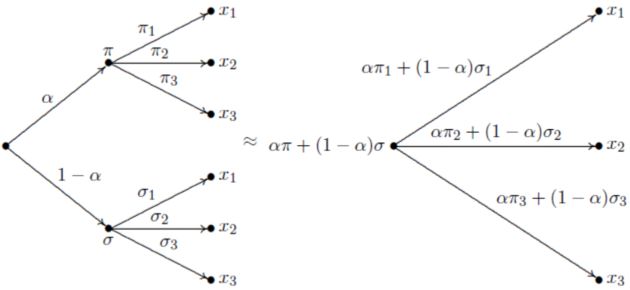

- One could assume the decision maker compounds correctly, but we cannot state this assumption since compound lotteries are not part of the primitives.

Compound Lotteries

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

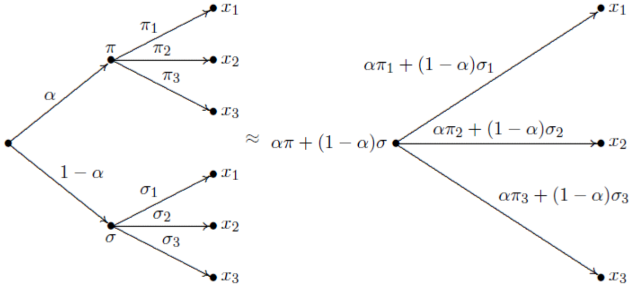

- One could assume the decision maker compounds correctly, but we cannot state this assumption since compound lotteries are not part of the primitives.
  - Suppose Z is a finite subset of ΔX. A lottery $\pi \in \Delta Z$ is a compound lottery, because it is a lottery over lotteries.

Compound Lotteries

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

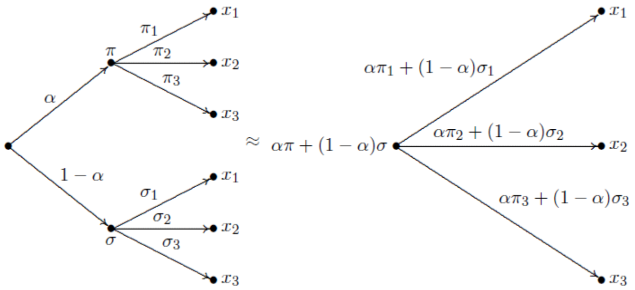

- One could assume the decision maker compounds correctly, but we cannot state this assumption since compound lotteries are not part of the primitives.
  - Suppose Z is a finite subset of ΔX. A lottery $\pi \in \Delta Z$ is a compound lottery, because it is a lottery over lotteries.
  - If one takes a $\pi \in Z$ then δ_π is an element of ΔX, a lottery over X.

Independence and Lotteries

independence: *for all* $\pi, \pi', \sigma \in \Delta X$ *and* $\alpha \in (0, 1)$,

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

Hence:

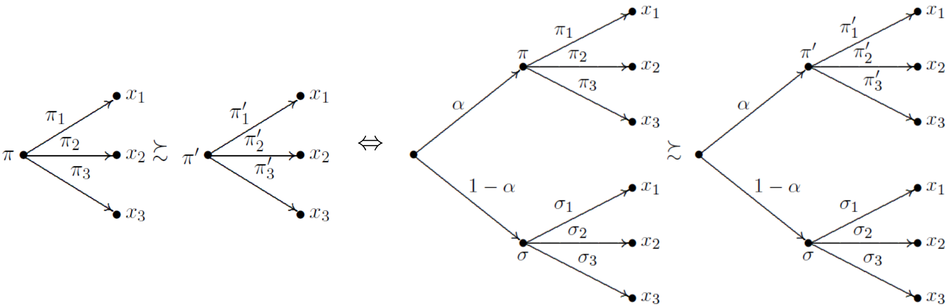

Independence and Lotteries

independence: *for all* $\pi, \pi', \sigma \in \Delta X$ *and* $\alpha \in (0,1)$,

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

Hence:

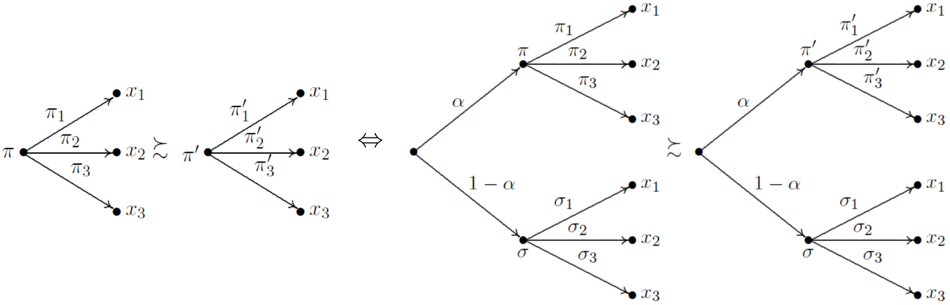

- The decision maker cares only about paths which differ.

Independence and Lotteries

independence: *for all* $\pi, \pi', \sigma \in \Delta X$ *and* $\alpha \in (0,1)$,

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

Hence:

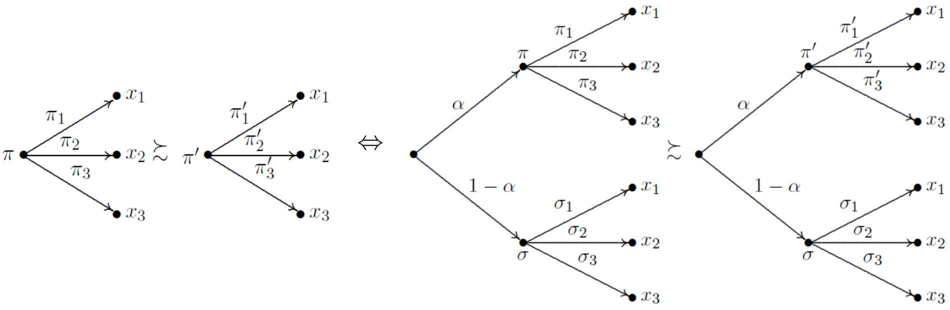

- The decision maker cares only about paths which differ.
- This is a 'normative' justification for the independence axiom on ΔX.

Expected Utility

Things we already know

- Under completeness, transitivity and continuity, there exists a continuous utility function representing the preferences.
- When X is convex, one can replace continuity with the archimedean axiom and add independence, and show that the utility function is affine.

Expected Utility

Things we already know

- Under completeness, transitivity and continuity, there exists a continuous utility function representing the preferences.
- When X is convex, one can replace continuity with the archimedean axiom and add independence, and show that the utility function is affine.

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

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

represents $\succsim$.

Expected Utility

Things we already know

- Under completeness, transitivity and continuity, there exists a continuous utility function representing the preferences.
- When X is convex, one can replace continuity with the archimedean axiom and add independence, and show that the utility function is affine.

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

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

  represents $\succsim$.
  - DM weights the utility of each outcome by the probability of receiving that outcome.

Expected Utility

Things we already know

- Under completeness, transitivity and continuity, there exists a continuous utility function representing the preferences.
- When X is convex, one can replace continuity with the archimedean axiom and add independence, and show that the utility function is affine.

- Under the extra structure given by ΔX, the representation theorem identifies a function $v : X \longrightarrow \mathbf{R}$ such that

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

  represents $\succsim$.
  - DM weights the utility of each outcome by the probability of receiving that outcome.
- Since the probability distribution over π is given, the theorem identifies via preferences the functional form of $U(\cdot)$ and the function $v(\cdot)$.

Expected Utility

Things we already know

- Under completeness, transitivity and continuity, there exists a continuous utility function representing the preferences.
- When X is convex, one can replace continuity with the archimedean axiom and add independence, and show that the utility function is affine.

- Under the extra structure given by ΔX, the representation theorem identifies a function $v : X \rightarrow \mathbf{R}$ such that
$$U(\pi) = \sum_{x \in X} \pi(x) v(x) = \sum_{i=1}^{n} \pi_i v(x_i)$$
represents $\succsim$.
  - DM weights the utility of each outcome by the probability of receiving that outcome.
- Since the probability distribution over π is given, the theorem identifies via preferences the functional form of $U(\cdot)$ and the function $v(\cdot)$.

Remark

- The function $v : X \rightarrow \mathbf{R}$ yields a vector $v \in \mathbf{R}^n$ by letting $v_i = v(x_i)$.
- The expected utility formula is the dot product of two vectors (v and π) in $\mathbf{R}^n$.

Expected Utility Theorem

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

Let ΔX be the set of all probability distributions 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.

Expected Utility Theorem

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

Let ΔX be the set of all probability distributions 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
$$\pi \succsim \rho \Leftrightarrow U(\pi) \geq U(\rho) \Leftrightarrow \sum_{x \in X} v(x)\pi(x) \geq \sum_{x \in X} v(x)\rho(x).$$

Expected Utility Theorem

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

Let ΔX be the set of all probability distributions 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
$$\pi \succsim \rho \Leftrightarrow U(\pi) \geq U(\rho) \Leftrightarrow \sum_{x \in X} v(x)\pi(x) \geq \sum_{x \in X} v(x)\rho(x).$$
- Uniqueness means: $U'(\pi) = \sum_x v'(x)\pi(x)$ 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 $x \in X$.

Expected Utility Theorem

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

Let ΔX be the set of all probability distributions 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
$$\pi \succsim \rho \Leftrightarrow U(\pi) \geq U(\rho) \Leftrightarrow \sum_{x \in X} v(x)\pi(x) \geq \sum_{x \in X} v(x)\rho(x).$$
- Uniqueness means: $U'(\pi) = \sum_x v'(x)\pi(x)$ 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 $x \in X$.
- The function v is called von Neumann & Morgenstern utility index or Bernoulli utility function.

Expected Utility Theorem

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

Let ΔX be the set of all probability distributions 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
$$\pi \succsim \rho \Leftrightarrow U(\pi) \geq U(\rho) \Leftrightarrow \sum_{x \in X} v(x)\pi(x) \geq \sum_{x \in X} v(x)\rho(x).$$
- Uniqueness means: $U'(\pi) = \sum_x v'(x)\pi(x)$ 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 $x \in X$.
- The function v is called von Neumann & Morgenstern utility index or Bernoulli utility function.

Remark

- $v : X \to \mathbf{R}$ is not a utility representation of $\succsim$; the domain of v is X, which is not equal to ΔX. The utility index v is a component of the utility representation U, which is defined on ΔX (the correct domain).

Next Class

- Proof Von Neumann & Morgentstern Expected UtilityTheorem
- Subjective vs. Objective Probability
- Anscombe and Aumann Acts
- State Independence