Robust Mechanism Design

Lecture 1: Introduction to Robust Mechanism Design

Lecturer: Zeyuan Allen Zhu *Scribe: Zeyuan Allen Zhu*

0 Background

This (series of) lecture notes intend to introduce the new concept of robust mechanism design recently proposed by a group of researchers in CSAIL. Unless otherwise noted, definitions and results all come from their published manuscripts, and due to the lack of time, large pieces of the original paragraphs are directly adopted here.

1 Introduction

Combinatorial Auction is a kind of auction in which bidders are allowed to bid for a union of goods, instead of only bidding for single good in traditional auction. Maximizing revenue or/and social welfare is one of the goals in designing auction mechanisms. Traditionally based on equilibria, mechanisms are found to be vulnerable to:

- **Collusion**. VCG is vulnerable when 2 players collude.
- **Complexity**. Some mechanisms require exponential number of rounds or communications.
- **Privacy**. Players publicly report their own utility functions
- **Equilibrium Selection**. There may be several Nash equilibria, while the property holds for just some of them.

In an unpublished manuscript of Micali and Valiant [1], they advocate overcoming the above weakness by designing mechanisms in a *resilient* way. This idea has then been formalized and well-discussed in a paper [2] submitted to STOC'10. Notice that several attempts have been made upon this new definition and they succeeded in designing a few mechanisms robust against collusion, complexity, privacy and equilibrium selection [3] [4] [5].

This new solution concept is not equilibrium-based. Only occasionally, the new concept coincides with traditional ones:

- **Strictly dominated strategy**. A special case when only a single strategy survives after the elimination.
- **The unique subgame-perfect equilibrium**. Another special case in games of *extensive form*.

REMARK 1:

- Games of *normal form*: players act simultaneously and the mechanism decides the outcome.
- Games of *extensive form*: a tree-structured mechanism and the player chooses among possible moves at each node.

Technically, the new concept is based on an elementary notion, *iterative elimination of distinguishably dominated strategies*. This new notion bridges a currently vast chasm: that between elimination of strictly dominated strategies and the elimination of weakly dominated strategies.

REMARK 2:

- B strictly dominates A: choosing B always gives a better outcome than choosing A, no matter what the other player(s) do.
- B weakly dominates A: There is at least one set of opponents' action for which B is superior, and all other sets of opponents' actions give B at least the same payoff as A.

1.1 Equilibrium Selection

As noticed in [2], elimination of weakly dominated strategies is unlike to be sufficiently meaningful. Therefore, the problem of equilibrium selection vanishes when the mechanism yields a (single) equilibrium:

- With strictly dominant-strategy. Such mechanism is rare and proved to be unable to guarantee even constant fraction of the property. [4] [6]
- With survived strategy through iterated elimination of strictly dominated strategies. Shown by Abreu and Matsushima [7], and further extended by Glazer and Perry [8], mechanisms of this kind are capable of achieving all desired properties, but under strong assumptions: *each player perfectly knows the utility functions of all players*.

Notice that in the second kind mentioned above, the strictly dominated strategy is not yet defined when players do not have the perfect knowledge of other players' utilities. To overcome such obstacle, they consider the iterated elimination of *distinguishably dominated strategies*.

2 Preliminaries

Combinatorial Auctions. In a combinatorial auction with n players and multiple goods for sale,

- The *true valuation* of a player i for a subset S of goods: $TV_i(S)$.
- An *allocation* A consists of a partition of goods $A = A_0, A_1, \ldots A_n$, where A_0 represents the unallocated goods
- An *outcome* Ω consists of an allocation A and a price profile P, where P_i is the payment by i.
- We consider *unrestricted* auction: $TV_i(S)$ is independent of $TV_j(S')$ for any $(i,S) \neq (j,S')$.
- The *utility* of a player i: $u_i(A,P) = TV_i(A_i) - P_i$.

Extensive-Form Public-Action Mechanisms. They mechanism must specify the *decision nodes*, the players' *acting* at each node, the set of actions available to each player at each node, and the auction outcome associated to each *terminal node* – leaf of the game tree. The mechanism is of *public action*.

A player i's *strategy* specifies i's action at each decision node, and let σ be a *play* consists of a profile of strategies. $H(\sigma)$ denotes the history of the play, and $M(\sigma)$ denotes the auction outcome (A,P) associated to $H(\sigma)$. Besides, a mechanism must provide an *opt-out* strategy OUT_i satisfying: $u_i(M(\mathrm{OUT_i} \sqcup \sigma_{-i}))$, for each strategy subprofile σ_{-i}. All definitions above may be probabilistic.

Generalized Contexts and Auctions. A game $G = (\mathcal{C}, M)$ has two components: a context C and a mechanism M. The context describes the players, possible outcomes, players' utilities for each outcome, and the players' knowledge. The mechanism describes which strategies are available to the players and how each profile yields an outcome.

DEFINITION 1: *A* (generalized) (auction) context $\mathcal{C} = (TV, (\mathbb{C}, I), EK)$ *consists of three components:*

- *The* true-valuation profile TV.

- *The* collusion structure $(\mathbb{C}, I)$*, where* $\mathbb{C}$ *is a partition of (collusive) players, and* I *is the set of players* i *s.t.* $\{i\} \in \mathbb{C}$*. We use agent to denote either an independent player or a collusive set and call a player in* I *as independent.*
- *In contrast to the* internal knowledge *of* A *which is* TV_A*, we define the* external-knowledge *vector* EK*: for each agent* $A \in C$*,* EK_A *is a set of valuation subprofiles for the players outside* A*, satisfying* $TV_{-A} \in EK_A$*.*

Now we define the relevant knowledge of an agent. Essentially this is the outcome with maximum welfare known to its members.

DEFINITION 2 ($\mathbb{MKW}$ **Benchmark**): *Given a context* $\mathcal{C}$ *and an agent* $\mathcal{A}$*, we define* $RK_{\mathcal{A}}$ *the* (total) relevant knowledge *of* $\mathcal{A}$*, to be the outcome with maximum revene among all outcomes* (A, P) *such that, for all player* j*:*

- *If* $j \in \mathcal{A}$*, then* $P_j = TV_j(A_j)$
- *If* $j \notin \mathcal{A}$*, then* $V_j(A_j) \geq P_j$ *for all* $V \in EK_{\mathcal{A}}$

The maximum known welfare *of* $\mathcal{A}$*,* $\mathbb{MKW}_{\mathcal{A}}$*, is the revenue of* $RK_{\mathcal{A}}$*. The maximum known welfare of* $\mathcal{C}$ *is denoted by* $\mathbb{MKW} = \max_{\mathcal{A} \in \mathbb{C}} \mathbb{MKW}_{\mathcal{A}}$*.*

The above two definitions come from [5]. There are actually two more types of benchmarks other than $\mathbb{MKW}$ defined in [5] and [3]. To avoid the ambiguity, here I define them in slightly different ways. For comes the $\mathbb{MKW}^{I}$, which denotes the maximal known welfare upon independent players.

DEFINITION 3 ($\mathbb{MKW}^{I}$ **Benchmark**): *Given a context* $\mathcal{C}$ *and the independent player set* I*:*

$$\mathbb{MKW}^{I} = \max_{i \in I} \mathbb{MKW}_{\{i\}}$$

Next comes the $\mathbb{MEW}$, which is defined over the external knowledge of independent players, and thus more restricted than $\mathbb{MKW}^{I}$.

DEFINITION 4 ($\mathbb{MEW}$ **Benchmark**): *Given a context* $\mathcal{C}$ *and an independent player* i*, we define* RK_i' *the* (external) relevant knowledge *of* i*, to be the outcome with maximum revenue among all outcomes* (A, P) *such that, for all player j:*

- *If* $j = i$*, then* $P_i = 0, A_i = \emptyset$*. (External Sale Only)*
- *If* $j \neq i$*, then* $V_j(A_j) \geq P_j$ *for all* $V \in EK_{\{i\}}$*.*[1]

The maximum external welfare *of* i*,* $\mathbb{MEW}_i$*, is the revenue of* RK_i'*. The maximum external welfare of* $\mathcal{C}$ *is denoted by* $\mathbb{MEW} = \max_{i \in I} \mathbb{MEW}_i$*.*

3 The Results

The Relationship Between Three Benchmarks. Notice that $\mathbb{MKW}^{I}$ is a benchmark more demanding than $\mathbb{MEW}$. Indeed, $\mathbb{MEW}$ is only defined over the external knowledge of independent players. By contrast, $\mathbb{MKW}^{I}$ allows any player i to assign goods to any player, including herself. Thus, $\mathbb{MKW}^{I}$ captures the total (i.e., both internal and external) relevant knowledge of all players, in a *collusion-resilient* way. We go one step forward. To leverage the total knowledge of not only just the independent players, but also the colluded players, is more demanding. This is the so-called *collusion-leveraging*. The above definition of $\mathbb{MKW}$ is coincident with such idea.

---

[1] There are actually more subtle requirements like $P_j = 0$ if $A_j = \emptyset$ in [3]. It is unknown to me at this moment that whether these requirements can be removed in their proofs.

Informal Statement of the Results. If we define the *total performance* to be the summation of revenue and social welfare:

1. There exists a mechanism guaranteeing a revenue of $\mathbb{MEW}/2$. [3]
2. There exists a mechanism guaranteeing a total performance of $\mathbb{MKW}^I/3$.
3. There exists a mechanism guaranteeing a total performance $\mathbb{MKW}/6$.[2] [5]
4. There exists a mechanism guaranteeing perfect revenue when players have perfect knowledge. [4]

Notice that result 2 is a direct corollary of result 1. For any c between 0 and 1, one can transform a collusion-resilient mechanism M guaranteeing revenue $\geq c\mathbb{MEW}$ into a (collusion-resilient) mechanism with: 1) a total performance $\geq \frac{c}{c+1}\mathbb{MKW}^I$; 2) a revenue no less than $\frac{c}{c+1}$ times the total relevant knowledge of the "second-best-informed independent player". Essentially the new mechanism M' runs M with probability $\frac{1}{c+1}$ and the "second-price" auction[3] with probability $\frac{c}{c+1}$.

4 Distinguishable Dominance and Rationally Robust Implementation

We first come to the concept of distinguishable dominance in replace of the old strictly/weakly dominance. Whenever we say S is a vector of strategy sets in auction $(\mathcal{C}, M)$, we assume that each S_A is a set of strategies for agent $A \in \mathbb{C}$. We also define the Cartesian closure of S as $\bar{S} = \prod_{A\in\mathbb{C}} S_A$, and define $\overline{S_{-A}} = \prod_{B\in\mathbb{C}, B\neq A} S_B$.

Definition 5 (**Distinguishable Strategies**): *In auction $G = (\mathcal{C}, M)$, let S be a vector of strategy sets, and let σ_A and σ_A' be two different strategies for some agent A. Then, we say that σ_A and σ_A' are* distinguishable *over S if $\exists \tau_{-A} \in \overline{S_{-A}}$ such that $H(\sigma_A \sqcup \tau_{-A}) \neq H(\sigma_A' \sqcup \tau_{-A})$. In this case, we say that τ_{-A}* distinguishes *σ_A and σ_A' over S; else, σ_A and σ_A' are* equivalent *over S.*

Definition 6 (**Distinguishably Dominated Strategies**): *In auction $G = (\mathcal{C}, M)$, let S be a vector of strategy sets, and let σ_A and σ_A' be two different strategies for some agent A. Then, we say that σ_A is* distinguishably dominated *by σ_A' over S (i.e. σ_A'* distinguishably dominates *σ_A over S), if:*

- *σ_A and σ_A' are distinguishable over S;*
- *$\mathbb{E}[u_A(M(\sigma_A \sqcup \tau_{-A}))] < E[u_A(M(\sigma_A' \sqcup \tau_{-A}))]$ for all strategy sub-vectors τ_{-A} distinguishing σ_A and σ_A' over S.*

Definition 7 (**Compatible Contexts**): *In auction $G = (\mathcal{C}, M)$, we say that context $\mathcal{C}'$ is* compatible with *agent $A \in \mathbb{C}$ if:*

- $\mathcal{C}'$ and $\mathcal{C}$ have the same set of players and the same set of goods;
- $TV_A^{\mathcal{C}} = TV_A^{\mathcal{C}'}$;
- $EK_A^{\mathcal{C}} = EK_A^{\mathcal{C}'}$

Notice that it can be implied that $\mathrm{RK}_\mathrm{A}^{\mathcal{C}} = \mathrm{RK}_\mathrm{A}^{\mathcal{C}'}$. In the following, the rationally robust play is defined in the way that distinguishably dominated strategies are eliminated through a two iterations, assuming the possible occurrence of any compatible contexts. In a high level, an agent chooses to give up some strategy if for all compatible contexts, this strategy is distinguishably dominated.

---

[2] This bound is said to be able to be improved to $1/(2+2\sqrt{2})$ in the footnote of [5].

[3] In this auction each player bids a value together with a subset of goods. The player with the highest bid wins but pays the second highest value. All other goods remain unallocated and other players pay nothing.

DEFINITION 8 (**L_1-Rationally Robust Plays**): *In auction $G = (\mathcal{C}, M)$, given agent A.*

1. *Let $\Sigma^0 = \prod \Sigma_i^0$ be a profile of strategy sets, such that Σ_i^0 is the set of all possible strategies of i according to M.*
2. *Define $\Sigma_{\mathcal{C},A}^1$ to be the set of strategies in Σ_A^0 that are not distinguishably dominated over Σ^0 in G, and $\Sigma_{\mathcal{C}}^1$ to be $\prod_{A \in \mathbb{C}} \Sigma_{\mathcal{C},A}^1$.*
3. *We say that a strategy $\sigma_A \in \Sigma_{\mathcal{C},A}^1$ is globally distinguishably dominated if there exists a strategy $\sigma_A' \in \Sigma_{\mathcal{C},A}^1$, such that for all contexts $\mathcal{C}'$ compatible with A, σ_A' distinguishably dominates σ_A over $\Sigma_{\mathcal{C}'}^1$, where $\Sigma_{\mathcal{C}'}^1$ is defined as $\Sigma_{\mathcal{C}}^1$ but for auction $(\mathcal{C}', M)$.*
4. *We denote by $\Sigma_{\mathcal{C},A}^2$ the set of all strategies in $\Sigma_{\mathcal{C},A}^1$ that are not globally distinguishably dominated.*
5. *We say that a strategy vector σ is an L_1-rationally robust play of auction G if $\sigma_A \in \Sigma_{\mathcal{C},A}^2$ for all agent A.*

REMARK 3: fixing the mechanism M, we have:

1. $\Sigma_{\mathcal{C},A}^1$ is the same for any $\mathcal{C}$ compatable with A. If fact, the set Σ^0 is independent of $\mathcal{C}$, and the strategies of A that are undominated over Σ^0 solely depends on A's own TV_A, EK_A. This makes the third item in Definition 8 well defined, i.e. $\sigma_A' \in \Sigma_{\mathcal{C},A}^1 = \Sigma_{\mathcal{C}',A}^1$.
2. $\Sigma_{\mathcal{C}}^1$ is dependent on $\mathcal{C}$. For example, given two contexts $\mathcal{C}$ and $\mathcal{C}'$ compatible with some agent A. Although we have that $\Sigma_{\mathcal{C},A}^1 = \Sigma_{\mathcal{C}',A}^1$, however, for some other agent $B \neq A$, it is very likely to have $\Sigma_{\mathcal{C},B}^1 \neq \Sigma_{\mathcal{C}',B}^1$, because even agent B's own true value TV_B differs in these two cases.

5 The 1/6 Mechanism on the Total Performance

In the description of the mechanism

- $\{1, \dots, n\}$ is assumed to be the set of players;
- ϵ, ϵ_1 and ϵ_2 are three arbitrarily small constants in $(0,1)$ s.t. $2n\epsilon_2 < \epsilon_1$.
- An outcome (A, P) is called reasonable if each P_j is non-negative;
- An allocation A is said to be for a set C of players if $A_j = \emptyset$ whenever $j \notin C$;
- Numbered steps refer to steps taken by players, bulleted ones to steps taken by the mechanism.

Mechanism *M*

- Set $A_i = \emptyset$ and $P_i = 0$ for each player i.
  (Outcome (A, P) will be the final outcome when exited)

1. Each player i, simultaneously with the others, publicly announces three things:
   a. a subset of players including i, C_i (allegedly the collusive set to which i belongs);
   b. an allocation for C_i, S^i (allegedly the allocation desired by C_i);
   c. a reasonable outcome, $\Omega^i = (\alpha^i, \pi^i)$ (allegedly the relevant knowledge of C_i).

- Set: $R_i = REV(\Omega^i)$, $\star = \text{argmax}_\text{i} R_i$ (ties broken lexicographically), and $R' = \max_{i \notin C_\star} R_i$.
  (We shall refer to player $\star$ the star player, and to R' as the "second highest announced revenue")
- For each player i for which C_i includes a player j such that $i \notin C_j$, do:
  - reset $P_i = P_i + R_\star + \epsilon_1$ (i.e. punish i with a fine of $R_\star + \epsilon_1$)
  - for each $j \in C_i$ s.t. $i \notin C_j$, reset $P_i = P_i + R_\star + \epsilon_1$ and $P_j = P_j - R_\star - \epsilon_1$.
    (i.e. let i pay j the amount of $R_\star + \epsilon_1$)
- If there is a punishment in the above step, HALT. (no good allocation)

- Publicly flip a biased coin c_1 with probability ϵ in Heads. If Heads, uniformly and randomly choose a player i, reset $\star = i$ and $R' = 0$. (This reset happens rarely)
- Publicly flip a fair coin c_2. If Heads, set $A = S^\star$ and HALT.

2. (If Tails) Each player i such that $i \notin C_\star$ and $\pi_i^\star \geq 1$ publicly, and simultaneously with others, announce YES or NO. (i.e. declares whether he wants to receive $\alpha_i^\star$ with price $\pi_i^\star - \epsilon_2$)

- Reset allocation and prices as follows:
  - $P_\star = R' - n\epsilon_2$;
  - for each player i such that either $i \in C_\star$ or $\pi_i^\star = 0$, reset $A_i = \alpha_i^\star$;
  - for each player i such that $i \notin C_\star$ and $\pi_i^\star \geq 1$: if i announced NO, then $P_\star = P_\star + \pi_i^\star$ (i.e., $\star$ is punished due to i announcing NO); else, $A_i = \alpha_i^\star$, $P_i = P_i + \pi_i^* - \epsilon_2$, and $P_\star = P_\star - (\pi_i^\star - \epsilon_2)$ (i.e., $\star$ is rewarded due to i announcing YES).
- Finally, reset $P_i = P_i - \epsilon_2(1 - \frac{1}{1+R_i})$ for each player i (i.e., to break "utility ties", a small reward is added to each player)

LEMMA 1. *For all agents C and all $\sigma_C \in \Sigma_C^1$, the following two properties hold in Step 1:*

- *for all $i \in C, C_i \subset C$*
- *for any two different $i_1, i_2 \in C$, $i_2 \in C_{i_1}$ iff $i_1 \in C_{i_2}$.*

LEMMA 2. *For all agents C and all $\sigma_C \in \Sigma_C^1$, if $\star \notin C$, then in Step 2, for all players i in $C \setminus C_\star$ s.t. $\pi_i^\star \geq 1$:*

- *i announces YES whenever $TV_i(\alpha_i^\star) \geq \pi_i^\star$;*
- *i announces NO whenever $TV_i(\alpha_i^\star) < \pi_i^\star$*

LEMMA 3. *For all agents C and all $\sigma_C \in \Sigma_C^1$, if $\star \in C$, then in Step 2, for all players i in $C \setminus C_\star$ s.t. $\pi_i^\star \geq 1$, i always announces YES.*

LEMMA 4. *For all agents C and all $\sigma_C \in \Sigma_C^2$, and player $j \in C$ such that j is the lexicographical first player among all with $REV(\Omega^i) = \max_{k \in C} REV(\Omega^k)$, we have that $HiddenV_C(\Omega_j) \geq REV(RK_C)$. Where:*

$$HiddenV_C(\Omega) = \sum_{k \in C} TV_k(A_k) + \sum_{k \notin C} P_k, \qquad s.t. \Omega = (A, P)$$

LEMMA 5. *For all agents C and all $\sigma_C \in \Sigma_C^2$, we have that $\max_{k \in C} REV(\Omega^k) \geq REV(RK_C)/3$.*

THEOREM 1. *For all contexts $\mathcal{C}$ and all L_1-rationally robust plays σ of $(\mathcal{C}, M)$, we have that*

$$\mathbb{E}[REV(M(\sigma))] + \mathbb{E}[SW(M(\sigma))] \geq \frac{(1-\epsilon)\mathbb{MKW}}{6} - \epsilon_1$$

6 Open Questions

- Consider the benchmark of $a \cdot SW + b \cdot REV$?
- To improve the proof of **LEMMA 5**? Very likely.
- Impossibility result? An revenue upper bound of ½ by Guang Yang, under the strong assumption that: 1) the final outcome is chosen using a single player's EK; 2) there exists some general knowledge – belief about other player's external knowledge.

7 References

[1] Silvio Micali and Paul Valiant, "Leveraging Collusion in Combinatorial Auctions," MIT, 2008.

[2] Jing Chen and Silvio Micali, "Rational Robustness for Mechanism Design," in *STOC*, 2010, Submitted.

[3] Jing Chen and Silvio Micali, "A New Approach to Auctions and Resilient Mechanism Design," in *STOC*, 2009.

[4] Jing Chen, Avinatan Hassidim, and Silvio Micali, "Robust Perfect Revenue From Perfectly Informed Players," in *ICS*, 2010, upgraded to a new version entitled "Resilient and Virtually Perfect Revenue From Perfectly Informed Players".

[5] Jing Chen, Silvio Micali, and Paul Valiant, "Robustly Leveraging Collusion in Combinatorial Auctions," in *ICS*, 2010.

[6] M. Satterthwaite, "Strategy-Proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions," *Journal of Economic Theory*, vol. 10, no. 2, pp. 187-217, Apr 1975.

[7] D. Abreu and H. Matsushima, "Virtual Implementation in Iteratively Undominated Strategies: Complete Information," *Econometrica*, vol. 60, no. 5, pp. 993-1008, Sep 1992.

[8] J. Glazer and M. Perry, "Virtual Implementation in Backwards Induction," *Games and Economic Behavior*, vol. 15, pp. 27-32, 1996.