Mechanism Design

5: Beyond the Euclidean setting

Egor Starkov

Københavns Unversitet
Fall 2024

This slide deck:

1. Implementability: linear and quasilinear setting
2. Implementability: general setting
3. Example 1: Voting with single-peaked preferences
4. Example 2: Delegation
5. Example 3: Communication
6. Example 4: Cheap talk with correlated senders

Testing Implementability

The running question for now:

How can we check whether a given s.c.f. $f(\theta)$ is implementable?

- We have seen an answer for the Euclidean setting (k should be monotone).
- As nice as it is, Euclidean setting is very restrictive, mainly due to the $u(x, \theta) = \theta k - t_i$ payoff function.
- What about more general settings?

Euclidean setting: refresher

- **Reminder**: monotonicity for Euclidean problems.
  - Note: only the players' preferences are required to be linear for this to hold; the principal can have non-linear prefs.
- In a Euclidean setting, $k(\theta)$ is implementable only if it is monotone.
- Turns out, this is a sharp characterization:
  if $k(\theta)$ is monotone, there exist transfers t such that $\Gamma = (\Theta, (k, t))$ is DSIC.
  - Monotonicity may require $k(\theta)$ to be either weakly increasing, or weakly decreasing, depending on the problem.
  - To prove: use the relevant ERP to construct all transfers; can then show that the resulting mechanism is DSIC/BIC as needed.

Quasilinear setting: result 1

Extension to quasilinear setting:

- Suppose we are fine with one-dimensional k_i and θ_i, but want to consider more general payoff functions $v_i(k_i, \theta_i)$.

Theorem (Monotonicity in qlin setting)

In the quasilinear setting: if $K_i, \Theta_i \subset \mathbb{R}$, Θ_i is bounded, and $\frac{\partial^2 v_i(k_i,\theta_i)}{\partial k_i \partial \theta_i} > 0$, then:

if $k_i(\theta)$ is weakly increasing in θ_i for all i then k is DSIC.

- In words: if θ_i and k_i are complements from i's standpoint, then it is sufficient for k to be monotone to be implementable.
  - (Necessity probably also holds in this formulation, but I haven't checked)
- See Börgers ch.5.6 for details (including what boundedness of Θ_i means).

Quasilinear setting: result 2

What follows is a couple more results for the quasilinear setting with no restrictions on K or Θ.

Definition (weak q-monotonicity)

Allocation k is weakly q-monotone if for all $i, \theta'_i, \theta''_i, \theta_{-i}$:

$$v_i(k(\theta'_i, \theta_{-i}), \theta'_i) - v_i(k(\theta''_i, \theta_{-i}), \theta'_i) \geq v_i(k(\theta'_i, \theta_{-i}), \theta''_i) - v_i(k(\theta''_i, \theta_{-i}), \theta''_i)$$

Theorem (Necessity of weak monotonicity in qlin setting)

In the quasilinear setting: if k is DSIC then k is weakly q-monotone.

So k must be weakly q-monotone to be implementable.

But weak q-monotonicity does not guarantee implementability.

But we can strengthen this...

Quasilinear setting: result 3

Definition (cyclical q-monotonicity)

Allocation k is cyclically q-monotone if for all i, θ_{-i}, and all sequences $(\theta_i^1, \theta_i^2, \ldots, \theta_i^M) \in \Theta_i^M$ of arbitrary length M s.t. $\theta_i^M = \theta_i^1$, the following holds:

$$\sum_{m=1}^{M-1} \left[v_i(k(\theta_i^m, \theta_{-i}), \theta_i^{m+1}) - v_i(k(\theta_i^m, \theta_{-i}), \theta_i^m) \right] \leq 0$$

Theorem (Rochet [1987])

In a quasilinear setting: k is DSIC if and only if k is cyclically q-monotone.

Note: "Weak q-monotonicity" = "cyclical q-monotonicity for $M = 3$". See Börgers, ch.5.3-5.4 for proofs or references to proofs (for $N = 1$). See rest of ch.5 for other kinds of monotonicity for quasilinear setting.

This slide deck:

General setting: monotonicity

Without transfers, interesting results come up...

Definition (outcome g-monotonicity)

In a general setting, outcome x is g-monotone if for all $\theta', \theta'' \in \Theta$ the following holds:

- if for all i and all $x' \in X$ s.t. $u_i(x(\theta'), \theta') \geq u_i(x', \theta')$ it holds that $u_i(x(\theta'), \theta'') \geq u_i(x', \theta'')$,
- then $x(\theta'') = x(\theta')$.

In words, if under θ'' everyone likes $x(\theta')$ more than under θ', then we give $x(\theta')$ under θ''.

Theorem (Necessity of monotonicity in general setting)

In the general setting: if x is DSIC and $x(\Theta) = X$ then x is g-monotone.

(This is not THE interesting part yet. The next result is.)

General setting: dictatorship 1

Assumption (Domain)

Assume type sets Θ_i are rich enough to contain all possible (ordinal) preferences over X for all i.

Definition (dictatorial s.c.f.)

S.c.f. f is called dictatorial if there exists $i \in N$ s.t. for all type profiles θ:
$f(\theta) \in \arg\max_x u_i(x, \theta_i)$.

General setting: dictatorship 2

Theorem (Gibbard [1973], Satterthwaite [1975])

In a general setting with $|X| \geq 3$: if $x(\Theta) = X$ and the domain assumption holds, then

x is DSIC if and only if x is dictatorial.

- To clarify, $x(\Theta) \equiv \{x \in X \mid \exists \theta : x(\theta) = x\}$ is the set of "outcomes that could be prescribed for some $\theta \in \Theta$".
- Note: restriction $x(\Theta) = X$ is irrelevant for this result.
  If $x(\Theta) \subset X$, then only preferences over alternatives in $x(\Theta)$ are relevant, and we will still have a dictatorship on $x(\Theta)$. This is something we'll come back to later.
- GS theorem is the mechanism design version of Arrow's theorem from social choice.

General setting: dictatorship 3

- The missing link between the two results above is this:

Theorem (Monotonicity implies dictatorship)

In a general setting with $|X| \geq 3$: if $x(\Theta) = X$ and the domain assumption holds, then if x is g-monotone then x is dictatorial.

- For a proof of GS thm, see Narahari ch.17
- All three results for the general setting hold with ordinal preferences ($\succsim_i$) too, they do not rely on cardinal utilities u_i.
- The GS thm is also extendable to infinite X.

General setting: Takeaways

- It *seems* like GS thm is a strong negative result saying "we can't implement anything unless it's dictatorial!". But we obviously can: we've seen examples (like VCG). Where is the contradiction?

General setting: Takeaways

- It *seems* like GS thm is a strong negative result saying "we can't implement anything unless it's dictatorial!". But we obviously can: we've seen examples (like VCG). Where is the contradiction?
- The source of evil in GS thm is the domain assumption ("any preference is possible"). We often know something about some players' preferences, so can restrict the set of possible preferences.
  - E.g., quasilinear prefs: "everyone always likes money/transfers". Then at least the efficient allocation rule $k^*(\theta)$ is implementable (and typically not dictatorial!)
  - Another common example is the single-peaked preferences, explored below.

General setting: Takeaways

- It *seems* like GS thm is a strong negative result saying "we can't implement anything unless it's dictatorial!". But we obviously can: we've seen examples (like VCG). Where is the contradiction?
- The source of evil in GS thm is the domain assumption ("any preference is possible"). We often know something about some players' preferences, so can restrict the set of possible preferences.
  - E.g., quasilinear prefs: "everyone always likes money/transfers". Then at least the efficient allocation rule $k^*(\theta)$ is implementable (and typically not dictatorial!)
  - Another common example is the single-peaked preferences, explored below.
- Further, dictatorship is not a sentence!
  - With $N = 1$, all mechanisms are by definition dictatorial – yet they can still be useful!

Next steps

- We will now look at a few examples of mechanisms without transfers, dictatorial and not.
- These will show what kind of mechanisms we can have if we relax the domain assumption (similar to assuming we have access to transfers) and what kinds of instruments we can use.

This slide deck:

Example 1: Voting with single-peaked preferences

Setup:

- There are M alternatives that are ordered in some sense: $X = (x_1, x_2, ..., x_M)$ with $x_1 < x_2 < ... < x_M$.
- There are N players with private single-peaked ordinal preferences $\succsim_i (\theta_i)$ over $x \in X$
  - An ordinal preference relation $\succsim_i (\theta_i)$ is called single-peaked if $\exists x^*(\theta_i)$ s.t. for any $x_k < x_l \leq x^*(\theta_i)$, $x_l \succsim_i x_k$, and for any $x^*(\theta_i) \leq x_k < x_l$, $x_k \succsim_i x_l$.
  - Think "utility function $u_i(x, \theta_i)$ that is increasing between x_1 and $x^*(\theta_i)$ and decreasing between $x^*(\theta_i)$ and x_M".
- Example: policy debate on a one-dimensional issue – corporate tax rate, level of the unemployment benefits, openness of the immigration policy, level of governmental oversight over media/internet.

Question:

- Can we implement any non-dictatorial s.c.f. $f(\theta)$?

Single-peaked: Pairwise majority voting

If we allowed arbitrary preferences over X, then GS thm says "no, only dictatorship is incentive compatible". But we assumed single-peakedness, so GS thm does not hold:

Theorem

In the setting defined above, if the number of players N is odd, then pairwise majority voting selects the peak of the median voter.

- **Pairwise Majority Voting:** for any pair x_k, x_l, if the majority of voters prefers x_k to x_l (according to their reported types $(\theta_1, ..., \theta_N)$), then say that x_k is *socially preferred* to x_l. After comparing all pairs, select the one that is socially preferred to all others.
- Without single-peakedness, Condorcet cycles $(x_1 \succ_S x_2 \succ_S x_3 \succ_S x_1)$ may arise in the social preference from PMV. But if individual prefs are single-peaked, soc pref is well defined, its most preferred alternative exists, and coincides with the median voter's most preferred alternative. (See MWG, theorems 21.D.1-2.)

Single-peaked: Takeaways

- We can often make some assumptions on players' prefs, which make some non-dictatorial s.c.f.s implementable.
- In the problem of social choice on a one-dimensional issue, the median voter's ideal option is preferred by the majority.

This slide deck:

Example 2: Delegation

Setup: (following Holmstrom [1980])

- Consider a principal-agent model with one principal/designer and one agent, with the respective payoffs given by quadratic loss functions:

$$u_P(x,\theta) = -\alpha(x-\theta)^2$$
$$u_A(x,\theta) = -\alpha(x-\theta-b)^2,$$

  where $\theta \sim U[0,1]$ is the state of the world only known by the agent; $\alpha \geq 0$ and $b \in [0,1/2]$ are commonly known parameters; and $x \in \mathbb{R}$ is the decision to be made.

- The principal can choose x directly or let the agent choose $x \in X$ while restricting the set of options $X \subseteq \mathbb{R}$ available to the agent.

Question: how should the principal act?

Delegation: Solution 1

- Note that this can be seen as a mechanism design problem, in which the agent reports θ, and the principal commits to some decision rule $x(\theta)$.
  The agent then effectively chooses x from the set $x(\Theta)$.

Solution:

- The agent's FOC for x given θ is

$$(-2\alpha)(x - \theta - b) = 0$$

  So the agent selects $x^*(\theta) = \theta + b$ if it is available; $x \in X$ closest to $x^*(\theta)$ otherwise.

- The principal would ideally prefer $x_P(\theta) = \theta$.

Delegation: Solution 2

Lemma

The optimal delegation set X is convex (i.e., an interval).

Proof: suppose not, and the principal's chosen X has $x_1, x_2 \in X$ and $(x_1, x_2) \notin X$. (Formally, there are more cases to consider, but we ignore those.) Consider an alternative delegation set $X' \equiv X \cup (x_1, x_2)$.

X and X' yield the same payoff to the principal when $\theta \notin (x_1 - b, x_2 - b)$, since then all $x \in (x_1, x_2)$ are dominated for the agent by either x_1, or x_2. I.e., the agent's choice is the same from X and X' for such θ.

Suppose $\theta \in (x_1 - b, x_2 - b)$. Under X', the agent plays $x^*(\theta)$ for all such θ, hence $x(\theta) - \theta = b$ for all such θ. Under X, the agent plays x_1 or x_2; can show that $\mathbb{E}[x(\theta) - \theta \mid \theta \in (x_1 - b, x_2 - b)] = b$ in that case as well.

The principal's payoff is a concave function of $x - \theta$, hence (by Jensen's inequality) the principal prefers constant $x - \theta$ to a lottery with the same mean – hence X' is better than X! $\square$

Delegation: Solution 3

So the optimal delegation set is $X = [\underline{x}, \bar{x}]$. Agent wants to take higher actions than the principal $\Rightarrow$ no sense restricting them from the bottom $\Rightarrow$ $\underline{x} = 0$ (or $\underline{x} = -\infty$). Actions $x \geq 1$ never optimal for the principal, hence the upper limit must be $\bar{x} \leq 1$. Principal's expected payoff is

$$\mathbb{E}\left[-\alpha\left(x(\theta) - \theta\right)^2\right] = -\alpha\left[\int_0^{\bar{x}-b} \left(x^*(\theta) - \theta\right)^2 d\theta + \int_{\bar{x}-b}^{1} (\bar{x} - \theta)^2 d\theta\right].$$

Maximizing over $\bar{x}$ yields the optimal upper limit $\bar{x} = 1 - b$.

Final answer: the optimal delegation set is $X = [0, 1 - b]$.
Or, in mechanism design terms, the optimal direct mechanism is $x(\theta) = \min\{\theta + b, 1 - b\}$.

Delegation: Takeaways

Takeaways regarding the problem:

- Delegation is a prominent problem in organizational econ.
- Makes sense to restrict the agent where there's conflict (high x), not where there's none (low x).

Broader takeaways:

- Even dictatorial mechanisms can be useful for the principal (again, have $N = 1$ in this problem, so the mechanism is by definition dictatorial).
- An example of the applicability of mechanism design (not immediate from the start how this is a mechdesign problem).
- An example on the usefulness of indirect mechanisms IRL (note that the actual delegation does not require the principal to commit to a decision rule).

This slide deck:

Example 3: Communication

- Another prominent question in org.econ is **“delegation vs communication”**: is it better to ask the agent to report the state to the principal, or should we just let the agent make the decision?
- While it’s not necessarily in the mechanism design realm, this relates to our broad question of “how to best extract private information from the players?”

Cheap talk: setup

Setup: (following Crawford and Sobel [1982])

- Same as before: one principal/designer, one agent, with payoffs given by quadratic loss functions:

$$u_P(x,\theta) = -\alpha(x-\theta)^2$$
$$u_A(x,\theta) = -\alpha(x-\theta-b)^2,$$

  where $\theta \sim U[0,1]$ is the state of the world privately known by the agent; $\alpha \geq 0$ and $b \in [0, 1/4]$ are commonly known parameters; and $x \in \mathbb{R}$ is the decision to be made.
- **NEW:** The principal asks the agent to report θ, then (principal) chooses x given the report to maximize their payoff, cannot commit to a decision rule.

Question: how much can the principal learn about the state? (In MD lingo: which s.c.f. $x(\theta)$ are implementable?)

Cheap talk: comments

- Due to no-commitment assumption, the principal will always choose $x(m) = \mathbb{E}[\theta|m]$ after message m.
    - Can see this as a MD problem with two players: the agent and the future principal – but it's a stretch.
    - More reasonable interpretation – MD problem (with 1 agent) with an additional constraint (principal's ex post compliance).
- This model is known as a game of *cheap talk* communication, since the agent can message about θ but has no evidence to back up their claim – in contrast to models of *verifiable disclosure*.

Cheap talk: solution 1

- Again, de facto dictatorship: principal offers a set of options X to the agent.
- Again, need to figure what the agent's IC conditions (together with the principal's ex post IC) imply for how $x(\theta)$ can and cannot look like.
  The principal's IC implies that for all θ, $x(\theta) = \mathbb{E}[\theta' \mid x(\theta') = x(\theta)]$.

Cheap talk: solution 1

- Again, de facto dictatorship: principal offers a set of options X to the agent.
- Again, need to figure what the agent's IC conditions (together with the principal's ex post IC) imply for how $x(\theta)$ can and cannot look like.
  The principal's IC implies that for all θ, $x(\theta) = \mathbb{E}[\theta' \mid x(\theta') = x(\theta)]$.
- **Statement:** all implementable X are finite: $X = \{x_1, ..., x_K\}$, with options at least b apart.
  Proof: suppose $x', x'' \in X$ with $x'' - x' \in (0, b)$. Then for all $\theta \geq x'$, the agent prefers x'' to x', so $x(\theta) \neq x'$ for such θ. But then $\mathbb{E}[\theta \mid x(\theta) = x'] < x'$, a contradiction with the principal's IC.

Cheap talk: solution 1

- Again, de facto dictatorship: principal offers a set of options X to the agent.
- Again, need to figure what the agent's IC conditions (together with the principal's ex post IC) imply for how $x(\theta)$ can and cannot look like.
  The principal's IC implies that for all θ, $x(\theta) = \mathbb{E}[\theta' \mid x(\theta') = x(\theta)]$.
- **Statement:** all implementable X are finite: $X = \{x_1, \dots, x_K\}$, with options at least b apart.
  Proof: suppose $x', x'' \in X$ with $x'' - x' \in (0, b)$. Then for all $\theta \geq x'$, the agent prefers x'' to x', so $x(\theta) \neq x'$ for such θ. But then $\mathbb{E}[\theta \mid x(\theta) = x'] < x'$, a contradiction with the principal's IC.
  - The exact opposite of delegation!
  - This is fully due to the principal's IC. What constraints does the agent's IC add?

Cheap talk: solution 2

- It is immediate that $x(\theta)$ must be increasing (by the standard monotonicity argument – just look at the pair of the agent's mutual IC conditions for some θ', θ'').
- So any option $x \in X$ is chosen on an interval of states (or nowhere): $\{\theta \mid x(\theta) = x_k\} = [\theta_{k-1}, \theta_k]$, where.
- Can we figure out the interval boundaries? Yes, easy: when $\theta = \theta_k$, the agent must be indifferent between x_k and x_{k+1} (so $\theta_k + b \in [x_k, x_{k+1}]$):

$$-\alpha(x_k - \theta_k - b)^2 = -\alpha(x_{k+1} - \theta_k - b)^2 \iff \theta_k + b - x_k = x_{k+1} - \theta_k - b.$$

- Further, from the principal's IC and the uniform distribution:

$$x_k = \mathbb{E}[\theta \mid x(\theta) = x_k] = \frac{\theta_{k-1} + \theta_k}{2}$$

Cheap talk: solution 3

- Combining the two conditions yields a recurrent equation on $\{\theta_k\}$:

$$\theta_{k-1} - 2\theta_k + \theta_{k+1} = 4b.$$

- **Final answer:** any sequence $\{\theta_k\}_{k=0}^{K}$ that satisfies the equation above and the boundary conditions $\theta_0 = 0$, $\theta_K = 1$ characterizes an implementable s.c.f.
- Can show: there is an upper bound $\bar{K}(b)$ such that for any $K \in \{1, ..., \bar{K}(b)\}$, there exists an implementable s.c.f. with K options in the choice set.
  - Alternative (original) interpretation: for every $K \in \{1, ..., \bar{K}(b)\}$, there exists a corresponding equilibrium of the communication game.
  - Reminder: in MD, we assume the designer has the power of equilibrium selection!
- Both the principal and the agent prefer the more informative/responsive s.c.f. (higher K)

Cheap talk: results and discussion

1. Can calculate the principal's payoffs and show that delegation is better than communication for the principal
   - Not immediate from the setup of the two respective games.
   - Counterintuitive that the manager gains from having fewer decision rights under delegation!
   - Obvious when framed as mechanism design problems: "communication" is "delegation + extra constraint".
     Lesson: framing problems in MD context can provide clarity!
2. We can deal with extra constraints (like the principal's ex post IC), no problem! The whole "MD problem" is exactly about figuring out what the constraints are and what they imply!

This slide deck:

Example 4: Cheap talk with correlated senders

- This is a throwback to the previous topic:
  - Cremer-McLean: "if players' info is correlated and have access to transfers, can make players bet on each other's info to extract all private info at no cost"
- That relied on transfers.
- Now: show that if players' info is correlated and preferences not the same, there's another channel: setting players against each other
- Can extract all info even without transfers! Or commitment!

Multi-sender cheap talk: Setup

Setup: (based on Battaglini [2002])

- Two agents, $i \in \{1, 2\}$:
  - *both* know state $\omega = (\omega^1, \omega^2) \in \mathbb{R}^2$;
  - each sends report $m \in \mathbb{R}^2$ to the principal.
- Principal ("designer") does not know ω, must choose action $a \in \mathbb{R}^2$ after hearing (m_1, m_2).
- Preferences: squared Euclidean distance between a and resp. bliss points
  - Principal: $u_p(a, \omega) = -(\|a - \omega\|)^2$;
  - Agent i: $u_i(a, \omega) = -(\|a - (\omega + b_i)\|)^2$;
  - where $\|x\| \equiv \sqrt{(x^1)^2 + (x^2)^2}$ for $x = (x^1, x^2) \in \mathbb{R}^2$.
  - Biases b_i commonly known.
- (Subscripts index i, superscripts stand for coordinates [in default basis] and exponents.)

Multi-sender cheap talk: Comments

- Same basic model as cheap talk of Crawford and Sobel [1982], but now:
  - two agents
  - two dimensions for state and actions
  - notation is different because I didn't have time to fix it. Actions/options are now a (used to be x).
- Cross-verification mechanism may not work because:
  1. there's no punishment that's universally bad for all ω
  2. the principal cannot commit to punishment!

Multi-sender cheap talk: Idea

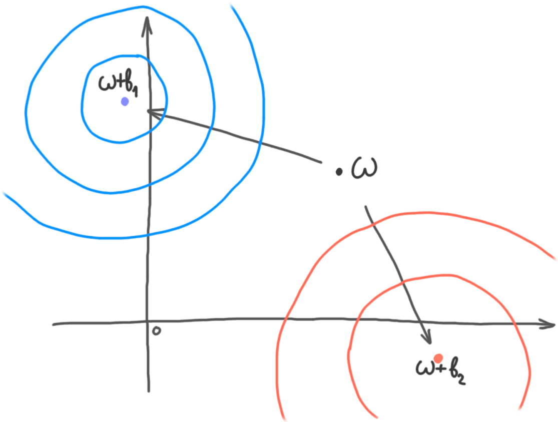

- Relative positions of bliss points and indifference curves are fixed; just the absolute location unknown.
- The circles on the graph represent the indifference curves of the two players.

Multi-sender cheap talk: Idea

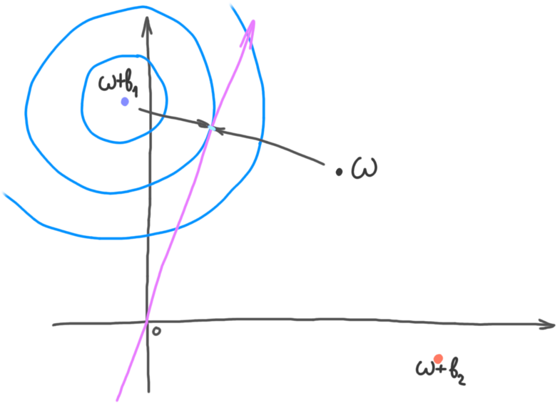

- Ask player i to project the state on some axis orthogonal to b_i and report the result.
- Will report honestly (i.e., report, which is a projection of $\omega + b_1$, coincides with the projection of the actual ω).
- So we learn one coordinate of the true state.

Multi-sender cheap talk: Idea

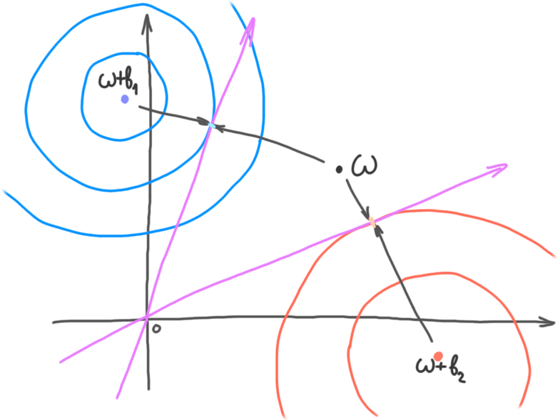

- Then with two players, we can learn state perfectly this way.
- (As long as b_i linearly independent.)
- More generally, two players are enough to learn the state of any dimensionality n, since asking either player allows to learn $n-1$ dimensions of state.
- See Battaglini (2002) for n dimensions and more general preferences.

Multi-sender cheap talk: Equilibrium strategies

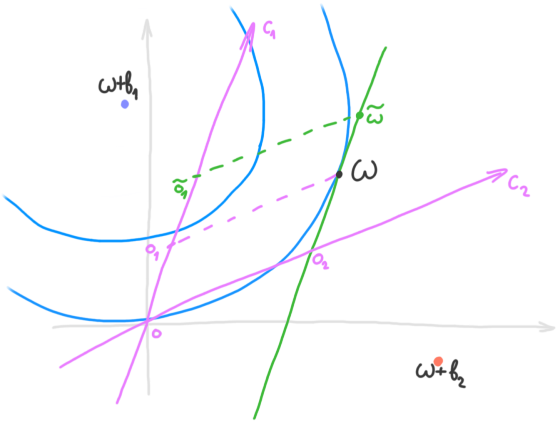

- Consider basis (c_1, c_2) where $b_i \perp c_i \in \mathbb{R}^2$ – i.e., $b_i \cdot c_i = 0$ ($\iff$ $b_i^1 c_i^1 + b_i^2 c_i^2 = 0$).
- State ω has unique coordinates (o_1, o_2) in this basis, i.e. $\omega = o_1 \cdot c_1 + o_2 \cdot c_2$.
- Ask A1 to report o_1. If A2 reports o_2 truthfully, A1 effectively chooses an action on the green line (see graph).
- So truthful reporting is optimal for A1 (green line is orthog to $b_1 \Rightarrow$ it is tangent to A1's circular indifference curve at $\omega \Rightarrow$ any lies $\tilde{o}_1$ puts the implemented action $\tilde{\omega}$ on a lower indiff curve).

Multi-sender cheap talk: Pesky details

- There are many vectors c_i that are orthogonal to a given b_i (due to scaling) – select any.
  - E.g., letting $c_i = \frac{1}{\|b_i\|_2}\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}\begin{pmatrix} b_i^1 \\ b_i^2 \end{pmatrix}$ would yield a vector c_i of unit length that is rotated 90 deg clockwise w.r.t. b_i.
- Formally, the problem setup says that messages are two-dimensional: $m_i \in \mathbb{R}^2$. So to be 100% formal you can say that, e.g., $m_i = (o^i, 0)$, and that the principal ignores the second coordinate of each message.
  - Alternatively: suppose the agents report ω, but then the principal calculates respective o_i and makes the final choice based on them (even (especially!) if reports do not coincide).

Multi-sender cheap talk: Takeaways

Correlated information is very easy to extract – even when you have no access to transfers and have commitment issues.

References I

M. Battaglini. Multiple Referrals and Multidimensional Cheap Talk. *Econometrica*, 70(4): 1379–1401, July 2002. ISSN 1468-0262. doi: 10.1111/1468-0262.00336.

V. P. Crawford and J. Sobel. Strategic Information Transmission. *Econometrica*, 50(6): 1431–1451, 1982.

A. Gibbard. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society*, pages 587–601, 1973.

B. Holmstrom. On the theory of delegation. Discussion Papers 438, 1980. URL `https://EconPapers.repec.org/RePEc:nwu:cmsems:438`.

J.-C. Rochet. A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of mathematical Economics*, 16(2):191–200, 1987.

References II

M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory*, 10 (2):187–217, 1975.