Aalto University
School of Science

Decision making and problem solving – Lecture 3

- Modeling risk preferences
- Stochastic dominance

Liesiö, Punkka, Salo, Vilkkumaa

Motivation

❑ Last time:

- Decisions should be based on expected value of the alternatives' outcomes (if and) only if the DM is risk neutral
- Under 4 axioms for the DM's preference relation between risky alternatives, there exists a real-valued function ("utility function") so that
  - The DM should choose the alternative with the highest expected utility
  - It is unique up to positive affine transformations -> we can normalize the utility function the way we want

❑ This time:

- What is this utility function and how to model the DM's preferences with it?
- We learn how these preferences correspond to the DM's attitude towards risk

Assessment of utility functions

- Utility functions are assessed by asking the DM to choose between a simple lottery and a certain outcome (i.e., a degenerate lottery)
  - X: Certain payoff t
  - Y: Payoff t^+ (t^-) with probability p (1-p)

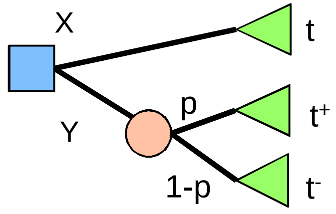

- General idea:
  - Vary the parameters (p,t,t^+, t^-) until the DM is indifferent between X and Y:

$$E[u(X)] = E[u(Y)] \Leftrightarrow u(t) = pu(t^+) + (1-p)u(t^-)$$

  - Repeat until sufficiently many points for the utility function have been obtained
- Because *u* is unique up to positive affine transformations, *u* can be fixed at two points
  - Usually, u is set at 1 at the most preferred level, and at 0 at the least preferred

Assessment: The certainty equivalence approach

- The DM assesses *t*
- Example: Assess utility function for the interval [-10,50] euros
  - Normalization: we can fix u(-10)=0 and u(50)=1

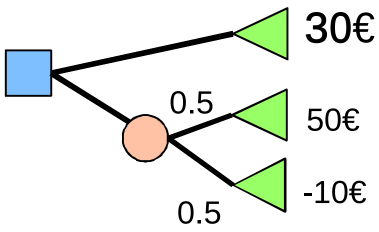

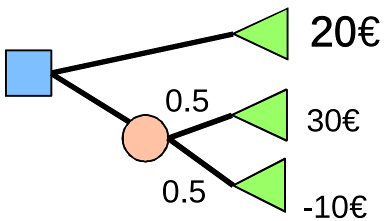

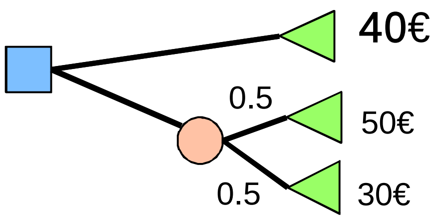

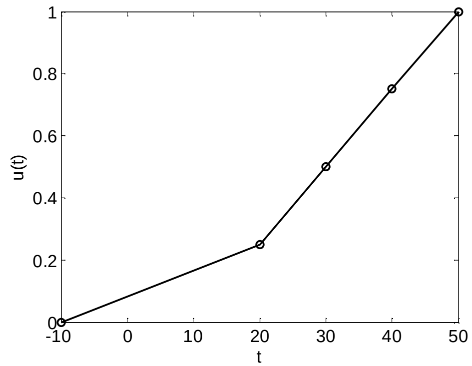

$u(30)$
$= 0.5u(-10) + 0.5u(50)$
$= 0.5 \cdot 0 + 0.5 \cdot 1 = 0.5$

$u(20)$
$= 0.5u(-10) + 0.5u(30)$
$= 0.5 \cdot 0 + 0.5 \cdot 0.5 = 0.25$

$u(40)$
$= 0.5u(30) + 0.5u(50)$
$= 0.5 \cdot 0.5 + 0.5 \cdot 1$
$= 0.75$

Other approaches to utility assessment

- Probability equivalence:
  - The DM assesses p
- Gain equivalence:
  - The DM assesses t^+
- Loss equivalence:
  - The DM assesses t^-

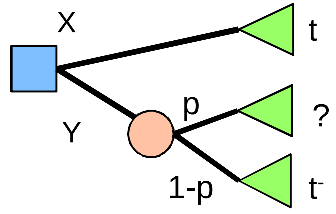

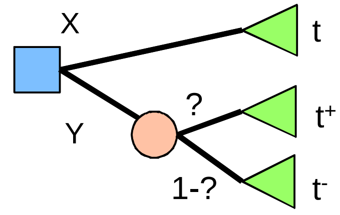

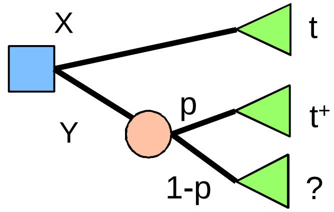

- Often in applications, the analyst chooses a family of utility functions and then asks the DM to compare lotteries to fix the parameter(s)
  - E.g., the exponential utility function (parameter ρ)

$$u(t) = 1 - e^{-\frac{t}{\rho}}, \rho > 0$$

Reference lottery revisited

- Assume that an expected utility maximizer with utility function *u* uses a reference lottery to assess the probability of event A
- She thus adjusts *p* such that she is indifferent between lottery X and reference lottery Y:

$$E[u(X)] = E[u(Y)]$$

$$\Leftrightarrow P(A)u(t^+) + \big(1 - P(A)\big)u(t^-) = pu(t^+) + (1-p)u(t^-)$$

$$\Leftrightarrow P(A)\big(u(t^+) - u(t^-)\big) = p\big(u(t^+) - u(t^-)\big)$$

$$\Leftrightarrow P(A) = p$$

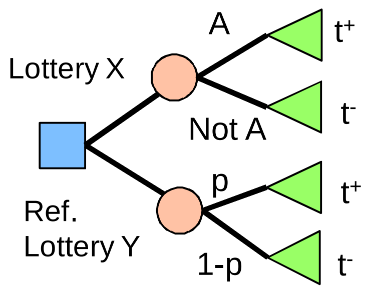

- Utility function *u* does not affect the result

Expected utility in decision trees

- Do everything in the usual way, but
  - Chance node: compute the expected <u>utility</u>
  - Decision node: select the alternative corresponding to maximum expected <u>utility</u>
  - Cf. the umbrella example, in which ‘some numbers’ represented preferences

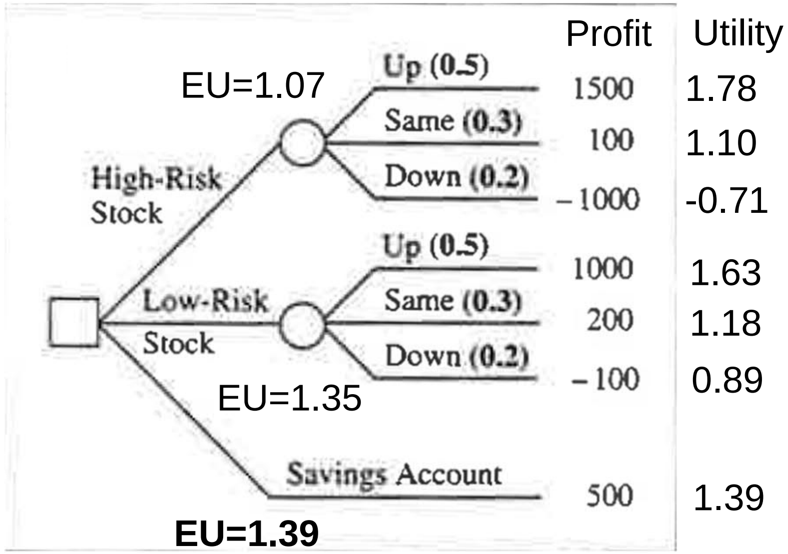

$$u(t) = 2 - e^{\frac{-t}{1000}}$$

Expected utility in Monte Carlo

- For each sample $x_1, \dots, x_n$ of random variable X, compute utility $u(x_i)$
- Mean of sample utilities $u(x_1), \dots, u(x_n)$ provides an estimate for $E[u(X)]$

=2-EXP(-F12/1000)

C	D	E	F	G	H
		Col.mean	Col.mean	Col.mean	
		0.502964	990.3014	1.580972	
	Sample	u	x	Utility	
	1	0.464077	954.9167	1.615156	
	2	0.704234	1268.308	1.718693	
	3	0.777865	1382.501	1.74905	
	4	0.534927	1043.831	1.647897	
	5	0.4426	927.8094	1.604581	
	6	0.916252	1690.147	1.815508	
	7	0.649453	1191.922	1.696363	
	8	0.65278	1196.418	1.697725	
	9	0.110887	389.0874	1.322325	
	10	0.189275	559.714	1.428628	
	11	0.902882	1649.073	1.807772	

EUT for normative decision support

- EUT is a normative theory: if the DM is rational, she **should** select the alternative with the highest expected utility
  - Not descriptive or predictive: EUT does not describe or predict how people actually do select among alternatives with uncertain outcomes

- The four axioms characterize properties that are required for rational decision support
  - Cf. probability axioms describe a rational model for uncertainty
  - The axioms are not assumptions about the DM's preferences

Question 1

http://presemo.aalto.fi/2134lec2

- Which of the below alternatives would you choose?

1. A sure gain of 1 M€
2. A gamble in which there is a
   - 1% probability of getting nothing,
   - 89% probability of getting 1M€, and
   - 10% probability of getting 5M€

Question 2

http://presemo.aalto.fi/2134lec2

- Imagine that a rare disease is breaking out in a community and is expected to kill 600 people. Two different programs are available to deal with the threat.
  - If Program A is adopted, 200 people will be saved
  - If Program B is adopted, there's a 33% probability that all 600 will be saved and a 67% probability that no one will be saved.

Which program will you choose?

1. Program A
2. Program B

Question 3

http://presemo.aalto.fi/2134lec2

- Which of the below alternatives would you choose?

1. A gamble in which there is a
   - 89% probability of getting nothing and
   - 11% probability of getting 1M€

2. A gamble in which there is a
   - 90% probability of getting nothing, and
   - 10% probability of getting 5M€

Question 4

http://presemo.aalto.fi/2134lec2

- Imagine that a rare disease is breaking out in some community and is expected to kill 600 people. Two different programs are available to deal with the threat.
  – If Program C is adopted, 400 of the 600 people will die,
  – If Program D is adopted, there is a 33% probability that nobody will die and a 67% probability that 600 people will die.

  Which program will you choose?

1. Program C
2. Program D

Allais paradox

- Which of the below alternatives would you choose?
  A. A sure gain of 1 M€
  B. A gamble in which there is a
     - 1% probability of getting nothing,
     - 89% probability of getting 1M€, and
     - 10% probability of getting 5M€

- Which of the below alternatives would you choose?
  C. A gamble in which there is a
     - 89% probability of getting nothing and
     - 11% probability of getting 1M€
  D. A gamble in which there is a
     - 90% probability of getting nothing, and
     - 10% probability of getting 5M€

Most people choose A; hence E[u(A)]>E[u(B)]:
$u(1) > 0.10u(5)+0.89u(1)+0.01u(0) \Rightarrow$

$$\boxed{0.11u(1) > 0.10u(5)+0.01u(0)}$$

Most people choose D; hence E[u(D)]>E[u(C)]:
$0.10u(5)+0.90u(0) > 0.11u(1)+0.89u(0) \Rightarrow$

$$\boxed{0.11u(1) < 0.10u(5)+0.01u(0)}$$

- **Actual choice behavior is not always consistent with EUT**

Framing effect

- Most people choose A and D
- People tend to be ”risk-averse” about gains and ”risk-seeking” about losses

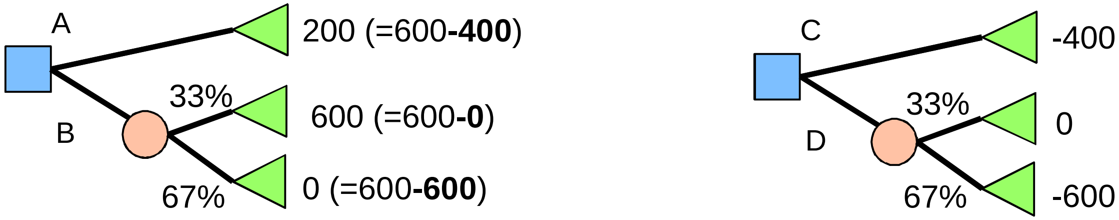

Risk and risk preferences

- Risk: possibility of loss (or some other unpreferred outcome)
  - Characterized by both the probability and magnitude of loss
- Risk preferences:
  - How does the riskiness of a decision alternative affect its desirability?
  - E.g., risk neutrality: choose the alternative with the highest expected (monetary) value, riskiness is not a factor
- Definition of risk preferences requires that outcomes T are quantitative and preferences among them *monotonic*
  - E.g., profits, costs, lives saved etc.
- Here, we assume that more is preferred to less, i.e., $u(t)$ is increasing (and differentiable) for all t

Certainty equivalent in Expected Utility Theory

- **Definition:** Certainty equivalent of a random variable X, denoted by CE[X], is an outcome in T such that

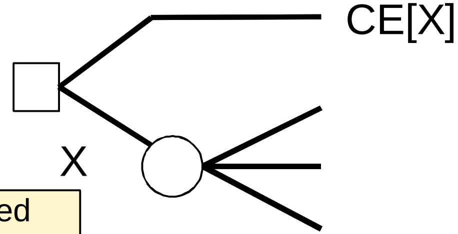

$$u(CE[X]) = E[u(X)] \Leftrightarrow$$
$$CE[X] = u^{-1}(E[u(X)])$$

Allowed because u is monotonic

– **IMPORTANT!** CE[X] is the certain outcome such that the DM is indifferent between alternatives X and CE[X]
  – CE[X] depends on both the DM's utility function u (preferences) and the distribution of X (uncertainty)
    - o My CE for roulette may be different from yours
    - o My CE for roulette may be different from my CE for one-armed bandit

Certainty equivalent - Example

- Consider a decision alternative X with $f_X(3) = 0.5$ and $f_X(5) = 0.5$ and three DMs with the below utility functions
- Compute each DM's certainty equivalent for X

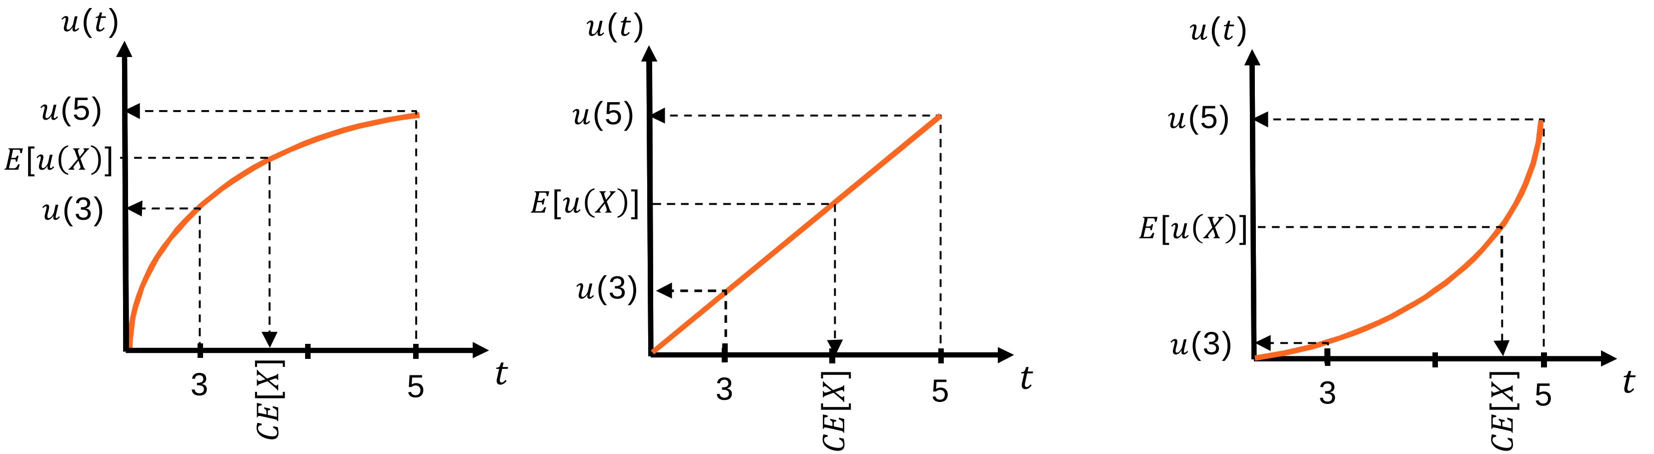

- The shape of the utility function seems to determine whether CE[X] is below, above, or equal to E[X]=4

Convex and concave functions

- **Definition:** *u* is concave, if for any $t_{1,}t_2$:

  $$\lambda u(t_1) + (1-\lambda)u(t_2) \leq u(\lambda t_1 + (1-\lambda)t_2) \quad \forall\lambda \in [0,1]$$

  – A line drawn between any two points $u(t_1)$ and $u(t_2)$ is below (or equal to) $u(t)$

  – $u''(t) \leq 0 \; \forall t \in T$, if the second derivative exists

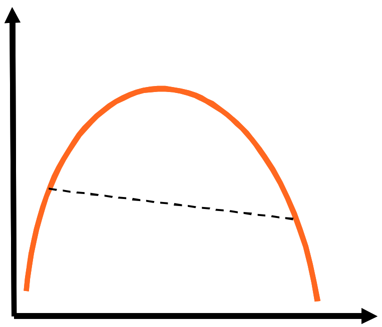

- **Definition:** *u* is convex, if for any $t_{1,}t_2$:

  $$\lambda u(t_1) + (1-\lambda)u(t_2) \geq u(\lambda t_1 + (1-\lambda)t_2) \quad \forall\lambda \in [0,1]$$

  – A line drawn between any two points $u(t_1)$ and $u(t_2)$ is above (or equal to) $u(t)$

  – $u''(t) \geq 0 \; \forall t \in T$, if the second derivative exists

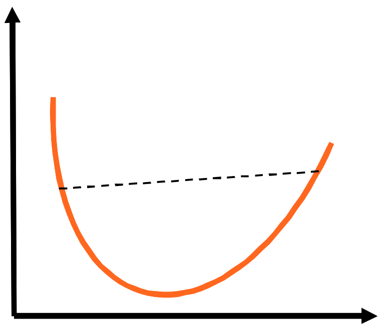

Convex utility functions

- For any utility function *u*, $E[u(X)] = \sum f_X(t_i)\, u(t_i)$ for X with discrete set of outcomes $t_i, i = 1, \dots, n$
  - Note: $\sum f_X(t_i) = 1$

- Let *u* be convex. Then
  - $\lambda u(t_1) + (1-\lambda)u(t_2) \geq u(\lambda t_1 + (1-\lambda)t_2)\ \ \forall \lambda \in [0,1]$ (by def., previous slide)
  - And, specifically, by applying this definition several times,

$$f_X(t_1)u(t_1) + \ \dots\ + f_X(t_n)u(t_n) = E[U(X)] \geq u\left(\sum f_X(t_i)t_i\right) = U(E[X])$$

- For convex *u*: Expected utility of X is higher than (expected) utility of E(X)

Jensen's inequality

- For any random variable *X*, if function *u* is
  I. Convex, then $E[u(X)] \geq u(E[X])$
  II. Concave, then $E[u(X)] \leq u(E[X])$

$\Rightarrow$

u concave

$$\Rightarrow E[u(X)] \leq u(E[X])$$
$$\Leftrightarrow u^{-1}(E[u(X)]) \leq u^{-1}(u(E[X]))$$
$$\Leftrightarrow CE[X] \leq E[X]$$

Allowed because u is increasing

u convex

$$\Rightarrow E[u(X)] \geq u(E[X])$$
$$\Leftrightarrow u^{-1}(E[u(X)]) \geq u^{-1}(u(E[X]))$$
$$\Leftrightarrow CE[X] \geq E[X]$$

Risk attitudes in Expected Utility Theory

I. *u* is concave iff CE[X] ≤ E[X] for all X
II. *u* is convex iff CE[X] ≥ E[X] for all X
III. *u* is linear iff CE[X]=E[X] for all *X*

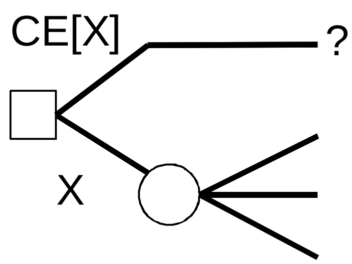

- A DM with a linear utility function is called *risk neutral*
  - Indifferent between uncertain outcome X and a certain outcome equal to E[X]
- A DM with a concave but not linear utility function is called *risk averse*
  - Prefers a certain outcome smaller than E[X] to uncertain outcome X
- A DM with a convex but not linear utility function is called *risk seeking*
  - Requires a certain outcome larger than E[X] to not choose uncertain outcome X

Risk premium in Expected Utility Theory

- **Definition:** Risk premium for random variable X is RP[*X*]=E[*X*]-CE[*X*]
  - RP[X] depends on both the DM's preferences (u) and the uncertainty in the decision alternative (distribution of X)
  - RP[X] is the premium that the DM requires on the expected value to change a certain outcome of CE[X] to an uncertain outcome X

I. DM is risk neutral, iff RP[X]=0 for all X

II. DM is risk averse, iff RP[X] $\geq$ 0 for all X

III. DM is risk seeking, iff RP[X] $\leq$ 0 for all X

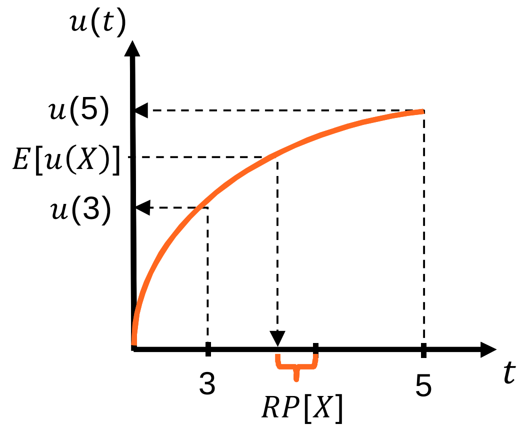

Computing CE and RP

1. Compute E[u(X)] and E(X)
2. Solve $u^{-1}(\cdot)$
3. Compute $CE[X] = u^{-1}(E[u(X)])$
4. Compute RP[*X*]=E[*X*]-CE[*X*]

- Step 2: if $u^{-1}(\cdot)$ cannot be solved analytically, solve it numerically from $u(CE[X]) = E[u(X)]$
  - Trial and error
  - Computer software

Example: Jane's $u(t) = t^2$ and her payoff is *Y~Uni*(3,5)

1. $E[u(X)] = \int_3^5 f_Y(t)u(t)dt = 16.33$
2. $v = u(t) = t^2 \Leftrightarrow t = u^{-1}(v) = \sqrt{v}$
3. $CE[X] = u^{-1}(16.33) = \sqrt{16.33} = 4.04$
4. RP[*X*] = 4 - 4.04 = -0.04

Prospect theory

- Expected Utility Theory assumes that people only care about the outcome in the *absolute* sense
- Yet, empirical evidence suggests that people tend to
  - think of possible outcomes relative to a certain reference point (often the status quo),
  - have different risk attitudes towards gains and losses with regard to the reference point,
  - overweight extreme, but unlikely events, but underweight "average" events.
- Prospect theory seeks to accommodate these empirical findings:

  Tversky, A. and D. Kahneman. "Advances in prospect theory: Cumulative representation of uncertainty." *Journal of Risk and uncertainty* 5.4 (1992): 297-323.

- **NOTE:**
  - EUT is a normative theory: tells what rational people should do
  - Prospect theory is a descriptive theory: tries to describe what people tend to do in real life

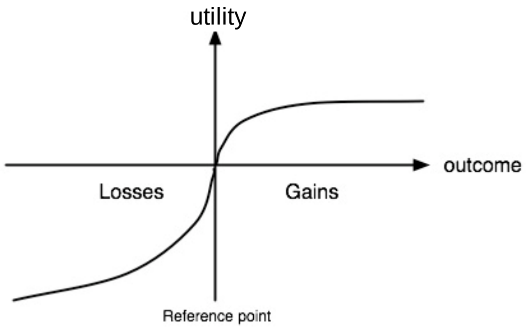

Stochastic dominance

https://presemo.aalto.fi/stocdom/

- **Question:** Which decision alternative would you choose?

1. X
2. Y

$$F_X(t) \leq F_Y(t) \ \ \forall t \in T$$

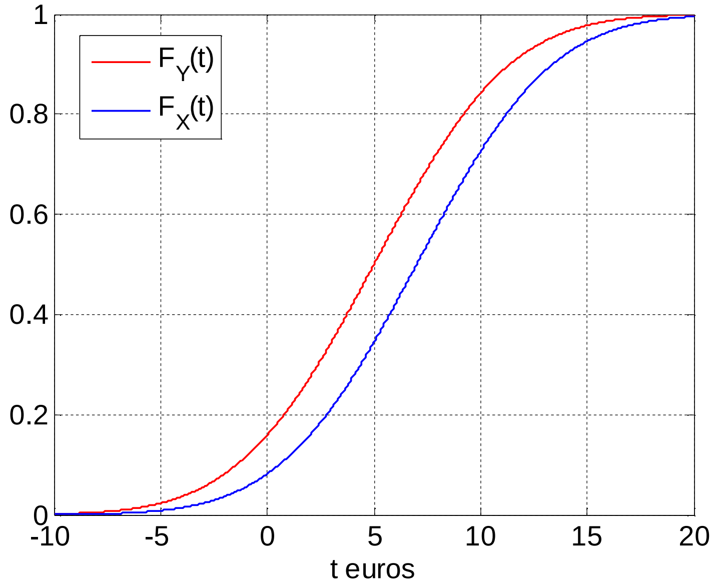

First-degree Stochastic Dominance

Definition: *X* dominates *Y* in the sense of First-degree Stochastic Dominance (denoted $X \succcurlyeq_{FSD} Y$), if

$$F_X(t) \leq F_Y(t) \;\; \forall t \in T$$

with strict inequality for some t.

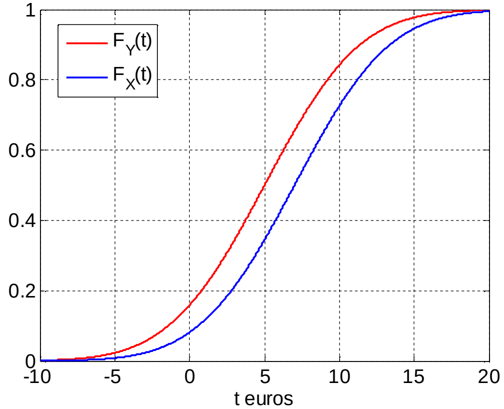

$X \succcurlyeq_{FSD} Y$

Theorem: $X \succcurlyeq_{FSD} Y$ if and only if

$$E[u(X)] \geq E[u(Y)] \;\; \forall u \in U^0,$$

where U^0 is the set of all strictly increasing functions

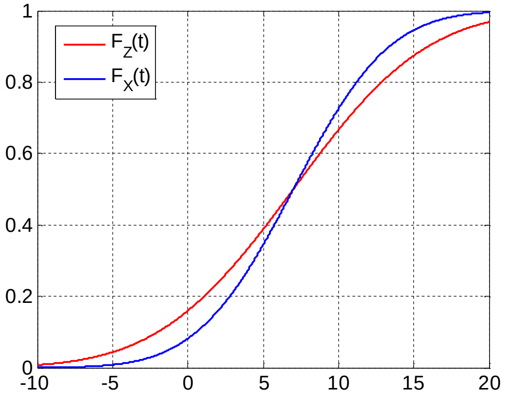

Neither alternative dominates the other in the sense of FSD

<u>Implication:</u> If an alternative is strictly dominated in the sense of FSD, then any DM who prefers more to less should not choose it.

FSD: Mining example

- A mining company has an opportunity to bid on two separate parcels of land
- Decisions to be made:
- Overall commitment of some $500 million
  - How much to bid?
  - Bid alone or with partner?
  - How to develop the site if the bid turns out successful?
- Large decision tree model built to obtain cumulative distribution functions of different strategies (= decision alternatives)

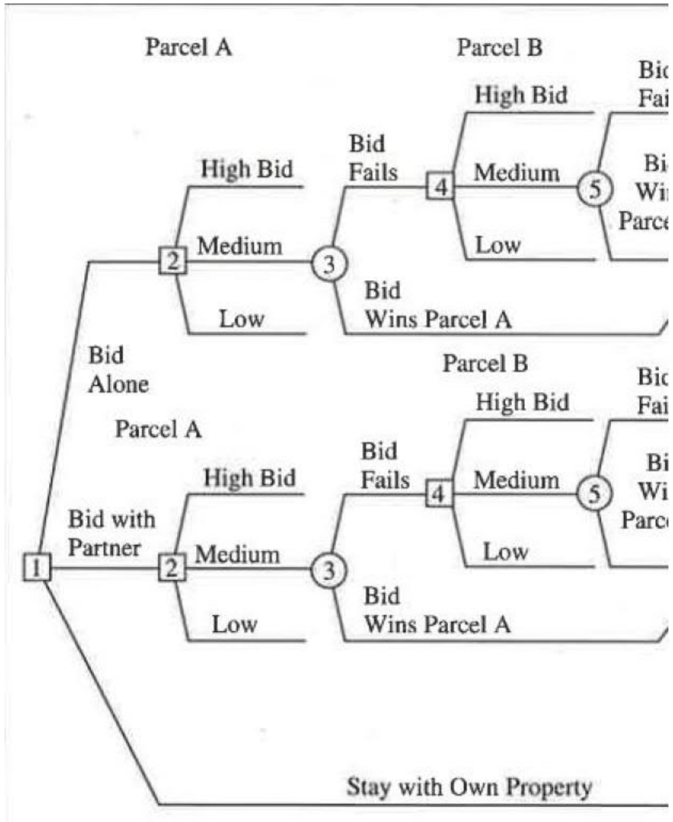

Source: Hax and Wing (1977): "The use of decision analysis in a capital investment probelm" In Bell, Keeney, and Raiffa (eds.): *Conflicting Objectives in Decisions*, Wiley.

FSD: Example (cont'd)

- Assume that the company prefers a larger net present value (NPV) to a smaller one

- Which strategies would you recommend?

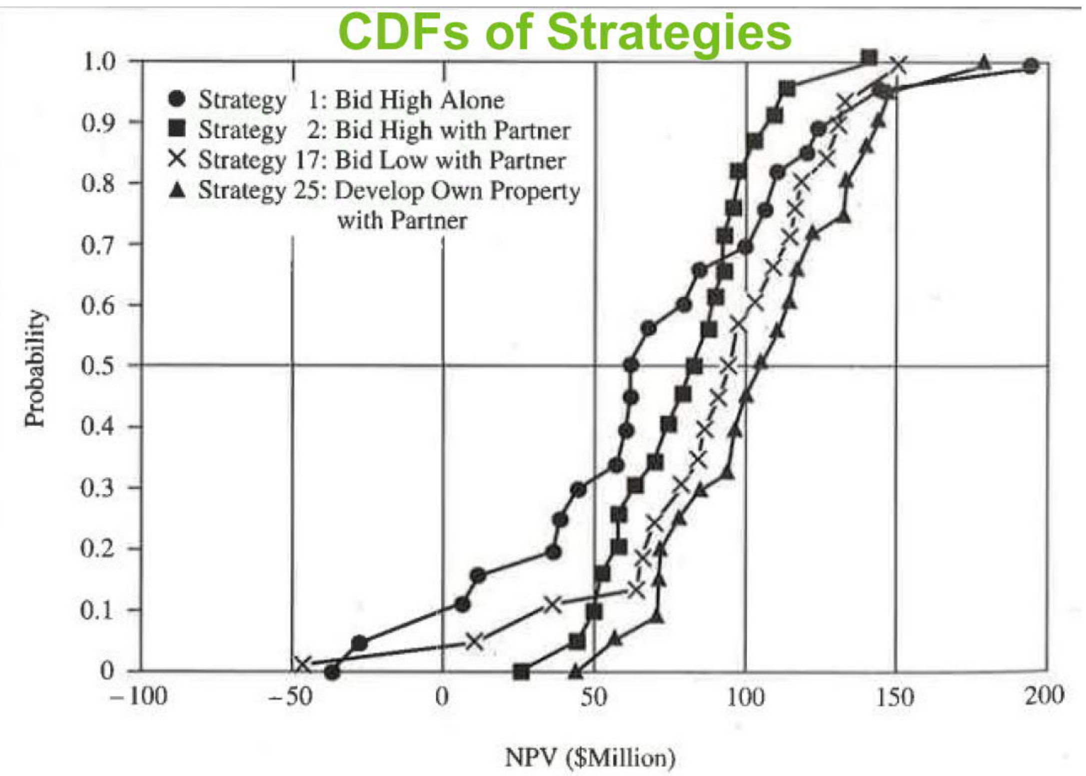

Source: Hax and Wing (1977): "The use of decision analysis in a capital investment probelm" In Bell, Keeney, and Raiffa (eds.): *Conflicting Objectives in Decisions*, Wiley.

Second-degree Stochastic Dominance

- **Theorem:**

$$E[u(X)] \geq E[u(Y)] \;\; \forall u \in U^{\,ccv} \Leftrightarrow \int_{-\infty}^{z} [F_X(t) - F_Y(t)]dt \leq 0 \;\; \forall z \in T,$$

  where $U^{\,ccv} = \{u \in U^0 | u \text{ is concave}\}$.

- **Definition:** *X* dominates *Y* in the sense of Second-degree Stochastic Dominance (denoted $X \succcurlyeq_{\text{SSD}} Y$), if

$$\int_{-\infty}^{z} [F_X(t) - F_Y(t)]dt \leq 0 \;\; \forall z \in T.$$

  with strict inequality for some *z.*

- <u>Implication:</u> If an alternative is strictly dominated in the sense of SSD, then any risk-averse or risk neutral DM who prefers more to less should not choose it.

SSD: graphical interpretation

$$\int_{-\infty}^{z} [F_X(t) - F_Y(t)]dt \leq 0 \;\; \forall z \in T$$

- Integral
  = the area between $F_X(t)$ and $F_Y(t)$ up to point *z*
  = the area between the $F_X(t)$-$F_Y(t)$ and the horizontal axis up to point *z*
- If it is non-positive for all *z*, then $X \succcurlyeq_{SSD} Y$
- Here: $X \succcurlyeq_{SSD} Y$, because area **A** is bigger than area **B**, and **A** is left of **B**

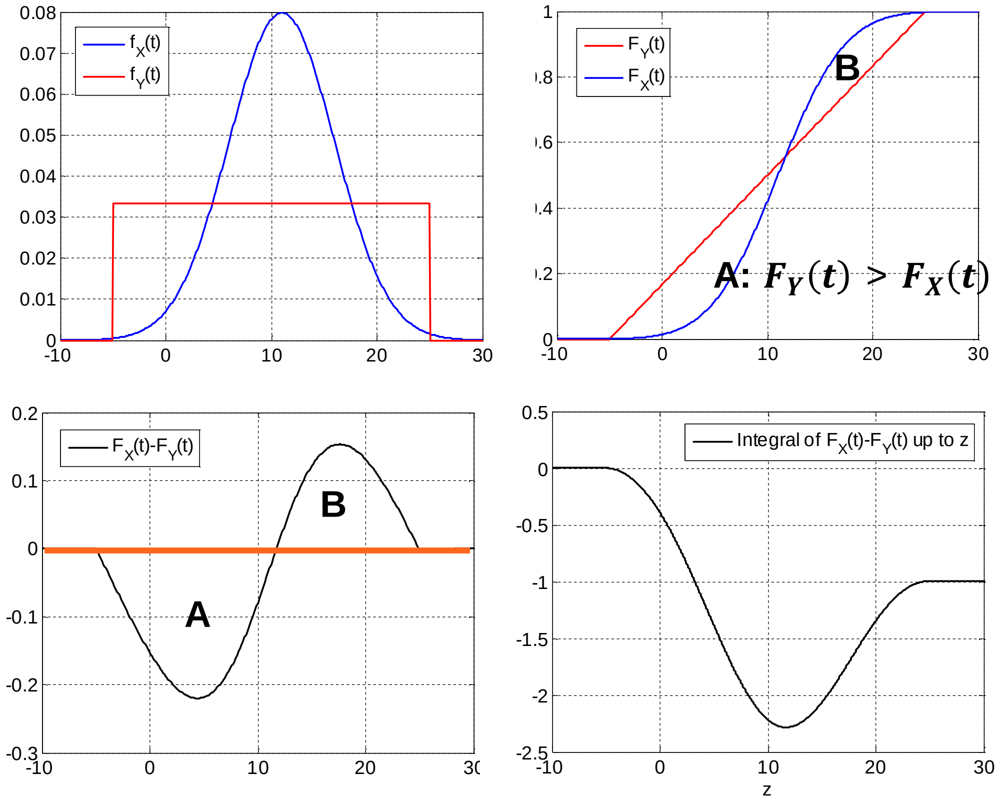

SSD: Mining example revisited

- Assume that the mining company is either risk-averse or risk-neutral

- Which strategies would you recommend?

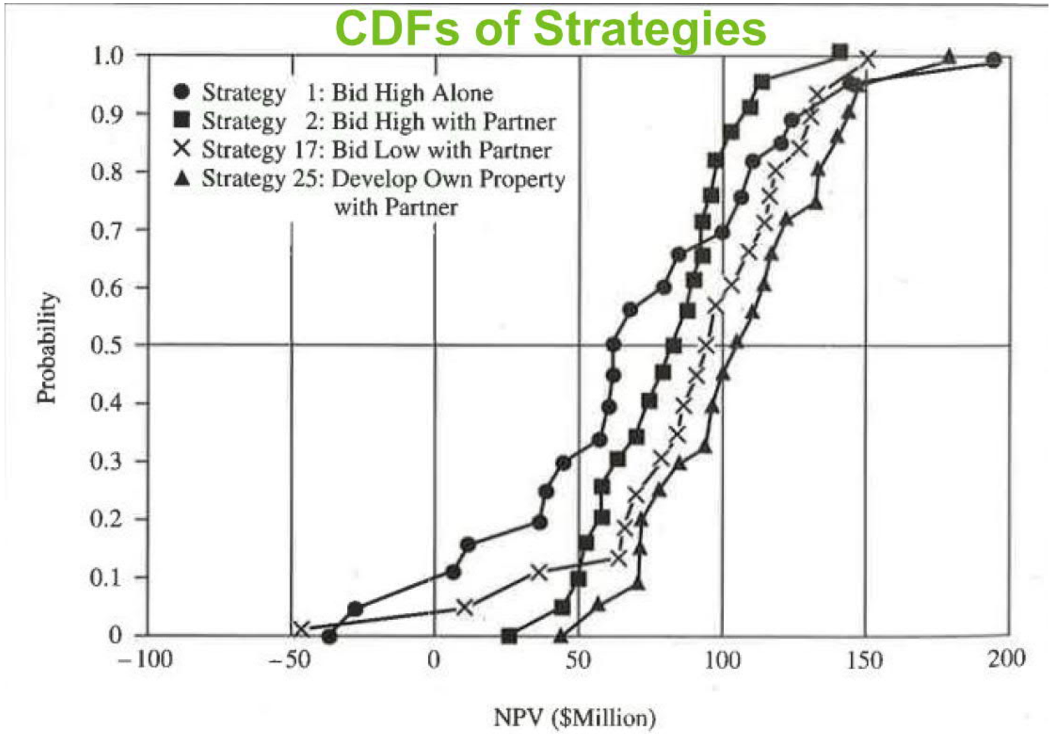

Properties of FSD and SSD

- Both FSD and SSD are transitive:
  - If X $\succcurlyeq_{FSD}$ Y and Y $\succcurlyeq_{FSD}$ Z, then X $\succcurlyeq_{FSD}$ Z
    - Why? Take any t. Then, $F_X(t) \leq F_Y(t) \leq F_Z(t)$.
  - If X $\succcurlyeq_{SSD}$ Y and Y $\succcurlyeq_{SSD}$ Z, then X $\succcurlyeq_{SSD}$ Z
    - Why? Take any $u \in U^{ccv}$. Then, $E[u(X)] - E[u(Z)] \geq E[u(Y)] - E[u(Z)] \geq 0$.
- FSD implies SSD:
  - If X $\succcurlyeq_{FSD}$ Y, then X $\succcurlyeq_{SSD}$ Y.
    - Why? Take any $u \in U^{ccv}$. Then, $u \in U^0$, and since X $\succcurlyeq_{FSD}$ Y, we have $E[u(X)] \geq E[u(Y)]$.
    - Or consider the definitions of FSD and SSD: If $F_X(t) \leq F_Y(t) \;\; \forall t \in T$, then

$$\int_{-\infty}^{z} [F_X(t) - F_Y(t)]dt \leq \int_{-\infty}^{z} 0dt \leq 0 \;\; \forall z \in T$$

Summary

- **Utility function** is elicited **through** specification of equally **preferred lotteries**
  - Then: expected utilities equal

- The **shape** of the utility function determines the DM's **risk attitude**
  - Linear utility function = risk neutral
  - Concave utility function = risk averse
  - Convex utility function = risk seeking

- Even if the utility function is not completely specified, **decision recommendations** may be implied by **stochastic dominance**
  - If the DM prefers more to less, she should not choose an FSD dominated alternative
  - If the DM is also risk averse, she should not choose an SSD dominated alternative