Quasilinear Mechanism Design

CMPUT 654: Modelling Human Strategic Behaviour

S&LB §10.3-10.4

Recap: Revelation Principle

Theorem: (Revelation Principle)
If there exists **any** mechanism that implements a social choice function C in dominant strategies, then there exists a **direct** mechanism that implements C in dominant strategies and is **truthful**.

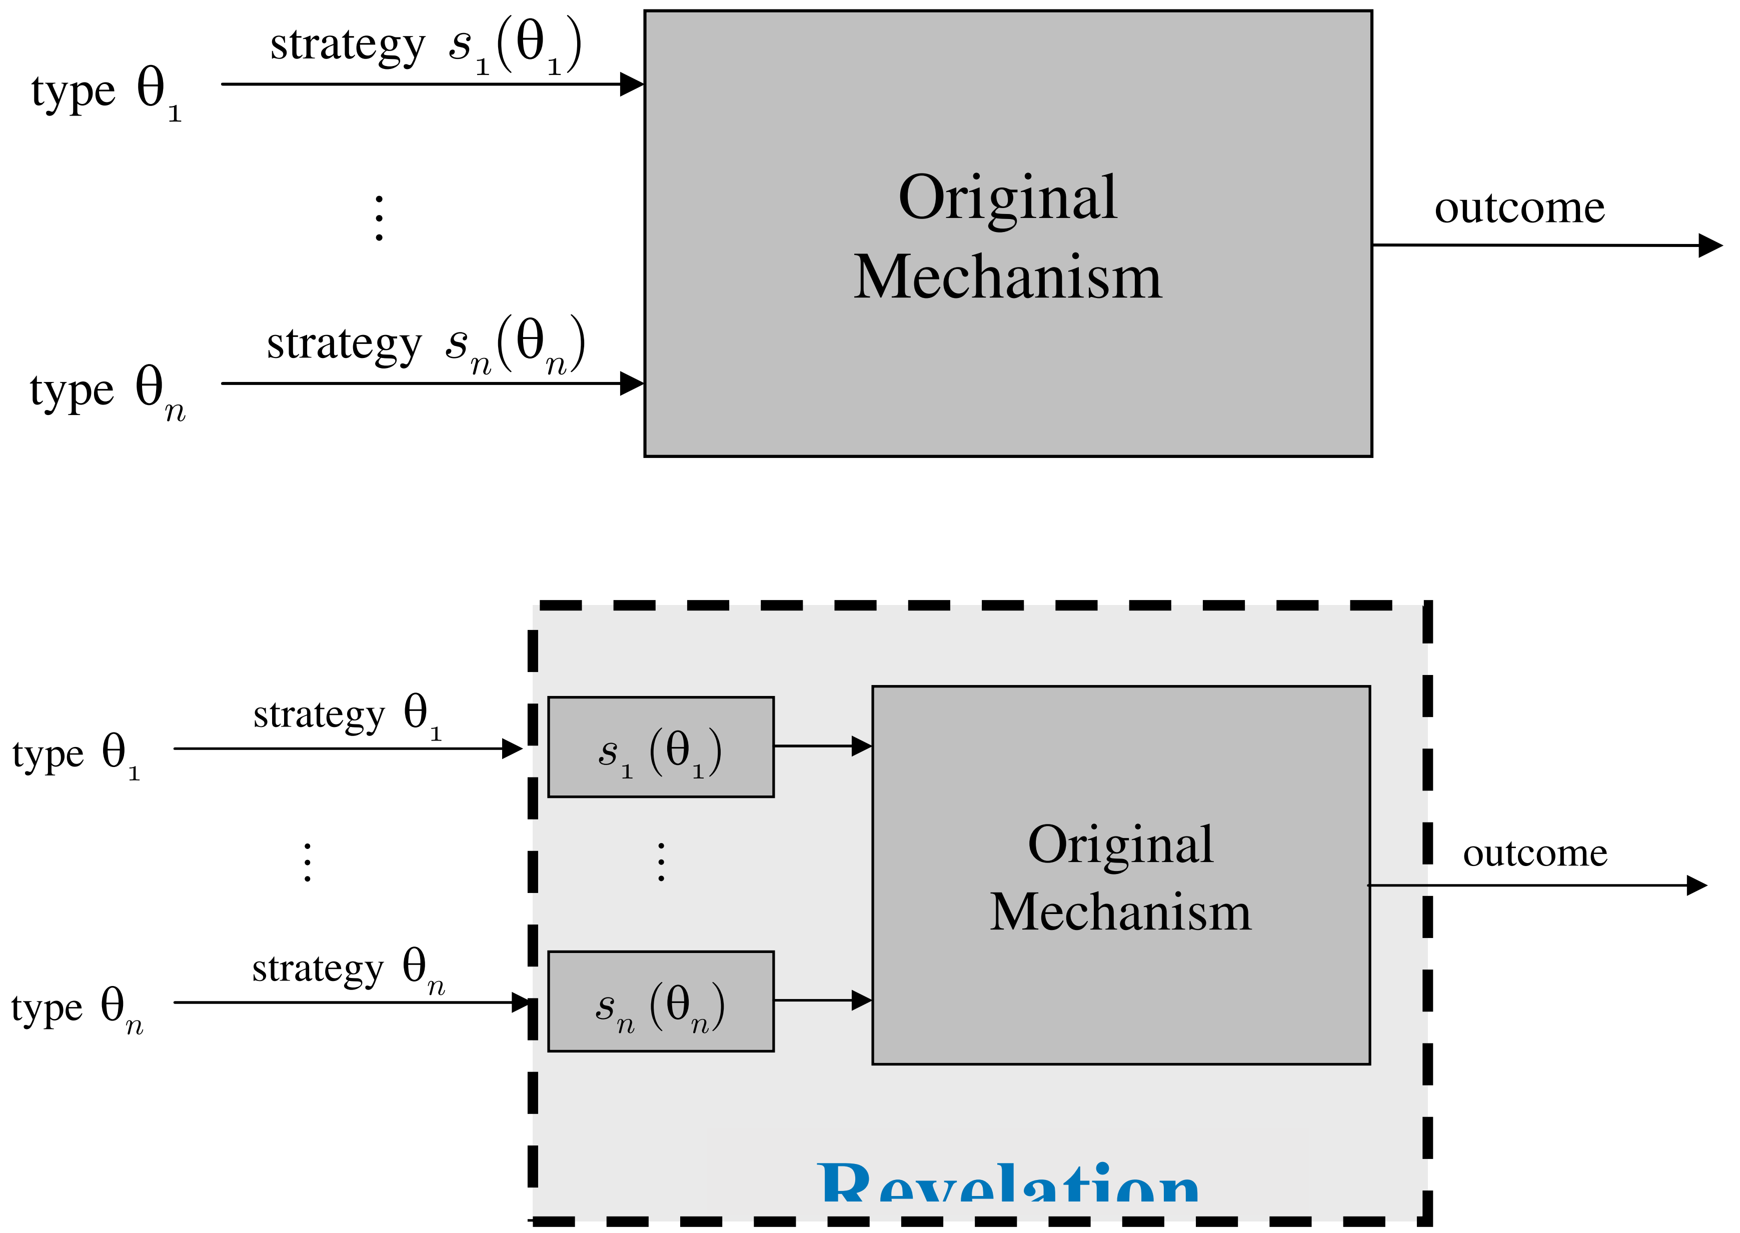

(Image: Shoham & Leyton-Brown 2008)

Recap: General Dominant-Strategy Implementation

Theorem: (Gibbard-Satterthwaite)
Consider any social choice function C over N and O. If $|O| > 2$ (there are at least **three** outcomes),

1. C is **onto**; that is, for every outcome $o \in O$ there is a preference profile $[\succ]$ such that $C([\succ]) = o$ (this is sometimes called **citizen sovereignty**), and

2. C is dominant-strategy **truthful**,

then C is **dictatorial**.

Recap:
Quasilinear Preferences

Definition:

Agents have **quasilinear preferences** in an n-player Bayesian game setting when

1. the set of outcomes is $O = X \times \mathbb{R}^n$ for a finite set X,
2. the utility of agent i given type profile θ for an element $(x, p) \in O$ is $u_i\big((x, p), \theta\big) = v_i(x, \theta) - f_i(p_i)$, where
3. $v_i : X \times \Theta \to \mathbb{R}$ is an **arbitrary** function, and
4. $f_i : \mathbb{R} \to \mathbb{R}$ is a **monotonically increasing** function.

Recap:
Direct Quasilinear Mechanism

Definition:

A **direct quasilinear mechanism** is a pair (χ, p), where

- $\chi : \Theta \to \Delta(X)$ is the **choice rule** (often called the **allocation rule**), which maps from a profile of reported types to a distribution over nonmonetary outcomes, and

- $p : \Theta \to \mathbb{R}^n$ is the **payment rule**, which maps from a profile of reported types to a payment for each agent.

Paper Presentations

Paper presentations start next week:

- There will be **1 or 2** presentations per class
    - (Rabin 2000 is rescheduled to Oct 31)
- Each paper is allocated **35 minutes** for talk + questions
    - Budget for about a 20-25 minute talk and 10-15 minutes for questions
- Summarize the **important parts** of the paper
- **Paper summaries** are due **before class** starts
    - Submit via Gradescope
    - See the course assignments page for details on what they should include

Lecture Outline

1. Recap & Logistics
2. Risk Attitudes
3. Efficient Quasilinear Mechanisms
4. Properties of Quasilinear Mechanisms

Value for Money

$$u_i\left((x,p),\theta\right) = v_i(x,\theta) - f_i(p_i)$$

- f_i represents agent i's **value for money**
  - **Question:** Why do we need a function instead of just a coefficient?
- The amount that you value $1 will typically depend on how much money you **already have**:
  - An extra $100 can change your life if you are starving
  - If you are a millionaire, you might not even notice the difference
- A **nonlinear** value for money can yield differing attitudes toward **risk**

Risk Aversion

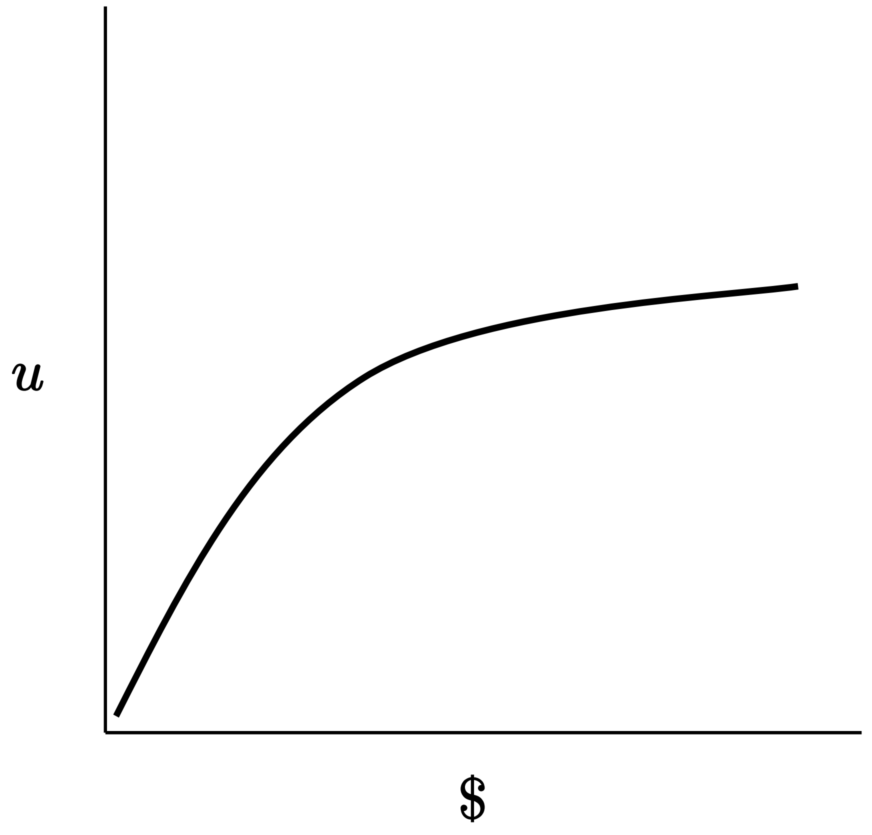

(c) Risk aversion

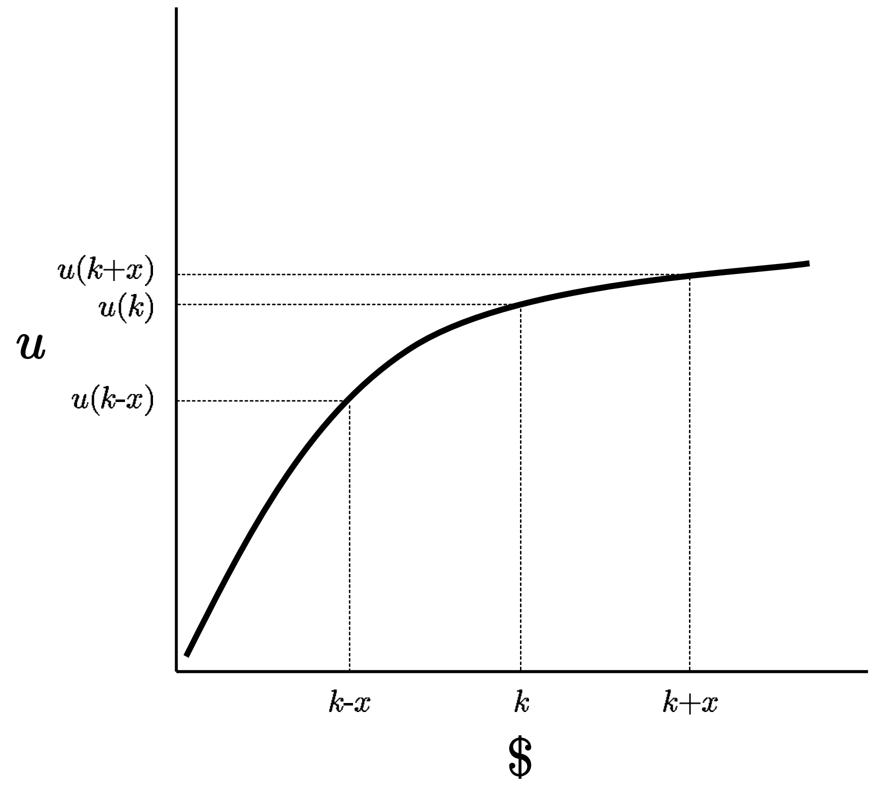

(d) Risk aversion: fair lottery

- A **concave** f_i models **decreasing marginal value** of money
- An agent with concave f_i is said to be **risk averse**, because they will **strictly prefer** to receive a lottery's **expected value** rather than to play the lottery
- **Question:** Is risk aversion irrational?

(Image: Shoham & Leyton-Brown 2008)

Risk Seeking

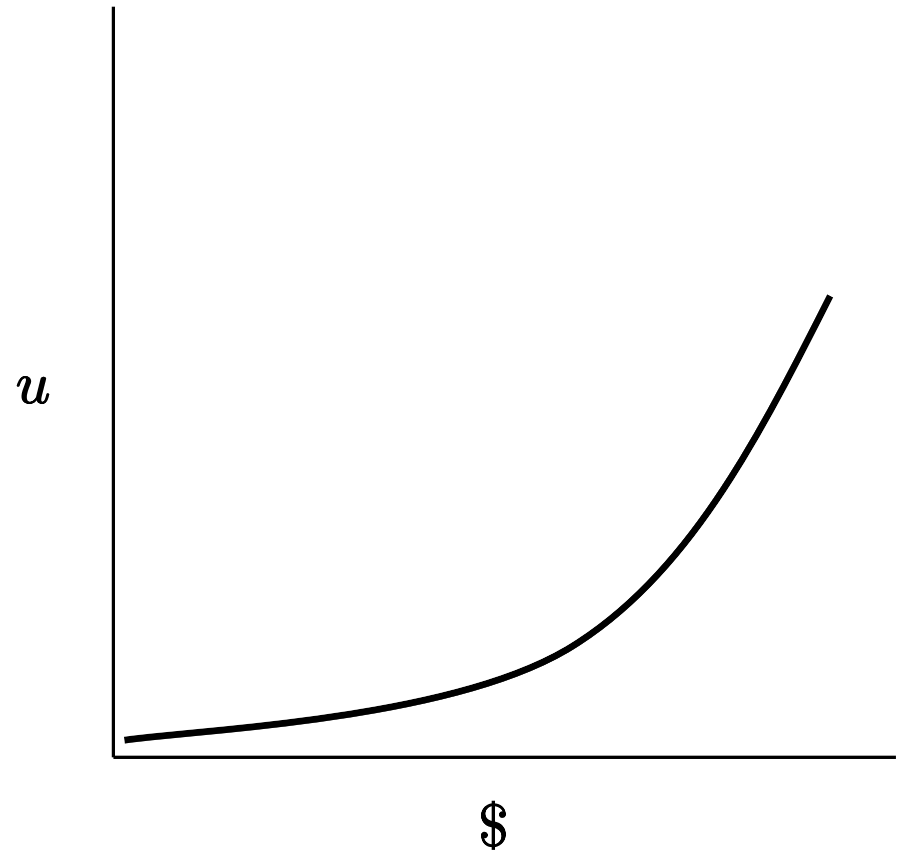

(e) Risk seeking

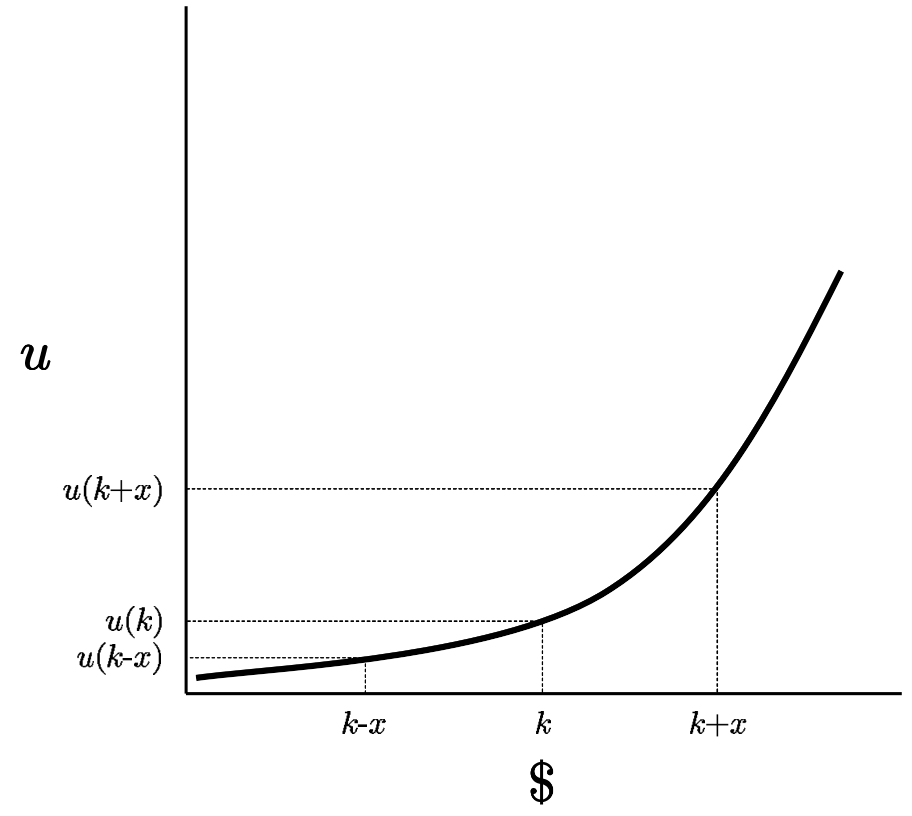

(f) Risk seeking: fair lottery

- A **convex** f_i models **increasing marginal value** of money
- An agent with convex f_i is said to be **risk seeking**, because they will **strictly prefer** to **play the lottery** rather than to receive a lottery's expected value
- **Question:** Is risk seeking irrational?

(Image: Shoham & Leyton-Brown 2008)

Risk Neutrality

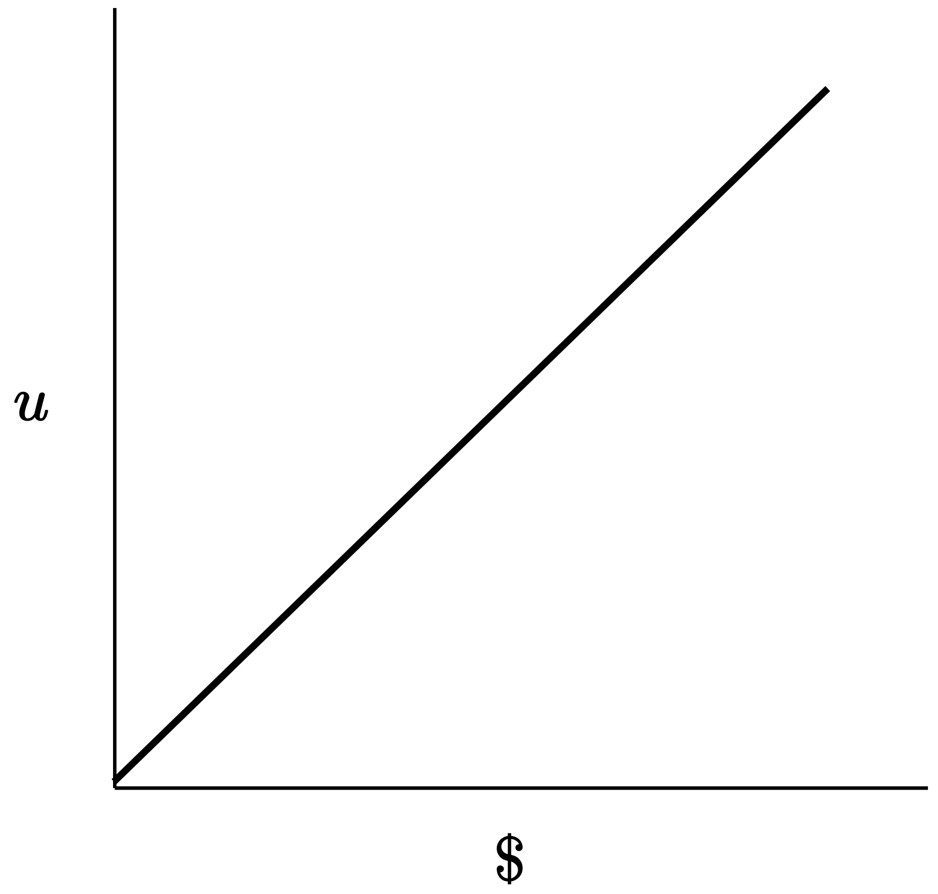

(a) Risk neutrality

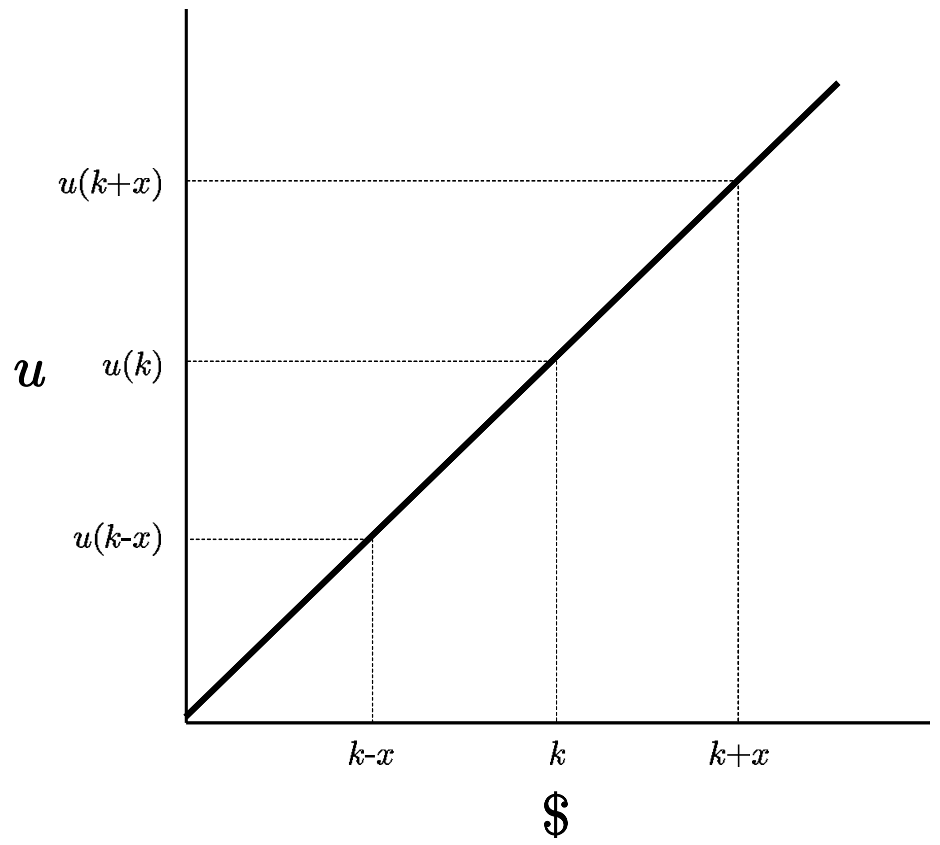

(b) Risk neutrality: fair lottery

- A **linear** f_i models **constant marginal value** of money
- An agent with linear f_i is said to be **risk neutral**, because they will be **indifferent** between receiving a lottery's **expected value** or playing the lottery

(Image: Shoham & Leyton-Brown 2008)

Transferable Utility

- Consider two agents i and j, who are both **risk-neutral**
- **Question:** Must they have the same value for money?
  - **No**, because they might have **different slopes:**
- When we additionally assume that $\beta_i = \beta_j$ for all $i, j \in N$, we say that the agents have **transferable utility**
  - Because I can increase i's utility by **exactly the amount** that I decrease j's utility, just by moving money from j to i
- Transferable utility is a **standard assumption** in quasilinear settings

$$f_i(x) = \beta_i x$$
$$f_j(x) = \beta_j x$$
$$\beta_i \neq \beta_j$$

Valuations

Definition:
A Bayesian game exhibits **conditional utility independence** if for all agents $i \in N$, all outcomes $o \in O$, and all pairs of joint types $\theta, \theta' \in \Theta$, it holds that $\theta_i = \theta'_i \implies u_i(o, \theta) = u_i(o, \theta')$.

- When this condition holds, we can write utility as $u_i(o, \theta_i)$
- Can equivalently refer to an agent's **valuation**: $v_i(x) = u_i(x, \theta_i)$.
- **Question:** When might this condition fail to hold?
- **Question:** Can we refer to an agent's valuation when this condition fails?

$$v_i(x) = u_i(x, \theta)$$

Groves Mechanisms

Definition:

Groves mechanisms are direct quasilinear mechanisms (χ, p) for which

$$\chi(\hat{v}) = \arg\max_x \sum_i \hat{v}_i(x)$$

$$p_i(\hat{v}) = h_i(\hat{v}_{-i}) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

- Where h_i is an **arbitrary** function of the reports of the other agents
- Groves mechanisms implement **any social welfare maximizing** choice function in dominant strategies

Proof Sketch: Dominant Strategies

1. Suppose that every other agent j declares arbitrary $\hat{v}_j$
2. Agent i wants to report $\hat{v}_i$ that solves $\max_{\hat{v}_i} \left(v_i\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) - p_i(\hat{v}_i, \hat{v}_{-i}) \right)$.
3. Substitute p_i: $\max_{\hat{v}_i} \left(v_i\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) \cancel{- h_i(\hat{v}_{-i})} + \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) \right)$
4. $h_i(\hat{v}_{-i})$ doesn't depend on $\hat{v}_i$

Proof Sketch #2

5. So i should report $\displaystyle \arg\max_{\hat{v}_i} \left(v_i\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) + \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) \right)$

6. But Groves will choose $\displaystyle \arg\max_{\chi(\hat{v}_i, \hat{v}_{-i})} \left(\hat{v}_i\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) + \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) \right)$

7. So i should report $\hat{v}_i = v_i$. ■

Dominant strategies, because this argument is for **arbitrary** $\hat{v}_{-i}$.

Vickrey-Clarke-Groves Mechanism

Definition:
The **Vickery-Clarke-Groves mechanism** is a direct quasilinear mechanism (χ, p), where

$$\chi(\hat{v}) = \arg\max_{x} \sum_{i} \hat{v}_i(x)$$

$$p_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}_{-i})) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

- i.e., it's a Groves mechanism with $h_i(\hat{v}_{-i}) = \Sigma_{j \neq i} \hat{v}_j(\chi(\hat{v}_{-i}))$.

- Each agent pays their **externality**: difference between *other* agents' utility **if i weren't there** and the *other* agents' utility given that ***i* is there.**

- **Question:** Why don't we use this for **everything**?

Second Price Auctions *Are* VCG

The **second price auction** is **VCG** in the quasilinear **single-item auction setting**:

- Agents are **not permitted unrestricted** preferences over the outcome space of allocations and payments
- Object is awarded to agent with **highest valuation**; this maximizes the sum of (non-monetary) agent valuations for the outcome
- Externality of **winning agent** is the value that **next-highest-valuation agent** could have gotten by winning the auction
- Externality of **losing agent** is nothing; if they weren't there, the outcome would be no different

Externalities: Example

1. Who wins the second-price auction?
   i.e., $\chi(\hat{v})$

2. Who would win if Alice weren't in the auction?
   i.e., $\chi(\hat{v}_{-Alice})$

3. How much does Alice pay?

4. What is the VCG payment?

$$v_{Alice}(\text{Alice gets object}) = 10$$
$$v_{Bob}(\text{Bob gets object}) = 6$$
$$v_{Carol}(\text{Carol gets object}) = 3$$
$$v_{Dave}(\text{Dave gets object}) = 1$$

$$\sum_{j \neq Alice} \hat{v}_j(\text{Bob}) - \sum_{j \neq Alice} \hat{v}_j(\text{Alice}) = (6 + 0 + 0) - (0 + 0 + 0) = 6$$

Mechanism Properties

Definition:
A quasilinear mechanism is **truthful** if it is direct and $\forall i \in N \, \forall v_i$, agent i's equilibrium strategy is to adopt the strategy $\hat{v}_i = v_i$.

Definition:
A quasilinear mechanism is **Pareto efficient**, or just **efficient**, if for all v in equilibrium it selects a choice x such that

$$\forall x' \sum_i v_i(x) \geq \sum_i v_i(x').$$

Budget Balance

Definition:
A quasilinear mechanism is **weakly budget balanced** when

$$\forall v, \sum_i p_i(s^*(v)) \geq 0,$$

where s^* is the equilibrium strategy profile.

Individual Rationality

Definition:
A quasilinear mechanism is ***ex-interim* individually rational** when

$$\forall i \forall v_i, \mathbb{E}_{v_{-i}|v_i}\left[v_i(\chi(s_i(v_i), s_{-i}(v_{-i}))) - p_i(s_i(v_i), s_{-i}(v_{-i}))\right] \geq 0.$$

All Efficient Dominant Strategy Mechanisms are Groves Mechanisms

Theorem: (Green-Laffont)

An **efficient** social choice function $C : \mathbb{R}^{X \times N} \to X \times \mathbb{R}^N$ can be implemented in dominant strategies for agents with **unrestricted quasilinear utilities only if**

$$p_i(\hat{v}) = h_i(\hat{v}_{-i}) - \sum_{j \neq i} \hat{v}_j \left(\chi(\hat{v})\right).$$

One Last Impossibility Result

Theorem: (Myerson-Satterthwaite)
No Bayes-Nash incentive-compatible mechanism is always simultaneously efficient, weakly budget-balanced, and ex-interim individually rational, even if agents are restricted to quasilinear utility functions.

- It does turn out to be possible to get any **two of the three**
- **Question:** Wait a minute, doesn't the second-price auction satisfy all three conditions?

Summary

- When agents are restricted to **quasilinear preferences**, social choice functions can be implemented in **dominant strategies**
- **Groves mechanisms** are the unique class of mechanisms that implement efficient social choice functions in dominant strategies
    - **VCG** is the pre-eminent Groves mechanism
    - Second-price auctions turn out to be VCG in the single-item auction setting
- You can only have **two** of **efficiency**, **weak budget balance**, and ***ex-interim* individual rationality**, even in the quasilinear setting