Informatics 2D: Reasoning and Agents

Alex Lascarides

School of informatics

THE UNIVERSITY OF EDINBURGH

Lecture 20b: Rational Decisions

Where are we?

Last time...

- Use probabilities to represent uncertainty in our beliefs about
  - current state, future outcomes of our actions
  - ignorance about cause and effect
- Useful for managing risk
- Now: **uncertainty and rational decisions**

Uncertainty and rational decisions

- Logical agent has a goal and executes any plan guaranteed to achieve it
- Different with degrees of belief: If plan P has a 90% chance of success, how about another P' with a higher probability? Or how about P'' with higher cost but same probability?
- Agent must have **preferences** over **outcomes** of plans
- **Utility theory** can be used to reason about those preferences
- Based on idea that every state has a degree of usefulness and agents prefer states with higher utility
- Utilities vary from one agent to another.

Decision theory

- A general theory of rational decision making
- **Decision theory** = probability theory + utility theory
- Foundation of decision theory:

  An agent is rational if and only if it chooses the action that yields the highest expected utility, averaged over all possible outcomes of the action

- Principle of **Maximum Expected Utility**
- Although we follow it here, some points of criticism:
  - Knowledge of preferences?
  - Consistency of preferences?
  - Risk-taking attitude?

Are We Rational?

A: 100% chance of £3000 C: 25% chance of £3000
B: 80% chance of £4000 D: 20% chance of £4000

- 84% of you chose lottery A over lottery B.
- 68% of you chose lottery D over lottery C.
- So lots of you chose A and D, which is irrational!
  - If $U(3000) > 0.8 * U(4000)$, then $0.25 * U(3000) > 0.2 * U(4000)$!!
- Our ability to MEU also affected by emotion, social relationships, relationships among our choices...
- In fact, we're predictably irrational.
- If we were always rational, we wouldn't have self-help, life coaches etc.

Utility of money (empirical study)

- For most people concave curve (a), showing that going into debt is considered disastrous relative to small gains in money—**risk averse**.

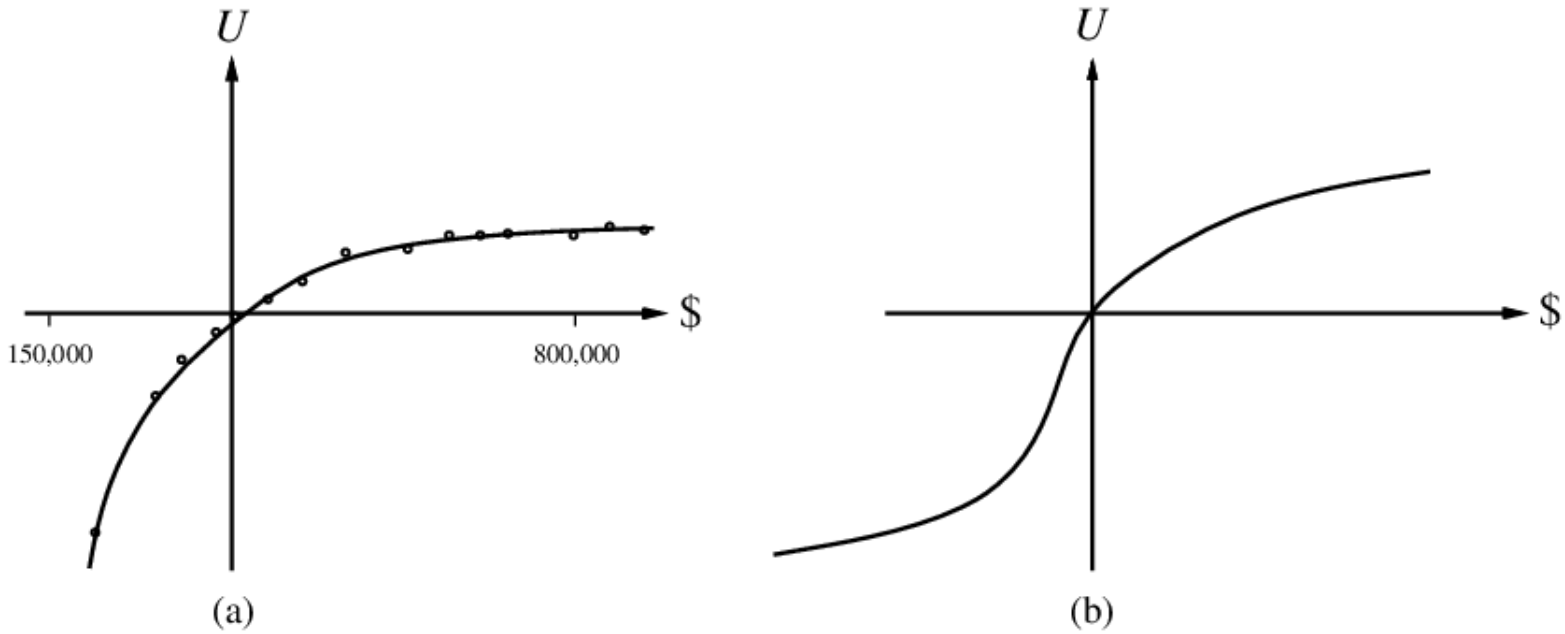

- But if you're already $10M in debt, your utility curve is more like (b)—**risk seeking** when desperate!

Design for a decision-theoretic agent

- For the time being, we will focus on probability and not utility.
- But still useful to have an idea of general abstract design for a decision-theoretic (utility-based) agent
- Characterised by basic perception-action loop as follows:
  1. Update belief state based on previous action and percept
  2. Calculate outcome probabilities for actions given action descriptions and belief states
  3. Select action with highest expected utility given probabilities of outcomes and utility information
- Very simple but broadly accepted as a general principle for building agents able to cope with real-world environments

Summary

- Probabilities represent degrees of belief
- Together with **utility theory**, we can model **rational decisions**:
  - Choose an action that maximises expected utility

  an optimal trade off between what you prefer and what you believe you can achieve
- Humans sometimes behave in a predictably irrational way, so rational agents will sometimes deviate from typical human behaviour.