CPS 570: Artificial Intelligence

Decision theory

Instructor: Vincent Conitzer

Risk attitudes

- Which would you prefer?
  - A lottery ticket that pays out $10 with probability .5 and $0 otherwise, or
  - A lottery ticket that pays out $3 with probability 1
- How about:
  - A lottery ticket that pays out $100,000,000 with probability .5 and $0 otherwise, or
  - A lottery ticket that pays out $30,000,000 with probability 1
- Usually, people do not simply go by expected value
- An agent is risk-neutral if she only cares about the expected value of the lottery ticket
- An agent is risk-averse if she always prefers the expected value of the lottery ticket to the lottery ticket
  - Most people are like this
- An agent is risk-seeking if she always prefers the lottery ticket to the expected value of the lottery ticket

Decreasing marginal utility

- Typically, at some point, having an extra dollar does not make people much happier (decreasing marginal utility)

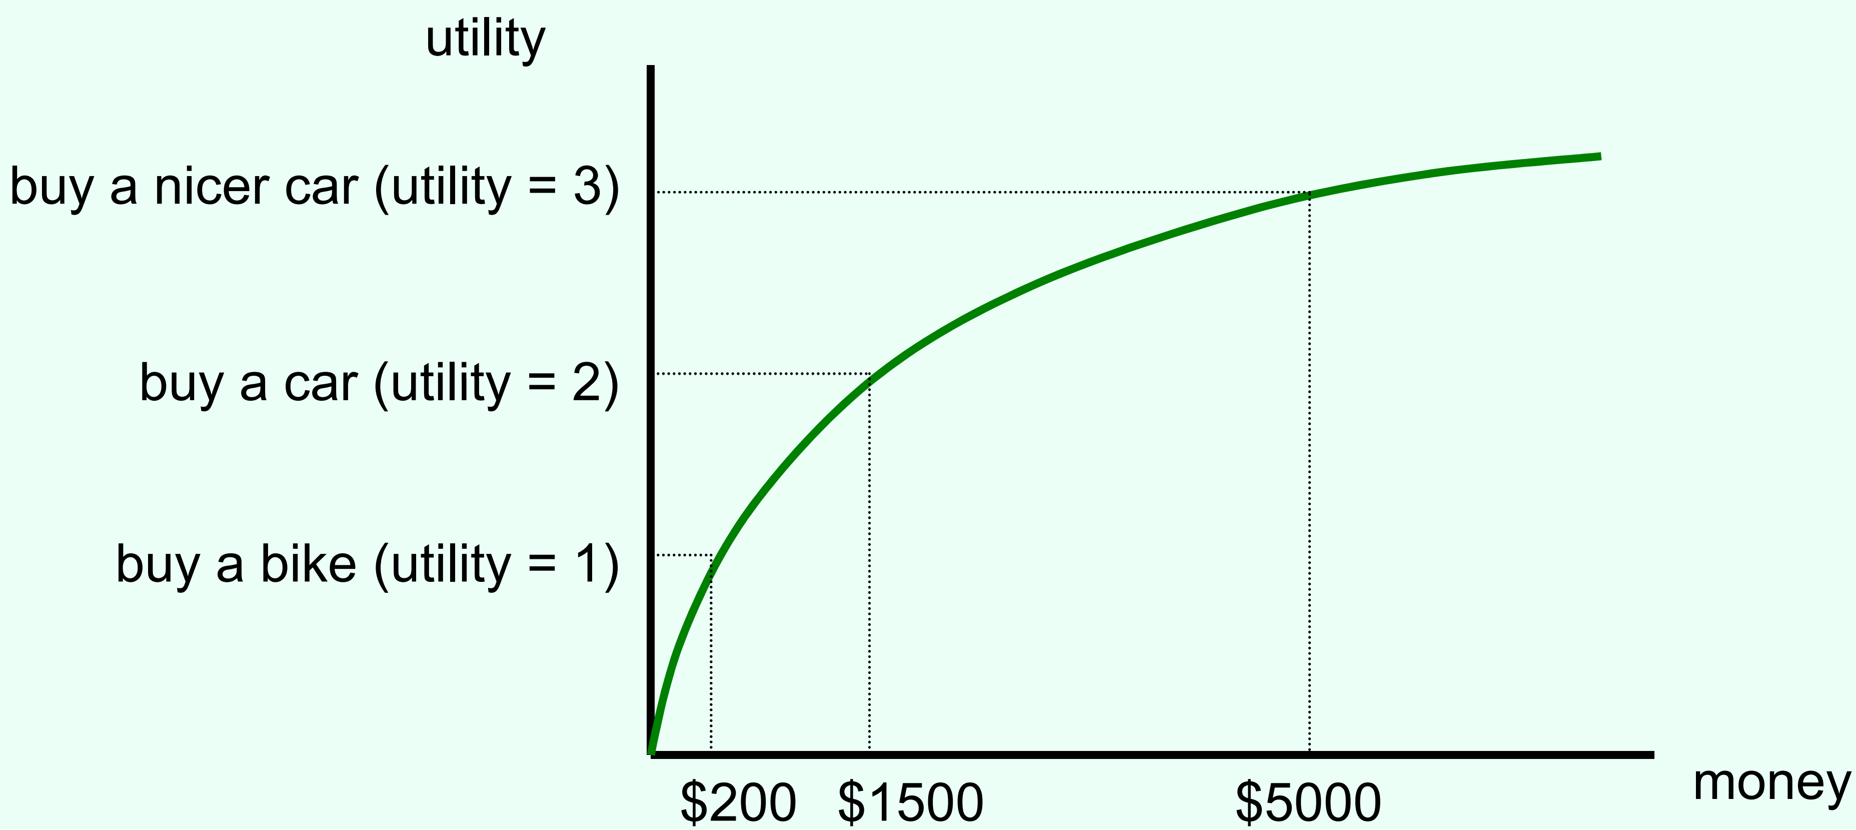

Maximizing expected utility

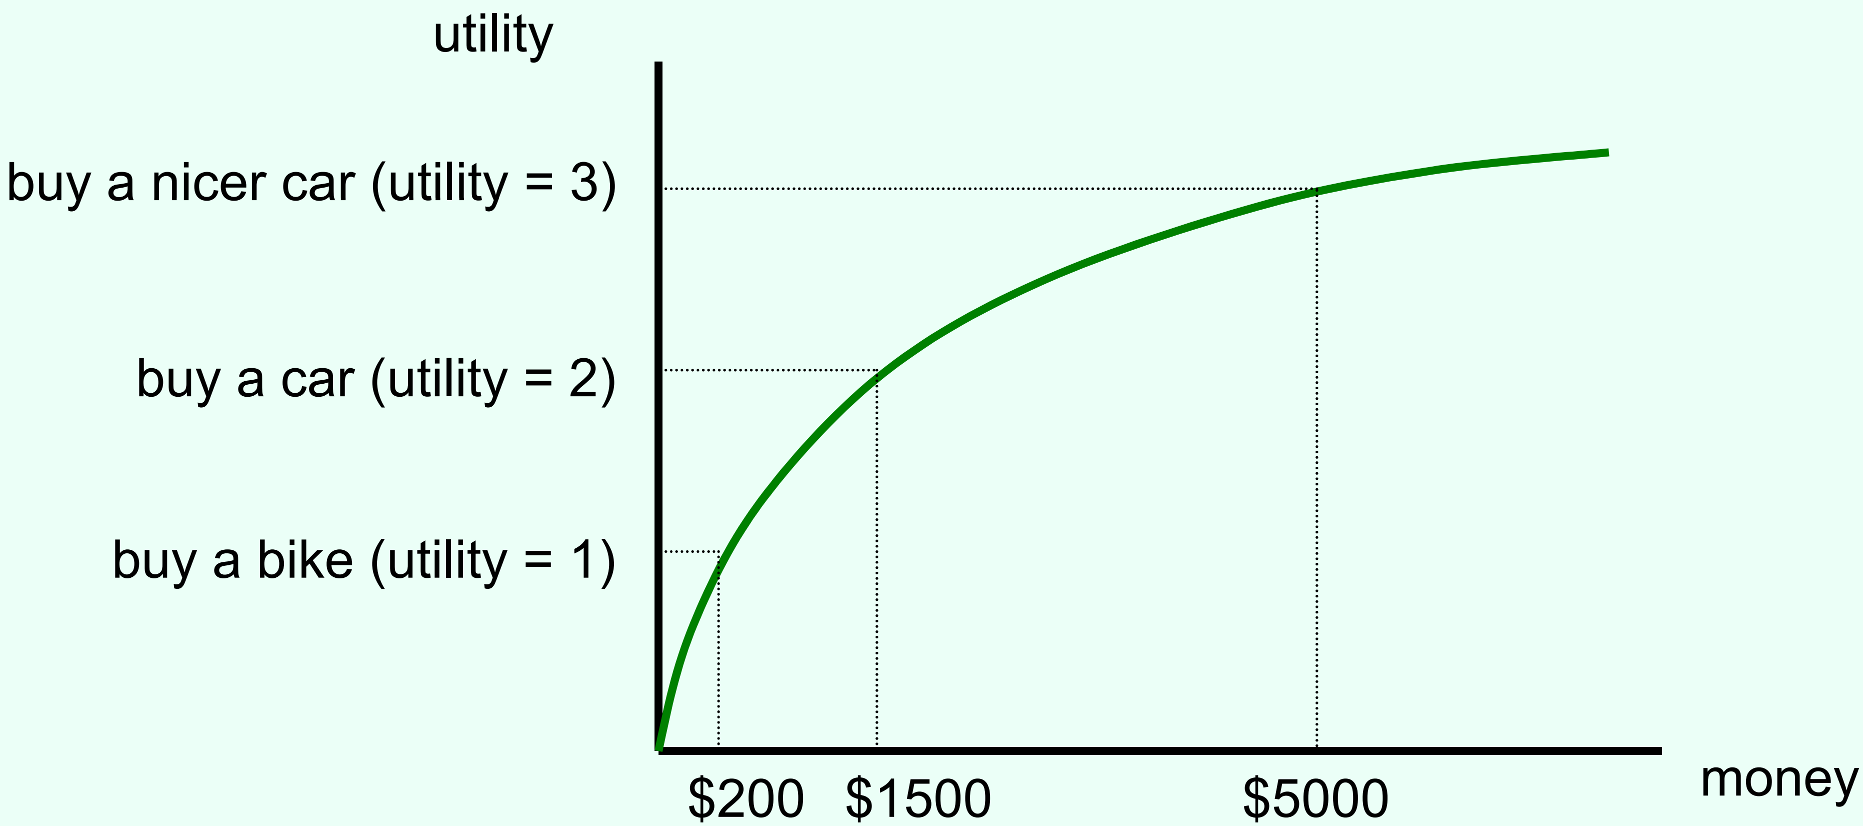

- Lottery 1: get $1500 with probability 1
  - gives expected utility 2
- Lottery 2: get $5000 with probability .4, $200 otherwise
  - gives expected utility .4*3 + .6*1 = 1.8
  - (expected amount of money = .4*$5000 + .6*$200 = $2120 > $1500)
- So: maximizing expected utility is consistent with risk aversion

Different possible risk attitudes under expected utility maximization

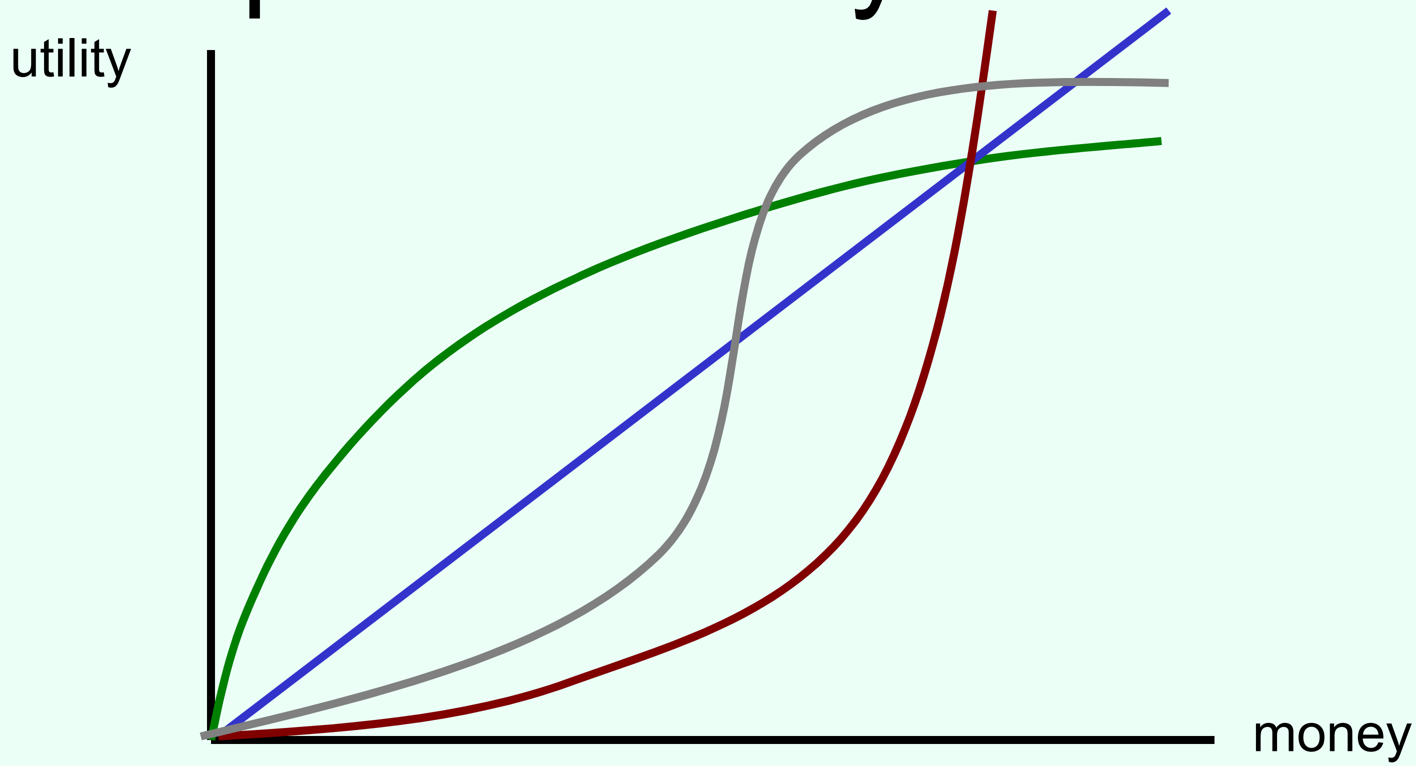

- Green has decreasing marginal utility → risk-averse
- Blue has constant marginal utility → risk-neutral
- Red has increasing marginal utility → risk-seeking
- Grey's marginal utility is sometimes increasing, sometimes decreasing → neither risk-averse (everywhere) nor risk-seeking (everywhere)

What is utility, anyway?

- Function u: O $\rightarrow \Re$ (O is the set of "outcomes" that lotteries randomize over)
- What are its units?
  - It doesn't really matter
  - If you replace your utility function by $u'(o) = a + bu(o)$, your behavior will be unchanged
- Why would you want to maximize expected utility?
  - This is a question about preferences over lotteries

Compound lotteries

- For two lottery tickets L and L', let pL + (1-p)L' be the "compound" lottery ticket where you get lottery ticket L with probability p, and L' with probability 1-p

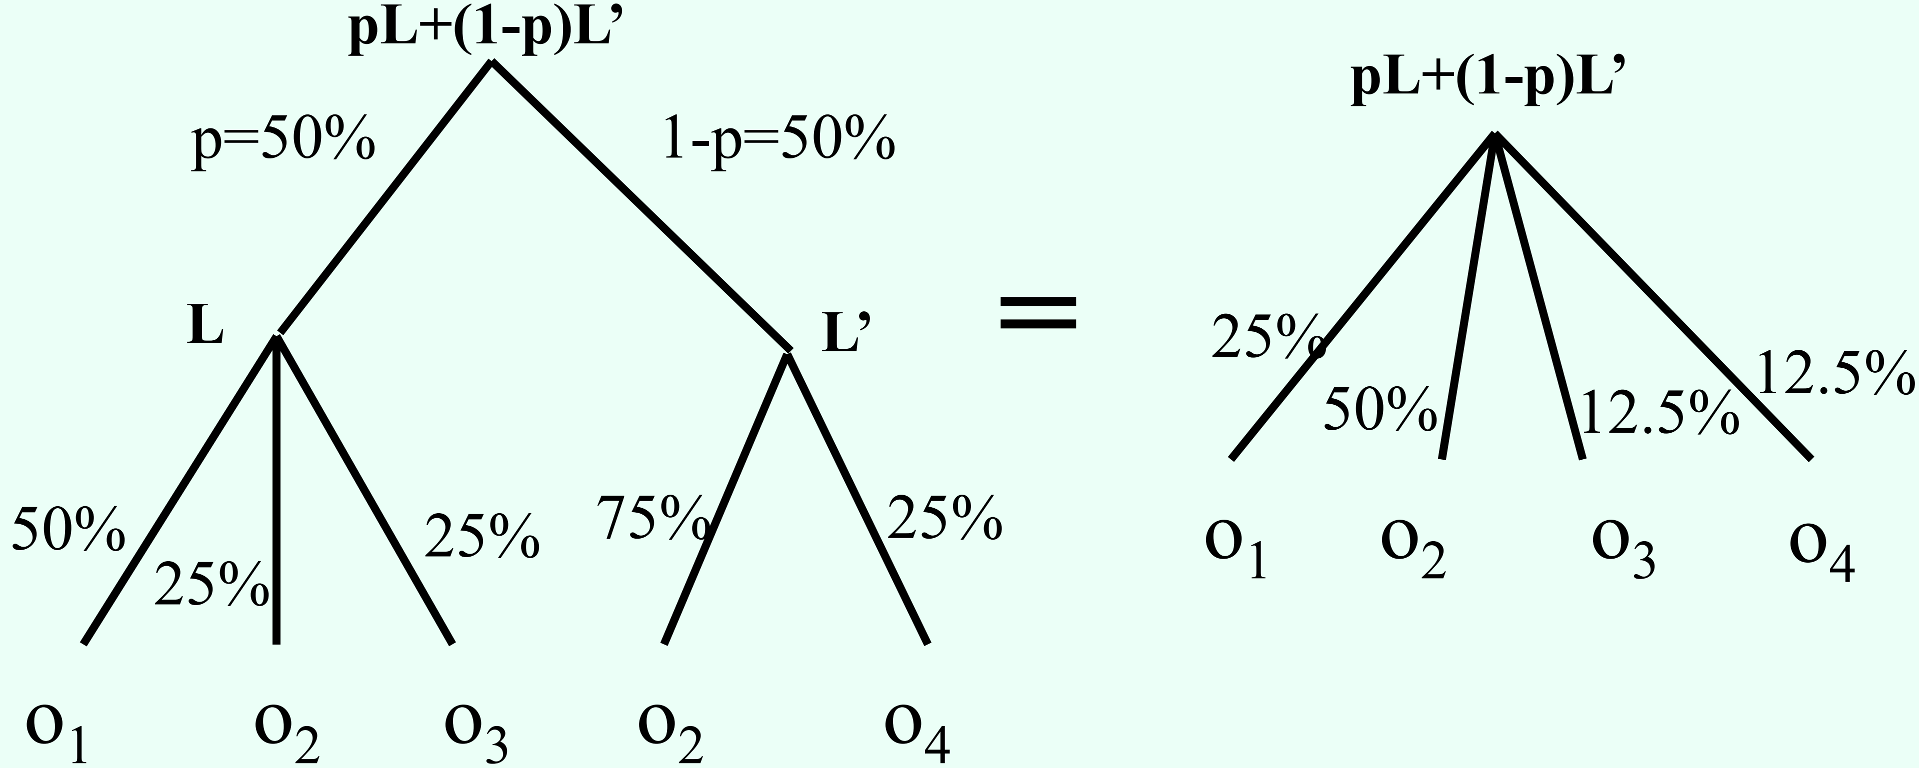

Sufficient conditions for expected utility

- L ≥ L' means that L is (weakly) preferred to L'
  - (≥ should be complete, transitive)
- **Expected utility theorem.** Suppose
  - (continuity axiom) for all L, L', L'', {p: pL + (1-p)L' ≥ L''} and {p: pL + (1-p)L' ≤ L''} are closed sets, and
  - (independence axiom – more controversial) for all L, L', L'', p > 0, we have L ≥ L' if and only if pL + (1-p)L'' ≥ pL' + (1-p)L''

Then, there exists a function u: O → $\Re$ so that L ≥ L' if and only if L gives a higher expected value of u than L'

Acting optimally over time

- Finite number of periods:
- Overall utility = sum of rewards in individual periods
- Infinite number of periods:
- … are we just going to add up the rewards over infinitely many periods?
  - Always get infinity!
- (Limit of) average payoff: $\lim_{n\to\infty}\Sigma_{1\leq t\leq n} r(t)/n$
  - Limit may not exist…
- Discounted payoff: $\Sigma_t \delta^t r(t)$ for some $\delta < 1$
- Interpretations of discounting:
  - Interest rate r: $\delta = 1/(1+r)$
  - World ends with some probability $1 - \delta$
- Discounting is mathematically convenient