Mechanism Deisgn with Multiple Informed Parties

Dr. Margaret Meyer

2021

Mechanism Design with multiple informed parties

- Optimal Contracting / Mechanism Design can be extended to analyze settings with more than one privately-informed party
- *Applications:* auctions; bilateral exchange; public good provision; mediation of disputes
- Suppose each agent i privately observes his own type s_i before participating in the game
- Mechanism designer **commits to** a direct revelation mechanism (DRM):

$$\{x_1(\hat{s}), \ldots, x_n(\hat{s}); t_1(\hat{s}), \ldots, t_n(\hat{s})\},$$

where $(x_i(\hat{s}), t_i(\hat{s}))$ is the (allocation, payment) pair for agent i, given vector $\hat{s} = (\hat{s}_1, \ldots, \hat{s}_n)$ of reported types.

Mechanism Design with multiple informed parties

Mechanism designer **commits to** a direct revelation mechanism (DRM):

$$\{x_1(\hat{s}), \ldots, x_n(\hat{s}); t_1(\hat{s}), \ldots, t_n(\hat{s})\},$$

where $(x_i(\hat{s}), t_i(\hat{s}))$ is the (allocation, payment) pair for agent i, given vector $\hat{s} = (\hat{s}_1, \ldots, \hat{s}_n)$ of reported types.

- Since $(x_i(\hat{s}), t_i(\hat{s}))$ depends on the whole vector $\hat{s}$, a DRM induces a game among agents, who choose reporting strategies $\hat{s}_i = S_i(s_i)$.
- Hence we need to specify an eqm concept, e.g. Bayesian Nash eqm or dominant strategy eqm.
- Given an eqm concept, the Revelation Principle generalizes to multi-agent settings: Any outcome achievable via an eqm of an arbitrary DRM can be achieved with a DRM for which truthful reporting is an eqm. Hence, can confine attention to truthful DRM's.
- Same qualitative conclusion as from single-agent analysis: if agents have hidden information when they accept/reject mechanism, then even though self-selection can typically be achieved, typically there are efficiency losses relative to full-information benchmark.

Multiple informed parties: Application

Mechanism to choose provider of a public good/service

- Based on Goldluecke and Troeger (2020), "The Multiple-Volunteers Principle", CEPR Disc. Pap. 15580.
- Exactly one of n risk-neutral individuals (agents) must be chosen
- Each can provide the public good/service at personal cost c
- Agents differ in the **quality** they provide, which **benefits everyone**: agent i's quality is $\theta_i \in \{L, H\}$
- If the provider is agent i, then i's payoff is $\theta_i - c$ and every other agent j's payoff is θ_i
- Each agent i **privately** observes his own θ_i (his "type")

Full-info benchmark: Choose a high-quality provider, as long as there is at least one such agent.

How should an uninformed mechanism designer P design a mechanism to maximize prob. of choosing a high-quality provider?

Mechanism to choose the provider of a public good/service

How should an uninformed mechanism designer P design a mechanism to maximize prob. of choosing a high-quality provider?

- Assume the mechanism cannot specify monetary transfers
- Ignore participation constraints

Consider the following naive task-assignment rule (mechanism):

- Agents are asked simultaneously whether or not they "volunteer".
- If there is at least one volunteer, task is assigned randomly *among the volunteers*.
- If no one volunteers, task is assigned randomly *among all agents*.
- This rule leads to poor outcomes: As long as c is not small, no one will volunteer, so task assignment will be completely random.

Mechanism to choose the provider of a public good/service

How can P encourage H-types to volunteer, in the absence of monetary transfers?

- *Answer:* By increasing the risk for H-types that, if they don't volunteer, an L-type will be (inefficiently) chosen as provider.
- NB: Since provider's quality affects everyone's payoff, self-selection is achievable because different types have different tradeoffs btw. volunteering and not volunteering.

What mechanism maximizes prob. of choosing a high-quality provider?

- Under what conditions does best mechanism allow P to do as well as under full info?
- If P cannot do as well as under full info, how costly in terms of efficiency is hidden information?

This model is, however, not yet fully specified...

Mechanism to choose the provider of a public good/service

This model is not yet fully specified:

1. How are $(\theta_1, \ldots, \theta_n)$ determined? What do players know about others' types?
   - Goldluecke and Troeger (2020) assumed θ_i's i.i.d.
   - We'll consider the much simpler case where θ_i's are identically distributed but *"negatively correlated"*, such that with prob. 1, there are exactly $\hat{k}$ H-types and $n - \hat{k}$ L-types, where $\hat{k}$ is between 1 and $n - 1$.

2. What equilibrium concept are we using in the reporting game among agents?
   - We will use Bayesian Nash eqm: We'll say a DRM is truthful if truthful reporting is a Bayesian Nash eqm for the agents.
   - Goldluecke and Troeger use Bayesian Nash eqm as well.

Mechanism to choose the provider of a public good/service

Given the Revelation Principle, and symmetry across agents, can focus on DRM's of the following form:

- Given agents' reports $(\hat{\theta}_1, \ldots, \hat{\theta}_n)$, let $k(\hat{\theta}_1, \ldots, \hat{\theta}_n)$ denote the number of H reports.
- For each $k \in \{1, \ldots, n-1\}$, the DRM specifies a probability p_k and **commits** P to behave as follows:
  - with prob. p_k, P randomly chooses as provider one of the k agents who reported H
  - with prob. $1 - p_k$, P randomly choosees as provider one of the $n - k$ agents who reported L
  - necessarily, $p_n = 1$ and $p_0 = 0$

Mechanism to choose the provider of a public good/service

Let $k(\hat{\theta}_1, \ldots, \hat{\theta}_n)$ denote the number of H reports.

DRM specifies probs. $\{p_1, \ldots, p_{n-1}\}$ and **commits** P to

- with prob. p_k, randomly choose as provider one of the k agents who reported H
- with prob. $1 - p_k$, randomly choose as provider one of the $n - k$ agents who reported L

Examples of such DRM's:

- $p_k = \frac{k}{n}$ for each k: This mechanism chooses the provider *completely randomly*— Regardless of reports, each agent is equally likely to be chosen.
- The naive mechanism which *always* chooses one of the volunteers (H-reporters) randomly, unless there are none, corresponds to $p_k = 1$ for all $k \geq 1$ and $p_0 = 0$.

Mechanism to choose the provider of a public good/service

Given our assumption that there are certain to be exactly $\hat{k}$ H-types, $k(\hat{\theta}_1, \ldots, \hat{\theta}_n) = \hat{k}$ if all agents report truthfully.

Hence, if we can find a DRM with $p_{\hat{k}} = 1$ that induces truthful reporting as a Bayesian Nash eqm, then efficiency is achievable, even with hidden information.

- Consider an L-type: If $p_{\hat{k}} = 1$ and $p_{\hat{k}+1} = 1$, then an L-type prefers to report L rather than H, if he expects all other agents to be truthful.
- Consider an H-type: Setting $p_{\hat{k}-1} = 0$ makes falsely reporting L *as unattractive as possible* for an H-type if all other agents are expected to be truthful, since H's lie then causes an L-type to be chosen (unless he himself is chosen).
- Thus, it remains to determine: If $p_{\hat{k}} = 1$ and $p_{\hat{k}-1} = 0$, is an H-type willing to report truthfully if he expects all other agents to be truthful?

Mechanism to choose the provider of a public good/service

If $p_{\hat{k}} = 1$ and $p_{\hat{k}-1} = 0$, is an H-type willing to report truthfully if he expects all other agents to be truthful?

- Expected payoff for an H-type who reports H is $H - \frac{1}{\hat{k}}c$.
- Expected payoff for an H-type who reports L is

$$\left(\frac{1}{n-\hat{k}+1}\right)(H-c) + \left(\frac{n-\hat{k}}{n-\hat{k}+1}\right)L.$$

- The ICC for an H-type can be rearranged to

$$\left(\frac{H-L}{c}\right) \geq \frac{n+1-2\hat{k}}{\hat{k}(n-\hat{k})} \qquad (ICC_H)$$

- *Thus, if this inequality is satisfied, then a DRM with $p_{\hat{k}} = p_{\hat{k}+1} = 1$ and $p_{\hat{k}-1} = 0$ yields efficient provision, even under hidden info.*

Mechanism to choose the provider of a public good/service

We've shown: Efficient provision of the public good is achievable (with $p_{\hat{k}} = p_{\hat{k}+1} = 1$ and $p_{\hat{k}-1} = 0$) if (ICC_H) holds:

$$\left(\frac{H-L}{c}\right) \geq \frac{n+1-2\hat{k}}{\hat{k}(n-\hat{k})} \qquad (ICC_H)$$

- (ICC_H) is more likely to be satisfied, the larger is $H-L$ relative to c; the smaller is n; and the larger is $\hat{k}$.
- Holding $\frac{\hat{k}}{n}$ fixed and increasing n, (ICC_H) will eventually be satisfied for n sufficiently large.
- Provide the intuition for these comparative statics results....

If (ICC_H) is violated, then there is no truthful DRM with $p_{\hat{k}} = 1$, so by the Revelation Principle, there is no DRM yielding efficient provision under hidden info.

Mechanism to choose the provider of a public good/service

When (ICC_H) is violated, what is the optimal mechanism?

By the Revelation Principle, a mechanism is optimal if it induces truthful reporting as a Bayesian Nash eqm and maximizes the probability of choosing an H-type provider.

Claim: *Suppose that, when $p_{\hat{k}} = 1$ and $p_{\hat{k}-1} = 0$, an H-type is not willing to report truthfully if he expects all other agents to be truthful. Then the* **optimal** *mechanism sets $p_{\hat{k}-1} = 0$, $p_{\hat{k}+1} = 1$, and $p_{\hat{k}}$ as large as possible subject to satisfying an H-type's truthtelling constraint. The optimal value of $p_{\hat{k}}$ is*

$$p_{\hat{k}} = \left(\frac{c - (H - L)}{c - \hat{k}(H - L)} \right) \left(\frac{\hat{k}}{n - \hat{k} + 1} \right).$$

Mechanism to choose the provider of a public good/service

Claim: *Suppose that, when* $p_{\hat{k}} = 1$ *and* $p_{\hat{k}-1} = 0$, *an H-type is not willing to report truthfully if he expects all other agents to be truthful. Then the* **optimal** *mechanism sets* $p_{\hat{k}-1} = 0$, $p_{\hat{k}+1} = 1$, *and* $p_{\hat{k}}$ *as large as possible subject to satisfying an H-type's truthtelling constraint.*

Proof:

1. Solve the "relaxed" problem of maximizing $p_{\hat{k}}$ subject to the H-type's ICC, ignoring the L-type's ICC.
2. The argument for setting $p_{\hat{k}-1} = 0$ is the same as above: the expected payoff of an H-type who reports L is minimized when $p_{\hat{k}-1} = 0$, so this makes the H-types's ICC as easy as possible to satisfy.
3. In the relaxed problem, it is now optimal to set $p_{\hat{k}}$ as large as possible subject to satisfying the H-type's ICC.
4. Under the claim's hypothesis, the expected payoff of an L-type who reports H is minimized when $p_{\hat{k}+1} = 1$, so this is the optimal $p_{\hat{k}+1}$.
5. It can be checked that with $p_{\hat{k}-1}$, $p_{\hat{k}}$, and $p_{\hat{k}+1}$ specified as described, the L-type's ICC is satisfied.