A brief introduction to economics
Part I

Tyler Moore

Computer Science & Engineering Department, SMU, Dallas, TX

Lecture 4

Notes

Why again are we studying economics?

- Economics is a *social science*
  - Studies behavior of individuals and firms in order to predict outcomes
  - Models of behavior based on systematic observation
  - Usually cannot run experiments as in bench science, but economics has developed ways to cope with differences inherent to observing the world
- Economics studies trade-offs between conflicting interests
  - Recognizes that people operate *strategically*
  - Have devised ways to model people's interests and decision making

Notes

Economics is not just about money

- Money helps to reveal preferences
- Money can serve as a common measure for costs and benefits
- As a discipline, economics examines much more than interactions involving money
  - Economics studies trade-offs between conflicting interests
  - Conflicting interests and incentives appear in many circumstances where money never changes hands

Notes

Notion of Model

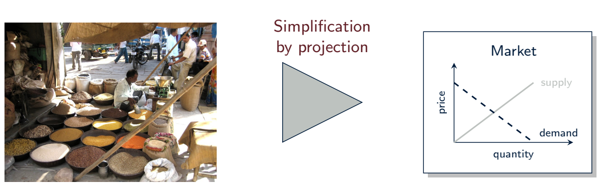

Reality

Model

- All models are wrong.
- Some are useful.

Notes

Types of models used in economics

1. **Analytical models**: state plausible assumptions about agent's behavior, then examine the implications
   + Good for theoretical analysis of individual behavior
   - When models disagree, ground truth can be elusive
2. **Empirical models**: observe relationships in aggregate, without explaining underlying individual decisions
   + Ground truth is achievable
   - Oversimplify, can't explain underlying mechanisms
3. **Measurement models**: collects data to compare deviations from predictions made by analytical models
   - Directly applying empirical analysis to analytical models usually fails
   + Offers feedback to analytical models to validate predictions

Notes

Model Complexity and Scientific Discovery

empirical observation

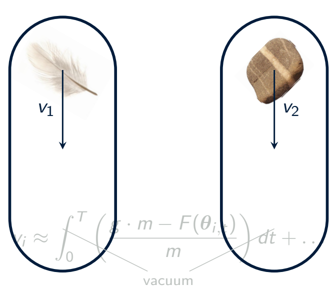

→ Drag is part of a complex modelReduction to simple model: drag causes measurement error

Notes

Model Complexity and Generalizability

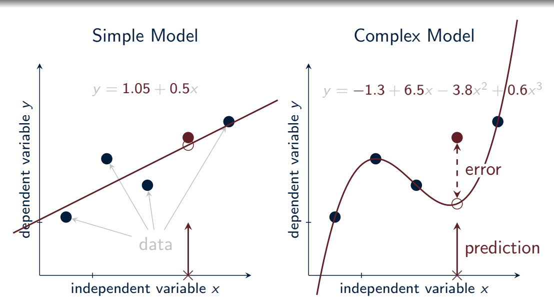

Measure of complexity for predictive models: number of estimated parameters

→ Risk of **overfitting** increases with model complexity

Notes

Trade-off on Model Complexity

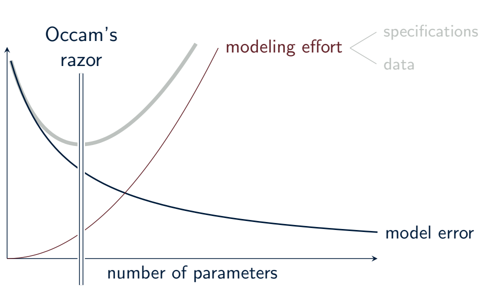

→ **William of Occam** († 1349)**: Principle of model parismony**

Notes

Occam's Razor

William of Occam, 1285–1349

entia non sunt multiplicanda praeter necessitatem

entities must not be multiplied beyond necessity

Notes

Our first model: rational choice theory

- Economics attempts to model the *decisions* we make, when faced with multiple choices and when interacting with other strategic agents
- Rational choice theory (RCT): model for decision-making
- Game theory (GT): extends RCT to model strategic interactions

Notes

Rationality defined

- Intuitive definition: a **rational** individual acts in his or her perceived best interest
- Rationality is what motivates a focus on *incentives*
- Question: can you think of scenarios when this definition does not hold in practice?
- To arrive at a precise definition: use rational choice theory to state available outcomes, articulate preferences among them, and decide accordingly

Notes

Model of preferences

- An agent is faced with a range of possible outcomes $o_1, o_2 \in \mathcal{O}$, the set of all possible outcomes
- Notation
  - $o_1 \succ o_2$: the agent is strictly prefers o_1 to o_2.
  - $o_1 \succeq o_2$: the agent weakly prefers o_1 to o_2;
  - $o_1 \sim o_2$: the agent is indifferent between o_1 and o_2;
- Outcomes can be also viewed as tuples of different properties $\hat{x}, \hat{y} \in \mathcal{O}$, where $\hat{x} = (x_1, x_2, \ldots, x_n)$ and $\hat{y} = (y_1, y_2, \ldots, y_n)$

Notes

Rational choice axioms

Rational choice theory assumes consistency in how outcomes are preferred.

Axiom

Completeness. For each pair of outcomes o_1 and o_2, exactly one of the following holds: $o_1 \succ o_2$, $o_1 \sim o_2$, or $o_2 \succ o_1$.

⇒ Outcomes can always be compared

Axiom

Transitivity. For each triple of outcomes o_1, o_2, and o_3, if $o_1 \succ o_2$ and $o_2 \succ o_3$, then $o_1 \succ o_3$.

⇒ People make choices among many different outcomes in a consistent manner

Notes

Example: trade-off between confidentiality and availability using cryptography

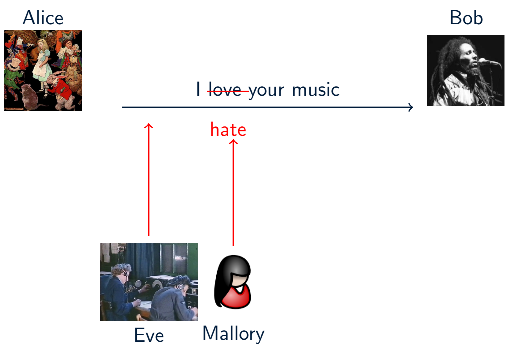

Notes

Example: trade-off between confidentiality and availability using cryptography

- Outcomes $\mathcal{O}$
  - $\mathfrak{c}_\oplus$: mechanism achieving high confidentiality
  - $\mathfrak{c}_\ominus$: mechanism achieving low confidentiality
  - $\mathfrak{a}_\oplus$: mechanism achieving high availability
  - $\mathfrak{a}_\ominus$: mechanism achieving low availability
- Preferences
  - $\mathfrak{c}_\oplus \succ \mathfrak{c}_\ominus$ and $\mathfrak{a}_\oplus \succ \mathfrak{a}_\ominus$
  - Taken together: $(\mathfrak{c}_\oplus, \mathfrak{a}_\oplus) \succ (\mathfrak{c}_\ominus, \mathfrak{a}_\ominus)$
  - Question: what about high availability and low confidentiality?
  - Indifferent: $(\mathfrak{c}_\oplus, \mathfrak{a}_\ominus) \sim (\mathfrak{c}_\ominus, \mathfrak{a}_\oplus)$.

Notes

Indifference curves

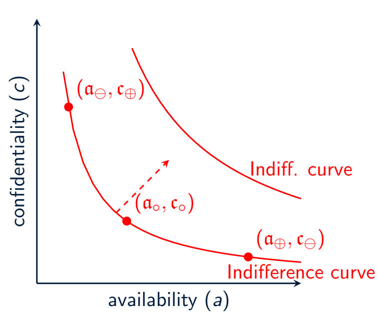

Notes

From preferences to utility

- It's great to express preferences, but to make mathematical analysis of decisions possible, we need to transform these preferences into numbers.
- We need a measure of utility, but what does that actually mean?

Notes

We do not mean utility according to Bentham

- Founder of utilitarianism: "fundamental axiom, it is the greatest happiness of the greatest number that is the measure of right and wrong"
- Utility: preferring "pleasure" over "pain"

Jeremy Bentham

Notes

Utility

Rational choice theory defines utility as a way of quantifying consumer preferences

Definition

(Utility function) A utility function U maps a set of outcomes onto real-valued numbers, that is, $U : \mathcal{O} \rightarrow \mathbb{R}$. U is defined such that $U(o_1) > U(o_2) \iff o_1 \succ o_2$.

Agents make a rational decision by picking the outcome with highest utility:

$$o^* = \arg\max_{o \in \mathcal{O}} U(o) \quad (1)$$

Notes

Example utility functions

- $U(o_1, o_2) = u \cdot o_1 + v \cdot o_2$
  - Useful when outcomes are *substitutes*
  - Example substitutes: processor speed and RAM
- $U(o_1, o_2) = \min\{u \cdot o_1, v \cdot o_2\}$
  - Useful when outcomes are *complements*
  - Example complements: operating system and third-party software

Notes

Returning to our crypto example

- First, we need a utility function
  - $U(\mathfrak{a}_i, \mathfrak{c}_i) = u \cdot \mathfrak{a}_i + v \cdot \mathfrak{c}_i$
  - Question: why is this a good choice?
- For simplicity, we assign $\mathfrak{a}_\oplus = 1$, $\mathfrak{a}_\ominus = -1$, $\mathfrak{c}_\oplus = 1$, and $\mathfrak{c}_\ominus = -1$
- Utility is in the eye of the beholder
- We consider two scenarios
  - Intelligence agency ($u = 1$ and $v = 3$)
  - First responders ($u = 3$ and $v = 1$)

Notes

Utility of different outcomes

Outcome	U_{FR} (first responder)	U_{intel} (intelligence)
$(\mathfrak{a}_\oplus, \mathfrak{c}_\oplus)$	4	4
$(\mathfrak{a}_\oplus, \mathfrak{c}_\ominus)$	2	-2
$(\mathfrak{a}_\ominus, \mathfrak{c}_\oplus)$	?	?
$(\mathfrak{a}_\ominus, \mathfrak{c}_\ominus)$	?	?

Notes

Why isn't utility theory enough?

- Only rarely do actions people take directly determine outcomes
- Instead there is uncertainty about which outcome will come to pass
- More realistic model: agent selects action a from set of all possible actions $\mathcal{A}$, and then outcomes $\mathcal{O}$ are associated with probability distribution

Notes

Lotteries

Definition

(Lottery) A lottery is a mapping from all outcomes $(o_1, o_2, \ldots, o_n) \in \mathcal{O}$ to probabilities corresponding to each outcome $(p_1, p_2, \ldots, p_n)$, where $\sum_1^n p_i = 1$. A lottery l_1 is represented as $l_1 = \langle o_1 \colon p_1, o_2 \colon p_2, \ldots, o_n \colon p_n \rangle$.

Notes

Where does randomness come from?

- Indeterminism in nature
- Lack of knowledge
- Incompleteness in the model
- Uncertainty concerns *which outcome* will occur
  - ⇒ Known unknowns, NOT unknown unknowns

Notes

Expected utility

Definition

(Expected utility (discrete)) The *expected utility* of an action $a \in \mathcal{A}$ is defined by adding up the utility for all outcomes weighed by their probability of occurrence:

$$E[U(a)] = \sum_{o \in \mathcal{O}} U(o) \cdot P(o|a) \tag{2}$$

Agents make a rational decision by maximizing expected utility:

$$a^* = \arg\max_{a \in \mathcal{A}} E[U(a)] \tag{3}$$

Notes

Example: process control system security

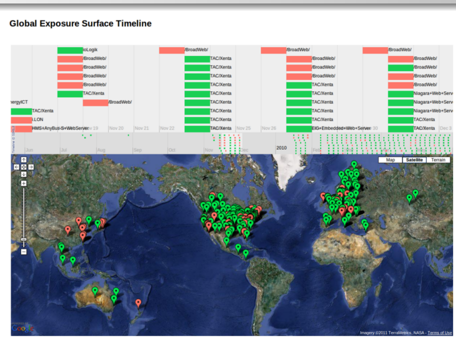

Figure 2.1: Example exposure time-map with red marking systems with known exploits

Source: `http://www.cl.cam.ac.uk/~fms27/papers/2011-Leverett-industrial.pdf`

Notes

Example: process control system security

- Actions available: $\mathcal{A} = \{\text{disconnect}, \text{connect}\}$
- Outcomes available: $\mathcal{O} = \{\text{attack}, \text{no attack}\}$
- If systems are connected, then the probability of successful attack is 0.01 ($P(\text{attack}|\text{connect}) = 0.01$)
- If systems are disconnected, then $P(\text{attack}|\text{disconnect}) = 0$

Notes

Example: process control system security

Action	attack U	attack $P(\text{attack}\mid\text{action})$	no attack U	no attack $P(\text{no attack}\mid\text{action})$	$E[U(\text{action})]$
disconnect	-100	0	5	1	?
connect	-100	0.01	10	0.99	?

$$E[U(a)] = \sum_{o \in \mathcal{O}} U(o) \cdot P(o|a)$$

$$\begin{aligned} E[U(\text{disconnect})] &= U(\text{attack}) \cdot P(\text{attack}|\text{disconnect}) \\ &+ U(\text{no attack}) \cdot P(\text{no attack}|\text{disconnect}) \\ &= -100 \cdot 0 + 5 \cdot 1 \\ &= 5 \end{aligned}$$

Notes

Example: process control system security

Action	attack U	attack $P(\text{attack}\mid\text{action})$	no attack U	no attack $P(\text{no attack}\mid\text{action})$	$E[U(\text{action})]$
disconnect	-100	0	5	1	5
connect	-100	0.01	10	0.99	?

$$E[U(a)] = \sum_{o \in \mathcal{O}} U(o) \cdot P(o|a)$$

$$\begin{aligned} E[U(\text{connect})] &= U(\text{attack}) \cdot P(\text{attack}|\text{connect}) \\ &+ U(\text{no attack}) \cdot P(\text{no attack}|\text{connect}) \\ &= -100 \cdot 0.01 + 10 \cdot 0.99 \\ &= 8.9 \end{aligned}$$

$\Rightarrow$ risk-neutral IT security manager chooses to connect since $E[U(\text{connect})] > E[U(\text{disconnect})]$.

Notes

Let's make a deal

- Option 1: Take $10
- Option 2: Get $20 with a 50% chance, $0 otherwise
- Which would you choose?
- $E[U] = 0.5 * \$20 + 0.5 * \$0 = \$10$
- Prefer option 1: you're risk-averse
- Prefer option 2: you're risk-seeking
- Are you indifferent? If so-you're risk-neutral

Notes

Let's make a deal (round 2)

- Option 1: Take $10
- Option 2: Get $150 with a 10% chance, $0 otherwise
- Which would you choose?
- $E[U] = 0.1 * \$150 + 0.5 * \$0 = \$15$
- Prefer option 1: you're risk-averse
- Prefer option 2: you're risk-neutral or seeking

Notes

Let's make a deal (round 3)

- Option 1: Take \$10
- Option 2: Get \$50 with a 10% chance, \$0 otherwise
- Which would you choose?
- $E[U] = 0.1 * \$50 + 0.5 * \$0 = \$5$
- Prefer option 1: you're risk-averse or risk-neutral
- Prefer option 2: you've got a gambling problem

Notes

Risk attitudes depend on the behavior of the utility function

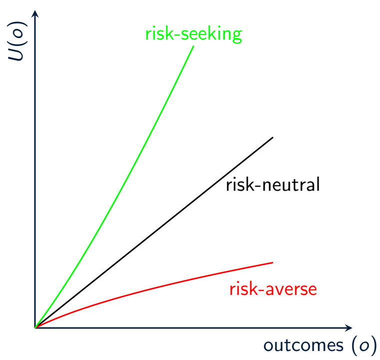

Notes

Risk-averse prefer utility of expected value over lottery

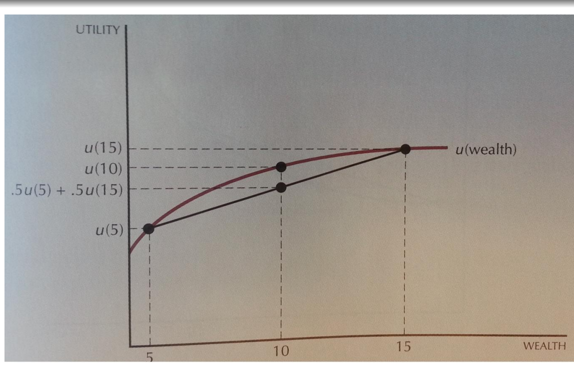

Source: Varian, *Intermediate Microeconomics*, p. 225

Notes

Risk-seekers prefer lottery over utility of expected value

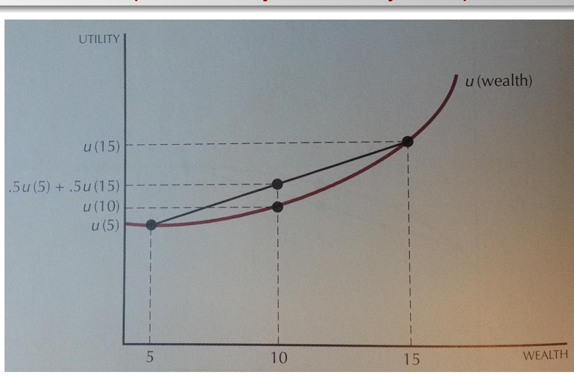

Source: Varian, *Intermediate Microeconomics*, p. 226

Notes

From attitudes to utility

- Suppose that outcomes are numeric $\mathcal{O} \in \mathbb{R}$
- When might that happen?
- Then we can define risk-attitudes by how the utility function behaves

Definition

(Risk neutrality) An agent is risk-neutral when $U(o)$ is a linear function on o.

Notes

From attitudes to utility

Definition

(Risk aversion) An agent is risk-averse when $U(o)$ is a concave function (i.e., $U''(x) < 0$ for a twice-differentiable function).

Definition

(Risk seeking) An agent is risk-seeking when $U(o)$ is a convex function (i.e., $U''(x) > 0$ for a twice-differentiable function).

Notes

Example: antivirus software

- Suppose you have \$10,000 in wealth. You have the option to buy antivirus software for \$75.
- Outcomes available:

$$\mathcal{O} = \{\text{hacked (decreases wealth by \$2,000)}, \text{not hacked (no change in wealth)}\}$$

- Without AV software, probability of being hacked is 0.05 ($P(\text{hacked}|\text{no antivirus}) = 0.05$)
- With AV software, probability of being hacked is 0 ($P(\text{hacked}|\text{antivirus}) = 0$)
- Exercise: compute the expected utility of both buying and not buying AV if you are risk-neutral (so that $U(o) = o$). Would you buy AV software?

Notes

Example: antivirus software

What if you are risk-averse (so that $U(o) = \sqrt{(o)}$)?

Risk-averse	hack		no hack		
Action	U	$P(\text{hack}|\text{action})$	U	$P(\text{no hack}|\text{action})$	$E[U(\text{action})]$
buy AV	$\sqrt{9,925}$	0	$\sqrt{9,925}$	1	99.6
don't buy	$\sqrt{8,000}$	0.05	$\sqrt{10,000}$	0.95	99.4

Exercise (on your own): How much would you pay for antivirus software if you were risk-neutral and the probability of getting hacked is 0.1 if you don't have AV installed?

Notes