Mechanism Design

2: Efficient Mechanisms

Egor Starkov

Københavns Unversitet
Fall 2023

This slide deck:

Quasilinear Preferences

Assumption: Quasilinear setting/Transferable utility

- Instead of allowing all possible preferences, adopt a special structure.
- Instead of $x \in X$ describing everything related to outcome, split it into:
  - $k(\theta) \in K$, "real/material outcome" a.k.a. allocation
  - $t(\theta) \in \mathbb{R}^N$, transfers/payments
- Instead of arbitrary $u_i(x, \theta)$ focus on quasilinear preferences:

$$u_i(x, \theta_i) = v_i(k, \theta_i) - t_i$$

- S.c.f. is $f(\theta) = (k(\theta), t_1(\theta), ..., t_N(\theta))$

Quasilinear Preferences

- Common interpretation: transfers=payments. This comes with a bunch of assumptions:
  - Monetary transfers always available,
  - individual utility is linear in money (risk-neutrality),
  - marginal social utility of money is constant across types and people and independent of allocation.
- All three are sometimes restrictive, the latter two especially.
- However, monetary payments are not necessary! Anything that i cares about that is not directly included in the allocation k can be used to adjust i's utility as needed!
  - i's time, i's effort, promises to i, etc

Efficient Implementation

- A frequent question: "Dr.Professor, how can we as society implement the efficient outcome?"
- Reminder: efficient outcome $x^*(\theta) = (k^*(\theta), t^*(\theta))$ is

$$x^*(\theta) = \arg\max_x \sum_{i=0}^{N} u_i(x, \theta_i) = \arg\max_{(k,t)} \sum_{i=0}^{N} [v_i(k, \theta_i) - t_i]$$

- Transfers just reallocate utility across agents, so focus on efficient allocation $k^*(\theta)$:

$$k^*(\theta) = \arg\max_k \sum_{i=0}^{N} v_i(k, \theta_i)$$

- Note that we can include $i = 0$ into welfare calculations. This can capture designer's preferences or any social costs/benefits not captured by individual agents (e.g., cost of implementing a public project)

Efficient Implementation

- How do we do that?
  - We already know that it's enough to consider direct revelation mechanisms.
  - We have the desired allocation rule k^*, but we can design the transfers t – the problem is not just "check whether s.c.f. x^* is IC", but "is there such t that k^* is IC?"
- What we as designers want:

$$\max \sum_{i=0}^{N} v_i(k, \theta_i)$$

- What agent i wants:

$$\max v_i(k, \theta_i) - t_i$$

- How to reconcile the two?

This slide deck:

VCG Mechanism: intro

- We now introduce the VCG mechanism that DSIC-implements the efficient allocation. We shall do it in a few steps.
- **NOTE**: while the broad idea behind the VCG mechanism is the same everywhere, the exact definition of the VCG mechanism differs in different sources (textbooks).

VCG Mechanism: Groves' Transfers

- More formally, the problem of agent i of type θ_i is:

$$\max_{\hat{\theta}_i} \left\{ v_i(k^*(\hat{\theta}_i, \theta_{-i}), \theta_i) - t_i(\hat{\theta}_i, \theta_{-i}) \right\}$$

- Try Groves' transfers:

$$t_i^G(\theta) \equiv -\left(\sum_{j \neq i} v_j(k^*(\theta_i, \theta_{-i}), \theta_j) \right) + h_i(\theta_{-i})$$

- Agent's problem is now

$$\max_{\hat{\theta}_i} \left\{ v_i(k^*(\hat{\theta}_i, \theta_{-i}), \theta_i) + \left(\sum_{j \neq i} v_j(k^*(\hat{\theta}_i, \theta_{-i}), \theta_j) \right) - h_i(\theta_{-i}) \right\}$$

VCG Mechanism: Groves' Transfers

- More formally, the problem of agent i of type θ_i is:

$$\max_{\hat{\theta}_i} \left\{ v_i(k^*(\hat{\theta}_i, \theta_{-i}), \theta_i) - t_i(\hat{\theta}_i, \theta_{-i}) \right\}$$

- Try Groves' transfers:

$$t_i^G(\theta) \equiv -\left(\sum_{j \neq i} v_j(k^*(\theta_i, \theta_{-i}), \theta_j) \right) + h_i(\theta_{-i})$$

- Agent's problem is now

$$\max_{\hat{\theta}_i} \left\{ v_i(k^*(\hat{\theta}_i, \theta_{-i}), \theta_i) + \left(\sum_{j \neq i} v_j(k^*(\hat{\theta}_i, \theta_{-i}), \theta_j) \right) - h_i(\theta_{-i}) \right\}$$

VCG Mechanism: Groves' Transfers

- More formally, the problem of agent i of type θ_i is:

$$\max_{\hat{\theta}_i} \left\{ v_i(k^*(\hat{\theta}_i, \theta_{-i}), \theta_i) - t_i(\hat{\theta}_i, \theta_{-i}) \right\}$$

- Try Groves' transfers:

$$t_i^G(\theta) \equiv -\left(\sum_{j \neq i} v_j(k^*(\theta_i, \theta_{-i}), \theta_j) \right) + h_i(\theta_{-i})$$

- Agent's problem is now

$$\max_{\hat{\theta}_i} \left\{ v_i(k^*(\hat{\theta}_i, \theta_{-i}), \theta_i) + \left(\sum_{j \neq i} v_j(k^*(\hat{\theta}_i, \theta_{-i}), \theta_j) \right) - h_i(\theta_{-i}) \right\}$$

VCG Mechanism: Groves' Transfers

- Agent's problem is now

$$\max_{\hat{\theta}_i} \left\{ \sum_{j \in N} v_j(k^*(\hat{\theta}_i, \theta_{-i}), \theta_j) - h_i(\theta_{-i}) \right\}$$

- Every agent i chooses report $\hat{\theta}_i$ to maximize welfare!
  - Optimal to report true $\hat{\theta}_i$,
  - for any θ_{-i}.
- Crucial that $h_i(\theta_{-i})$ does not depend on i's report.

VCG Mechanism: Example

Example (Moon Base)

- N citizens decide whether to build a Moon base which costs c
- citizen i has private valuation θ_i for the base and quasilinear utility (so if base built then $v_i = \theta_i$, otherwise $v_i = 0$)

- What are Groves' transfers? (Take $h_i(\theta_{-i}) \equiv 0$.)
- The incentives are there... but at what cost?

VCG Mechanism: Clarke Term

- A suggestion for $h_i(\theta_{-i})$ made by Clarke ("pivot mechanism"):

$$h_i(\theta_{-i}) = \sum_{j \neq i} v_j(k^{-i}(\theta_{-i}), \theta_j),$$

$$\text{where } k^{-i}(\theta_{-i}) \in \arg\max_k \sum_{j \neq i} v_j(k, \theta_j).$$

- **NOTE**: it is the default allocation rules $k^{-i}(\theta_{-i})$ that each textbook defines differently (or replaces with other things). My version is quite robust, but you can use other default rules if they make more sense in a given setting (so long as the rule for i is independent of θ_i)

- Resulting VCG transfers:

$$t_i^{VCG}(\theta) \equiv -\left(\sum_{j \neq i} v_j(k^*(\theta_i, \theta_{-i}), \theta_j) \right) + \sum_{j \neq i} v_j(k^{-i}(\theta_{-i}), \theta_j)$$

VCG Mechanism: Final Transfers

$$t_i^{VCG}(\theta) = -\left(\sum_{j\neq i} v_j(k^*(\theta_i,\theta_{-i}),\theta_j)\right) + \sum_{j\neq i} v_j(k^{-i}(\theta_{-i}),\theta_j)$$

- What's the big idea?
  - Agent i receives the externality his report imposes on others (mind the signs).
  - i's transfer is non-zero only if his presence affects the decision k.
  - Note that i cannot misreport θ_i and get lower transfer without also changing k.
- What are VCG transfers in the Moon Base question?

VCG Mechanism: Example

Example (Auction)

- One indivisible item to be allocated among N bidders.
- Bidder i's valuation is θ_i (private info).
- What is the VCG mechanism?

- VCG mechanism is the second-price auction (efficient and DSIC).
- Also known as the Vickrey auction (the V in VCG).

VCG aftermath

- We have an easy recipe to implement the efficient outcome in dominant strategies.
- Any problems?

This slide deck:

Feature example: bilateral trade

Example (Bilateral Trade)

- One indivisible good.
- Two agents: buyer and seller.
- Private valuations $\theta_b, \theta_s \in [0, 1]$ resp.
- Find the VCG transfers (take no trade as efficient when $\theta_s = \theta_b$).

Feature example: bilateral trade

- If you did everything correctly, you'll get

$$t_b^{VCG}(\theta) = \theta_s \cdot \mathbb{I}\{\theta_s < \theta_b\}$$
$$t_s^{VCG}(\theta) = \theta_b \cdot \mathbb{I}\{\theta_s \geq \theta_b\}$$

Feature example: bilateral trade

- If you did everything correctly, you'll get

$$t_b^{VCG}(\theta) = \theta_s \cdot \mathbb{I}\{\theta_s < \theta_b\}$$
$$t_s^{VCG}(\theta) = \theta_b \cdot \mathbb{I}\{\theta_s \geq \theta_b\}$$

- The seller pays to keep the good and doesn't get anything from selling it. Good deal?

Individual rationality

- In many settings can't force players to participate in mechanism:

Definition (IR)

A mechanism Γ is:

- interim individually rational if $\mathbb{E}_{\theta_{-i}}\left[u_i(\theta_i, \theta_{-i})\right] \geq \underline{U}_i(\theta_i)$ for all θ_i;
- ex post individually rational if $u_i(\theta_i, \theta_{-i}) \geq \underline{U}_i(\theta_i)$ for all θ.

- $\underline{U}_i(\theta_i)$ is the outside option of type θ_i
  - (in bilateral trade: $\underline{U}_s(\theta_s) = \theta_s$)
- expectation means that distribution of θs now matters!
  - (whether a mechanism is DSIC does not depend on the distr-n; but whether it is IR does)

Detour – brief review

- ex ante $=$ i knows nothing;
- ex interim $=$ i knows θ_i;
- ex post $=$ i knows θ_i and θ_{-i}.
- We'll mostly work with interim IR;
- ex post IR is also sometimes used in the literature.

Budget balance

- VCG for bilateral trade example is not IR for seller (outside option $=$ keep the good).

Budget balance

- VCG for bilateral trade example is not IR for seller (outside option = keep the good).
- If we want mechanism to be IR, easy solution is to decrease $t_i(\theta)$ by a lot, for all θ.
- But that's expensive – want mechanism to be budget balanced:

Budget balance

- VCG for bilateral trade example is not IR for seller (outside option $=$ keep the good).
- If we want mechanism to be IR, easy solution is to decrease $t_i(\theta)$ by a lot, for all θ.
- But that's expensive – want mechanism to be budget balanced:

Definition (BB)

- Mechanism Γ is ex ante budget balanced if $\mathbb{E}_\theta\left[\sum_{i=1}^N t_i(\theta)\right] \geq 0$;
- Mechanism Γ is ex post budget balanced if $\sum_{i=1}^N t_i(\theta) \geq 0$ for all θ.

- Mechanism is exactly BB if the above hold with equalities.
- If Γ is ex post BB then it is ex ante BB (prove).

IR vs BB

- Fundamental tension between IR and BB.
- We want to ask the following question (within our bilateral trade example, in particular):

Does there exist a mechanism that is: efficient, DSIC, IR, BB?

IR vs BB

- Fundamental tension between IR and BB.
- We want to ask the following question (within our bilateral trade example, in particular):

> Does there exist a mechanism that is: efficient, DSIC, IR, BB?

- We know VCG was not IR, but that's just one mechanism. Can we say whether any other mechanisms satisfy all requirements?
  - Not in most general case*, but all examples (trade, auction, pub.project) fit a much narrower model where we can.
  - *though see Prop 23.C.5 in MWG

This slide deck:

The Euclidean model

Assumption: Euclidean setting

Make the following assumptions on top of quasilinearity:

- $\theta_i \in \Theta_i = [\underline{\theta}_i, \bar{\theta}_i] \subseteq \mathbb{R}$, full support;
- $k \in K \subseteq \mathbb{R}^N$, K compact, convex set;
- $u_i(x, \theta_i) = \theta_i k_i - t_i$.

- I'll call the above the Euclidean model (not standard name).
- We'll derive Payoff-equivalence as a necessary condition for Γ to be DSIC in Euclidean model. It's a cool property on its own and will help answer the question about BB/IR mechanisms later.
- Given Γ, denote $U_i(\theta_i, \theta_{-i}) \equiv u_i\left(x(\theta_i, \theta_{-i}), \theta_i\right)$.

Monotonicity

- Assume Γ is a direct mechanism (or consider its direct equivalent).
- Play a bit with i's IC (truthtelling constraint): for any $i, \theta_i, \hat{\theta}_i, \theta_{-i}$,

$$U_i(\theta_i, \theta_{-i}) \geq u_i\left(x(\hat{\theta}_i, \theta_{-i}), \theta_i\right)$$

Monotonicity

- Assume Γ is a direct mechanism (or consider its direct equivalent).
- Play a bit with i's IC (truthtelling constraint): for any $i, \theta_i, \hat{\theta}_i, \theta_{-i}$,

$$\begin{aligned} U_i(\theta_i, \theta_{-i}) &\geq u_i\left(x(\hat{\theta}_i, \theta_{-i}), \theta_i\right) \\ &\equiv \theta_i k_i(\hat{\theta}_i, \theta_{-i}) - t_i(\hat{\theta}_i, \theta_{-i}) \end{aligned}$$

Monotonicity

- Assume Γ is a direct mechanism (or consider its direct equivalent).
- Play a bit with i's IC (truthtelling constraint): for any $i, \theta_i, \hat{\theta}_i, \theta_{-i}$,

$$
\begin{aligned}
U_i(\theta_i, \theta_{-i}) \geq u_i\left(x(\hat{\theta}_i, \theta_{-i}), \theta_i\right) \\
&\equiv \theta_i k_i(\hat{\theta}_i, \theta_{-i}) - t_i(\hat{\theta}_i, \theta_{-i}) \\
&= \hat{\theta}_i k_i(\hat{\theta}_i, \theta_{-i}) - t_i(\hat{\theta}_i, \theta_{-i}) + \left(\theta_i - \hat{\theta}_i\right) k_i(\hat{\theta}_i, \theta_{-i})
\end{aligned}
$$

Monotonicity

- Assume Γ is a direct mechanism (or consider its direct equivalent).
- Play a bit with i's IC (truthtelling constraint): for any $i, \theta_i, \hat{\theta}_i, \theta_{-i}$,

$$\begin{aligned}
U_i(\theta_i, \theta_{-i}) &\geq u_i\left(x(\hat{\theta}_i, \theta_{-i}), \theta_i\right) \\
&\equiv \theta_i k_i(\hat{\theta}_i, \theta_{-i}) - t_i(\hat{\theta}_i, \theta_{-i}) \\
&= \hat{\theta}_i k_i(\hat{\theta}_i, \theta_{-i}) - t_i(\hat{\theta}_i, \theta_{-i}) + \left(\theta_i - \hat{\theta}_i\right) k_i(\hat{\theta}_i, \theta_{-i}) \\
&= U_i(\hat{\theta}_i, \theta_{-i}) + \left(\theta_i - \hat{\theta}_i\right) k_i(\hat{\theta}_i, \theta_{-i})
\end{aligned}$$

Monotonicity

- Assume Γ is a direct mechanism (or consider its direct equivalent).
- Play a bit with i's IC (truthtelling constraint): for any $i, \theta_i, \hat{\theta}_i, \theta_{-i}$,

$$\begin{aligned}
U_i(\theta_i, \theta_{-i}) &\geq u_i\left(x(\hat{\theta}_i, \theta_{-i}), \theta_i\right) \\
&\equiv \theta_i k_i(\hat{\theta}_i, \theta_{-i}) - t_i(\hat{\theta}_i, \theta_{-i}) \\
&= \hat{\theta}_i k_i(\hat{\theta}_i, \theta_{-i}) - t_i(\hat{\theta}_i, \theta_{-i}) + \left(\theta_i - \hat{\theta}_i\right) k_i(\hat{\theta}_i, \theta_{-i}) \\
&= U_i(\hat{\theta}_i, \theta_{-i}) + \left(\theta_i - \hat{\theta}_i\right) k_i(\hat{\theta}_i, \theta_{-i})
\end{aligned}$$

- In the end:

$$U_i(\theta_i, \theta_{-i}) \geq U_i(\hat{\theta}_i, \theta_{-i}) + \left(\theta_i - \hat{\theta}_i\right) k_i(\hat{\theta}_i, \theta_{-i}).$$

- Similarly, type $\hat{\theta}_i$ should not want to report θ_i:

$$U_i(\hat{\theta}_i, \theta_{-i}) \geq U_i(\theta_i, \theta_{-i}) + \left(\hat{\theta}_i - \theta_i\right) k_i(\theta_i, \theta_{-i}).$$

Monotonicity

- Combining the two for $\theta_i > \hat{\theta}_i$, we get

$$k_i(\theta_i, \theta_{-i}) \geq \frac{U_i(\theta_i, \theta_{-i}) - U_i(\hat{\theta}_i, \theta_{-i})}{\theta_i - \hat{\theta}_i} \geq k_i(\hat{\theta}_i, \theta_{-i}), \tag{1}$$

Monotonicity

- Combining the two for $\theta_i > \hat{\theta}_i$, we get

$$k_i(\theta_i, \theta_{-i}) \geq \frac{U_i(\theta_i, \theta_{-i}) - U_i(\hat{\theta}_i, \theta_{-i})}{\theta_i - \hat{\theta}_i} \geq k_i(\hat{\theta}_i, \theta_{-i}), \tag{1}$$

- meaning $k_i(\theta_i, \theta_{-i}) \geq k_i(\hat{\theta}_i, \theta_{-i})$ – allocation rule must be monotone.
- DSIC: "Those who value things more should get more things."
- Monotonicity is necessary for f to be DSIC in Euclidean settings.
- From monotonicity we can build up to payoff equivalence, the second cool result in mechanism design (after revelation principle, not monotonicity).

Payoff Equivalence

- $k_i(\theta_i, \theta_{-i})$ is monotone in θ_i, hence continuous a.e.: $\lim_{\hat{\theta}_i \to \theta_i} k_i(\hat{\theta}_i, \theta_{-i}) = k_i(\theta_i, \theta_{-i})$.

Payoff Equivalence

- $k_i(\theta_i, \theta_{-i})$ is monotone in θ_i, hence continuous a.e.: $\lim_{\hat{\theta}_i \to \theta_i} k_i(\hat{\theta}_i, \theta_{-i}) = k_i(\theta_i, \theta_{-i})$.
- Together with the big inequality (1) this means that a.e. we have

Payoff Equivalence

- $k_i(\theta_i, \theta_{-i})$ is monotone in θ_i, hence continuous a.e.: $\lim_{\hat{\theta}_i \to \theta_i} k_i(\hat{\theta}_i, \theta_{-i}) = k_i(\theta_i, \theta_{-i})$.
- Together with the big inequality (1) this means that a.e. we have

$$\frac{\partial U_i(\theta_i, \theta_{-i})}{\partial \theta_i} = \lim_{\hat{\theta}_i \to \theta_i} \frac{U_i(\theta_i, \theta_{-i}) - U_i(\hat{\theta}_i, \theta_{-i})}{\theta_i - \hat{\theta}_i} = k_i(\theta_i, \theta_{-i}).$$

Payoff Equivalence

- $k_i(\theta_i, \theta_{-i})$ is monotone in θ_i, hence continuous a.e.: $\lim_{\hat{\theta}_i \to \theta_i} k_i(\hat{\theta}_i, \theta_{-i}) = k_i(\theta_i, \theta_{-i})$.
- Together with the big inequality (1) this means that a.e. we have

$$\frac{\partial U_i(\theta_i, \theta_{-i})}{\partial \theta_i} = \lim_{\hat{\theta}_i \to \theta_i} \frac{U_i(\theta_i, \theta_{-i}) - U_i(\hat{\theta}_i, \theta_{-i})}{\theta_i - \hat{\theta}_i} = k_i(\theta_i, \theta_{-i}).$$

- So if $k(\theta)$ is integrable in θ_i (e.g. if it's bounded) then for all θ_i

$$U_i(\theta_i, \theta_{-i}) = U_i(\underline{\theta}_i, \theta_{-i}) + \int_{\underline{\theta}_i}^{\theta_i} k_i(s, \theta_{-i}) ds$$

  (Note that the lower limit does not need to be $\underline{\theta}_i$ – it can be any other type.)
- This is the envelope representation of payoffs a.k.a. Mirrlees condition. From it we can immediately get revenue equivalence.

Payoff Equivalence

Theorem (Payoff Equivalence for DSIC Euclidean mechanisms)

In the Euclidean setting, for any two DSIC DRMs with $x = (k, t)$ *and* $x' = (k', t')$ *respectively: if* $k(\theta) = k'(\theta)$ *for all* θ *then* $t_i(\theta) = t'_i(\theta) + c_i(\theta_{-i})$ *for all* θ *for some* $c_i(\theta_{-i})$*.*

Payoff Equivalence

Theorem (Payoff Equivalence for DSIC Euclidean mechanisms)

In the Euclidean setting, for any two DSIC DRMs with $x = (k, t)$ and $x' = (k', t')$ respectively: if $k(\theta) = k'(\theta)$ for all θ then $t_i(\theta) = t_i'(\theta) + c_i(\theta_{-i})$ for all θ for some $c_i(\theta_{-i})$.

Proof. Given the envelope representation, invoke the definition of U_i:

$$v_i(k^*(\theta_i, \theta_{-i}), \theta_i) - t_i(\theta_i, \theta_{-i}) = U_i(\underline{\theta}_i, \theta_{-i}) + \int_{\underline{\theta}_i}^{\theta_i} k_i(s, \theta_{-i}) ds.$$

The above holds in any DSIC DRM. Take the two mechanisms in the statement, fix some θ, i, express $t_i(\theta)$ and $t_i'(\theta)$ from the above, and you will get that

$$t_i(\theta) = t_i'(\theta) - U_i(\underline{\theta}_i, \theta_{-i}) + U_i'(\underline{\theta}_i, \theta_{-i}),$$

where U_i and U_i' denote the eqm utilities in the two mechanisms. The last two terms only depend on i and θ_{-i} (but not θ_i), hence denoting them as $c_i(\theta_{-i})$ concludes the proof. $\square$

Payoff Equivalence: intuition

- Given allocation k (doesn't have to be efficient), utility of one type (usually "lowest" type) pins down utils of all types of player i given fixed θ_{-i}.

Payoff Equivalence: intuition

- Given allocation k (doesn't have to be efficient), utility of one type (usually "lowest" type) pins down utils of all types of player i given fixed θ_{-i}.
- Equivalently, have only one degree of freedom for i's transfers given θ_{-i}.
- Reminds you of anything?

Payoff Equivalence of Efficient Mechanisms

- Recall Groves' transfers: efficient k^* can be impl-d in DS by

$$t_i(\theta) = -\left(\sum_{j \neq i} v_j(k^*(\theta_i, \theta_{-i}), \theta_j)\right) + h_i(\theta_{-i})$$

$$= -\left(\sum_{j \neq i} \theta_j k_j^*(\theta_i, \theta_{-i})\right) + h_i(\theta_{-i})$$

for some $h_i(\theta_{-i})$.

Payoff Equivalence of Efficient Mechanisms

- Recall Groves' transfers: efficient k^* can be impl-d in DS by

$$
\begin{aligned}
t_i(\theta) &= -\left(\sum_{j\neq i} v_j(k^*(\theta_i, \theta_{-i}), \theta_j)\right) + h_i(\theta_{-i}) \\
&= -\left(\sum_{j\neq i} \theta_j k_j^*(\theta_i, \theta_{-i})\right) + h_i(\theta_{-i})
\end{aligned}
$$

  for some $h_i(\theta_{-i})$.

- Payoff equivalence implies that efficient k^* in a Euclidean model can ONLY be DS-implemented by some Groves' mechanism.

Payoff Equivalence: Beyond Euclidean

Both monotonicity and payoff equvalence hold beyond the Euclidean setting:

- various forms of monotonicity are necessary and sufficient for DSIC in quasilinear setting;
- payoff equivalence of DSIC mechanisms is generalizable beyond Euclidean (but you cannot get to general quasilinear setting);
- see Börgers for details:
  - ch.5: single-player problems,
  - ch.7: DSIC results building on ch.5.

Payoff Equivalence (DSIC): Conclusion

- So what does payoff equivalence tell us?
- Efficient allocation k^* can **only** be implemented using Groves' transfers...
  - ...but $h_i(\theta_{-i})$ still provides a lot of flexibility!
  - So it's hard to know whether VCG is the best mechanism or there are others.
- So let us *weaken* our implementation concept to obtain a **stronger** version of payoff equivalence.

Back to bilateral trade

- Remember how this detour started?

Example (Bilateral Trade)

- One indivisible good.
- Two agents: buyer and seller.
- Private valuations $\theta_b, \theta_s \in [0, 1]$ resp.
- Is there an efficient, DSIC, ex post IR, ex post BB mechanism?

Can we answer this question now?

This slide deck:

Payoff Equivalence in BIC

We now show payoff equivalence for BIC mechanisms in the Euclidean setting.
New assumption: types are independent across players: $\theta_i \perp \theta_{-i}$ for all i.

Theorem (Payoff Equivalence for BIC mechanisms)

For any two BIC DRMs with $x = (k, t)$ and $x' = (k', t')$ resp.:
if $\mathbb{E}_{\theta_{-i}} k_i(\theta_i, \theta_{-i}) = \mathbb{E}_{\theta_{-i}} k'_i(\theta_i, \theta_{-i})$ *for all* i, θ_i,
then $\mathbb{E}_{\theta_{-i}} t_i(\theta_i, \theta_{-i}) = \mathbb{E}_{\theta_{-i}} t'_i(\theta_i, \theta_{-i}) + h_i$ *for all* i, θ_i *for some* h_i.

- As before, implies that for given k (any, not just the efficient) we only have one degree of freedom for $t_i(\theta)$,
  - now "just one" instead of "just one given θ_{-i}".

Payoff Equivalence in BIC. Proof

- Let

$$\begin{aligned}\tilde{U}_i(\hat{\theta}_i, \theta_i) &\equiv \mathbb{E}_{\theta_{-i}}\left[u_i\left(x(\hat{\theta}_i, \theta_{-i}), \theta_i\right)|\theta_i\right] \\ &= \mathbb{E}_{\theta_{-i}}\left[\theta_i k_i(\hat{\theta}_i, \theta_{-i}) - t_i\left(\hat{\theta}_i, \theta_{-i}\right)|\theta_i\right].\end{aligned}$$

(do not confuse with U_i in Euclidean model for DS.)

Payoff Equivalence in BIC. Proof

- Let

$$\tilde{U}_i(\hat{\theta}_i, \theta_i) \equiv \mathbb{E}_{\theta_{-i}} \left[u_i \left(x(\hat{\theta}_i, \theta_{-i}), \theta_i \right) | \theta_i \right]$$
$$= \mathbb{E}_{\theta_{-i}} \left[\theta_i k_i(\hat{\theta}_i, \theta_{-i}) - t_i \left(\hat{\theta}_i, \theta_{-i} \right) | \theta_i \right].$$

(do not confuse with U_i in Euclidean model for DS.)

- Take full derivative w.r.t θ_i at $\hat{\theta}_i = \theta_i$:

$$\frac{d}{d\theta_i} \tilde{U}_i(\theta_i, \theta_i) = \frac{\partial}{\partial \hat{\theta}_i} \tilde{U}_i(\hat{\theta}_i, \theta_i) \bigg|_{\hat{\theta}_i = \theta_i} + \frac{\partial}{\partial \theta_i} \tilde{U}_i(\hat{\theta}_i, \theta_i) \bigg|_{\hat{\theta}_i = \theta_i}$$
$$= 0 + \mathbb{E}_{\theta_{-i}} \left[k_i(\theta_i, \theta_{-i}) | \theta_i \right]$$

The first term is zero because truthful report $\hat{\theta}_i = \theta_i$ maximizes $\tilde{U}_i(\hat{\theta}_i, \theta_i)$.

Payoff Equivalence in BIC. Proof

- Then by the Fundamental Theorem of Calculus (with $\bar{U}_i(\theta_i) \equiv \tilde{U}_i(\theta_i, \theta_i)$)

$$\begin{aligned}\bar{U}_i(\theta_i) &= \bar{U}_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\theta_i} \frac{d\bar{U}_i}{d\theta_i}(s,s)ds \\ &= \bar{U}_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\theta_i} \mathbb{E}_{\theta_{-i}}\left[k_i(s,\theta_{-i})|\theta_i\right] ds,\end{aligned}$$

  meaning that k and $\bar{U}_i(\underline{\theta}_i)$ pin down utilities $\bar{U}_i(\theta_i)$ for all θ_i. $\square$

- Remark: here we used a different argument to get $\frac{d\bar{U}_i(\theta_i,\theta_i)}{d\theta_i} = \mathbb{E}_{\theta_{-i}}\left[k(\theta_i,\theta_{-i})|\theta_i\right]$ compared to DSIC proof. Either argument can be used in both proofs.

Payoff Equivalence in BIC. Generalization

The proof is nice and illustrative in Euclidean setting.
Krishna and Maenner [2001] present a more general result for the following setting:

- quasilinear setting;
- independent types;
- $\Theta_i \subseteq \mathbb{R}^{K_i}$ is a convex set for every i
- $v_i(k, \theta_i)$ is convex in θ_i for all i.

Ex ante and ex post revenue in BIC

One cool thing about BIC mechanisms is that ex post BB is free if you have ex ante BB:

Theorem

In a quasilinear setting, for every direct mechanism $\Gamma = (\Theta, (k, t))$ that is BIC and ex ante BB, there exists a direct mechanism $\Gamma' = (\Theta, (k', t'))$ which is:

- *BIC,*
- *ex post BB,*
- *equivalent to Γ in the sense of: $k'(\theta) = k(\theta)$ for all θ and $\mathbb{E}_{\theta_{-i}} t'_i(\theta_i, \theta_{-i}) = \mathbb{E}_{\theta_{-i}} t_i(\theta_i, \theta_{-i})$ for all i and θ_i.*

For proof, see Prop 6.3 & Prop 3.6 in Börgers.

This slide deck:

Generalized VCG

- As it turns out, VCG can be (interim) IR with a slight modification... And it will be (ex ante) revenue-maximizing among all such mechanisms in that case.
- Enter generalized VCG [Krishna and Perry, 2000].

Generalized VCG

- As it turns out, VCG can be (interim) IR with a slight modification... And it will be (ex ante) revenue-maximizing among all such mechanisms in that case.
- Enter generalized VCG [Krishna and Perry, 2000].
- Define least charitable type $\tilde{\theta}_i$ as

$$\tilde{\theta}_i \in \arg\min_{\theta_i \in \Theta_i} \mathbb{E}_{\theta_{-i}} \left[\sum_{j=0}^{N} v_j(k^*(\theta_i, \theta_{-i}), \theta_j) - \underline{U}_i(\theta_i) \right]$$

(expectation taken w.r.t the common prior $\phi \in \Delta(\Theta)$)

Generalized VCG

GVCG mechanism is a DRM with the efficient allocation $k^*(\theta)$ and payments

$$t_i^{GVCG}(\theta) \equiv \sum_{j \neq i} v_j(k^*(\tilde{\theta}_i, \theta_{-i}), \theta_j) + v_i(k^*(\tilde{\theta}_i, \theta_{-i}), \tilde{\theta}_i) - \\ - \sum_{j \neq i} v_j(k^*(\theta_i, \theta_{-i}), \theta_j) - \underline{U}_i(\tilde{\theta}_i)$$

Generalized VCG

GVCG mechanism is a DRM with the efficient allocation $k^*(\theta)$ and payments

$$t_i^{GVCG}(\theta) \equiv \sum_{j \neq i} v_j(k^*(\tilde{\theta}_i, \theta_{-i}), \theta_j) + v_i(k^*(\tilde{\theta}_i, \theta_{-i}), \tilde{\theta}_i) - \\ - \sum_{j \neq i} v_j(k^*(\theta_i, \theta_{-i}), \theta_j) - \underline{U}_i(\tilde{\theta}_i)$$

Has the usual Groves' term (the third one); the other three guarantee IR.

- The first term is similar to Clarke's term, but with $k^*(\tilde{\theta}_i, \theta_{-i})$ instead of $k^{-i}(\theta_{-i})$
- 2^{nd} and 4^{th} is the net utility that LCT $\tilde{\theta}_i$ gets from participating in the mechanism – need to also pay it to all other types

Generalized VCG

Theorem (gVCG, part 1)

In a quasilinear model, gVCG is:

- *efficient (by construction),*
- *DSIC,*
- *interim IR.*
-

Prove DSIC on your own (analogous to VCG).

Generalized VCG. Proof: IIR

Interim expected utility for θ_i is

$$\mathbb{E}_{\theta_{-i}}\left[\sum_{j=0}^{N} v_j(k^*(\theta),\theta_j) - \sum_{j=0}^{N} v_j(k^*(\tilde{\theta}_i,\theta_{-i}),\theta_j)\bigg|_{\theta_i=\tilde{\theta}_i} |\theta_i\right] + \underline{U}_i(\tilde{\theta}_i)$$

Generalized VCG. Proof: IIR

Interim expected utility for θ_i is

$$\mathbb{E}_{\theta_{-i}}\left[\sum_{j=0}^{N} v_j(k^*(\theta),\theta_j) - \sum_{j=0}^{N} v_j(k^*(\tilde{\theta}_i,\theta_{-i}),\theta_j)\Big|_{\theta_i=\tilde{\theta}_i} |\theta_i\right] + \underline{U}_i(\tilde{\theta}_i) \geq \underline{U}_i(\theta_i)$$

The inequality above is the IIR constraint, and it holds since

$$\tilde{\theta}_i \in \arg\min_{\theta_i\in\Theta_i} \mathbb{E}_{\theta_{-i}}\left[\sum_{j=0}^{N} v_j(k^*(\theta_i,\theta_{-i}),\theta_j) - \underline{U}_i(\theta_i)|\theta_i\right]$$

Generalized VCG

Theorem (gVCG, part 2)

In a Euclidean model with independent players' types, gVCG is:

- *efficient (by construction),*
- *DSIC,*
- *interim IR;*
- *maximizes expected revenue among all mechanisms that are BIC, IIR, and implement the efficient k^*.*

If gCVG is not ex ante budget balanced, there does not exist a $\{$EFF + BIC + IIR + ex ante BB$\}$ mechanism (so no ex post BB either).

Generalized VCG. Proof: revenue maximizing in Euclidean

- Given revenue equivalence, just need to show we cannot decrease h_i for any player w/o violating IIR.

Generalized VCG. Proof: revenue maximizing in Euclidean

- Given revenue equivalence, just need to show we cannot decrease h_i for any player w/o violating IIR.
- Decreasing h_i only possible if IR slack for *all* types of i.

Generalized VCG. Proof: revenue maximizing in Euclidean

- Given revenue equivalence, just need to show we cannot decrease h_i for any player w/o violating IIR.
- Decreasing h_i only possible if IR slack for *all* types of i.
- But IR binds for $\tilde{\theta}_i$: $\bar{U}_i(\tilde{\theta}_i) = \underline{U}_i(\tilde{\theta}_i)$ (verify). $\square$

Application: Bilateral Trade

Example (Bilateral Trade (revisited))

- One indivisible good.
- Two agents: buyer and seller.
- Private valuations $\theta_b, \theta_s \sim$ i.i.d. $U[0, 1]$ resp.
- Is there an ~~efficient, DSIC, ex post IR, ex post BB~~ efficient, BIC, interim IR, ex ante BB mechanism?

- No, because gVCG is not BB. (This is the Myerson-Satterthwaite Theorem)

This slide deck:

AGV mechanism

- One last mechanism before we go – in case you care about BB, but not IR.
- Let

$$\tilde{t}_i(\theta_i) \equiv \mathbb{E}_{\theta_{-i}}\left[\sum_{j \neq i} v_j(k^*(\theta_i, \theta_{-i}), \theta_j) | \theta_i\right]$$

  be the "expected externality" imposed by i on everyone else.
- AGV transfers are given by

$$t_i^{AGV}(\theta) \equiv \frac{1}{N-1} \sum_{j \neq i} \tilde{t}_j(\theta_j) - \tilde{t}_i(\theta_i).$$

AGV mechanism

- One last mechanism before we go – in case you care about BB, but not IR.
- Let

$$\tilde{t}_i(\theta_i) \equiv \mathbb{E}_{\theta_{-i}} \left[\sum_{j \neq i} v_j(k^*(\theta_i, \theta_{-i}), \theta_j) | \theta_i \right]$$

  be the "expected externality" imposed by i on everyone else.
- AGV transfers are given by

$$t_i^{AGV}(\theta) \equiv \frac{1}{N-1} \sum_{j \neq i} \tilde{t}_j(\theta_j) - \tilde{t}_i(\theta_i).$$

- The second term is the averaged version of Groves' transfer,
- the first term is $h_i(\theta_{-i})$ which balances the budget.

AGV mechanism

Theorem (AGV)

In a quasilinear model, AGV is:

- *efficient (by construction),*
- *exactly ex post BB,*
- *BIC.*

Not necessarily IR. :(

AGV mechanism. Proof: budget balance.

- Observe that

$$\sum_i t_i^{AGV}(\theta) = \sum_i \left[\frac{1}{N-1} \sum_{j \neq i} \tilde{t}_j(\theta_j) - \tilde{t}_i(\theta_i) \right].$$

- For any j, RHS has:
  - $N-1$ terms of the form $\frac{1}{N-1}\tilde{t}_j(\theta_j)$, and
  - 1 term of the form $-\tilde{t}_j(\theta_j)$.
- These cancel out and exhaust all terms in the sum. Therefore, $\sum_i t_i^{AGV}(\theta) = 0$ for all $\theta =$ ex post exact budget balance.
  - **Note:** if the mechanism needs to raise some fixed sum *for any* θ, it can be treated as $\tilde{t}_0$. If the mechanism needs to raise some sum that is *dependent on* θ (e.g. fund a public project iff it is built), AGV cannot handle that.

AGV mechanism. Proof: BIC.

- If i reports $\hat{\theta}_i$ then receives utility

$$\mathbb{E}_{\theta_{-i}}\left[v_i(k^*(\hat{\theta}_i,\theta_{-i}),\theta_i)+\sum_{j\neq i}v_j(k^*(\hat{\theta}_i,\theta_{-i}),\theta_j)|\theta_i\right]-\frac{1}{N-1}\sum_{j\neq i}\tilde{t}_j(\theta_j)$$

- Last term indep of $\hat{\theta}_i$;

  bracket max-d by $\hat{\theta}_i=\theta_i$ for every θ_{-i} (since k^* efficient),

  so max-d by $\hat{\theta}_i=\theta_i$ in expectation as well.

- Reporting truth is a best response to $-i$ reporting truthfully

  $\Rightarrow$ truthful reporting is a BNE of the mechanism. $\square$

Bottom line

- We have two mechanisms that implement the efficient k^* in quasilinear model:
  - AGV: BIC + BB,
  - gVCG: DSIC + IIR.
- In the Euclidean model, gVCG is also [ex ante-]revenue-maximizing among BIC+IIR mechanisms.
- Revenue equivalence is powerful, but needs more structure than just quasilinear model.

References I

V. Krishna and E. Maenner. Convex potentials with an application to mechanism design. *Econometrica*, 69(4):1113–1119, 2001.

V. Krishna and M. Perry. Efficient Mechanism Design. *SSRN Electronic Journal*, 2000. ISSN 1556-5068. doi: 10.2139/ssrn.64934.