Game Theory
Lecture #16 – Mechanism Design

Outline:

- Social Choice
- Mechanism Design
- VCG Mechanism

1 Introduction

This final lecture will revisit the problem of social choice that we considered at the beginning of the course, which involved the analysis and design of mechanisms for deriving a reasonable set of societal preferences. In this lecture we will consider a relaxation of this problem where we will instead focus on the analysis and design of mechanisms for deriving the best societal alternative. However, we will now explicitly focus on the impact of strategic decision-making. That is, the societal planner will no longer have access to the preferences of the individual societal members. How can a societal planner incentivize individuals to reveal their true preferences? Is this even possible?

2 Influencing Equilibrium Behavior

The last few lectures focused on various mechanisms for improving societal behavior, e.g., signaling strategies at intersections and taxation mechanisms in transportation networks. The root of this analysis involved analyzing the equilibria, either pure Nash equilibria or (coarse) correlated equilibria, and rigorously arguing about the performance of these equilibria relative to a performance metric at hand. The central goal was to instill such a mechanism to incentivize the players to select an outcome that optimizes social welfare.

There are many challenges associated with such a direction. For example, a social planner may have minimal means with respect to enticing players. Or, each player could have private information that is not accessible by the social planner. A further complication resided on the proposed solution concept. Here, our primary goal has been to augment players' utility functions so that result Nash equilibrium corresponds to socially optimal joint action profile. However, is a Nash equilibrium a reasonable prediction of behavior in one-shot setting?

There is much work centered on whether or not a Nash equilibrium represents a reasonable prediction of societal behavior. While there may be reason to doubt the viability of Nash equilibrium for this purpose, a dominant strategy equilibrium is far less susceptible to such critiques as a player should never play a dominated strategy. However, as we have seen it is rare that a game actually has a dominant strategy. When we think about our objective associated with strategic social choice we are going to aim high. That is, our goal is to augment players' utility functions so that each player's *dominant strategy* results in the optimal joint action profile. Clearly, shaping dominant strategy will be far more challenging than shaping Nash equilibria.

The following will formalize our model and objective.

2.1 Social Choice Setup

We begin by recalling the social choice setup we studied early in the course. The elements of this model are as follows:

- Set of individuals: $N = \{1, \ldots, |N|\}$
- Set of alternatives: $X = \{x_1, \ldots, x_m\}$
- Preferences: For each individual $i \in N$ is associated with a ranking over the set of alternatives X. We denote this ranking of individual i by q_i.

The goal in this social choice problem is to derive a social choice function $SC(\cdot)$ of the form:

$$q_N = SC(q_1, \ldots, q_n), \tag{1}$$

which takes in the preferences of the individuals and returns a single societal preferences of the form q_N. Note that the outcome of this social choice function is a full ranking of the alternatives.

2.2 Strategic Social Choice Setup

We will now introduce the problem of strategic social choice. The elements of this model are as follows:

- Set of individuals: $N = \{1, \ldots, |N|\}$
- Set of alternatives: $X = \{x_1, \ldots, x_m\}$
- Private valuations: Each individual i has a valuation for each alternative $x \in X$, which we denote by $v_i(x)$. These valuations are private to each of the individuals.

The goal in such a social choice problem is to select the outcome that optimizes social welfare, i.e.,

$$x^* \in \arg\max_{x \in X} \sum_{i \in N} v_i(x).$$

However, note that the social planner is unable to use the true valuations $v_1, \ldots, v_n$ for that purpose as those are private information. Are there any available mechanisms that the system operator can employ so that the players' will reveal their true valuations in dominant strategy?

2.3 Revisiting Auctions

Auctions represent a classic example of a strategic social choice setup. The specifics of an auction are as follows:

- Set of individuals: $N = \{1, \ldots, |N|\}$
- Set of alternatives: $X = \{1, \ldots, |N|\}$ where the alternative $i \in X$ corresponds to the event where agent i is awarded the good.
- Private valuations: Each individual i has a private valuation for each alternative $x \in X$, which we denote by $v_i(x)$. Specifically, these valuations are of the form $v_i(i) = v_i$ and $v_i(j) = 0$ for all $j \neq i$. For valuations of this form, we say that the individuals do not have externalities, i.e., they only derive a benefit or cost if they are awarded the item.

There are many different auctions which prescribe a process through which the object of interest is awarded to an agent, i.e., an alternative is chosen. Here, we recall two specific one – the first price and second price auction.

Example 2.1 (First Price Auction) *In a first price auction each individual $i \in N$ has its own valuation $v_i \geq 0$ for the good and is tasked with making a single bid $b_i \geq 0$ for the item. The auction specifies a protocol that determines the winner and monetary transfer associated with a given collection of bids $b = (b_1, \ldots, b_{|N|})$. Informally, in a first price auction the winner is the individual with the highest bid and the winner is charged her bid. More specifically, given a bidding profile $b = (b_1, \ldots, b_n)$ with no ties (i.e., $b_i \neq b_j$ for any $i \neq j$), we have*

- *Selection of alternative: The selected alternative $x^* \in X$ is chosen by*

$$x^*(b) \in \arg\max_{i \in X} \sum_{j \in N} b_i(j),$$

where $b_i(j) = 0$ for any $j \neq i$ and $b_i(i) = b_i$.

- *Monetary payments: If alternative $x^* \in X$ is chosen, the monetary payments to the players is of the form*

$$t_i(b_i, b_{-i}) = \begin{cases} b_i & \text{if} \quad b_i = \max_{j \in N} b_j \\ 0 & \text{if} \quad b_i < \max_{j \in N} b_j \end{cases} \tag{2}$$

Accordingly, the payoff to each player $i \in N$ is of the form

$$U_i(b_i, b_{-i}) = v_i(x^*(b_i, b_{-i})) - t_i(b_i, b_{-i}). \tag{3}$$

Previously, we demonstrated that such games do not have a dominant strategy.

Example 2.2 (Second Price Auction) *Different auctions can vary by the protocol that determines the winner and monetary transfers associated with a given collection of bids* $b = (b_1, \ldots, b_{|N|})$. *Informally, in a second price auction the winner is the individual with the highest bid and the winner is charged the second highest bid. More specifically, given a bidding profile* $b = (b_1, \ldots, b_n)$ *with no ties (i.e.,* $b_i \neq b_j$ *for any* $i \neq j$*), we have*

- *Selection of alternative: The selected alternative* $x^* \in X$ *is chosen by*

$$x^*(b) \in \arg\max_{i \in X} \sum_{j \in N} b_i(j),$$

  where $b_i(j) = 0$ *for any* $j \neq i$ *and* $b_i(i) = b_i$.

- *Monetary payments: If alternative* $x^* \in X$ *is chosen, the monetary payments to the players is of the form*

$$t_i(b_i, b_{-i}) = \begin{cases} \max_{j \in N} b_j & \textit{if} \quad b_i = \max_{j \in N} b_j \\ 0 & \textit{if} \quad b_i < \max_{j \in N} b_j \end{cases} \tag{4}$$

Accordingly, the payoff to each player $i \in N$ *is of the form*

$$U_i(b_i, b_{-i}) = v_i(x^*(b_i, b_{-i})) - t_i(b_i, b_{-i}). \tag{5}$$

Unlike first price auctions, in second price auctions we demonstrated that the bidding strategy $b_i = v_i$ *is a weakly dominant strategy. Furthermore, at this weakly dominant equilibrium* b*, the chosen alternative* $x^*(b)$ *is the alternative that optimizes social welfare. Hence, this auction meets all the objectives we set forth previously. This auction is known as the* Vickrey Auction*, and the conclusion still holds for similar* English Auctions *where bids are continually updated.*

3 Mechanism Design

The previous section demonstrated that there are several different ways to implement an auction, and the implementation can significantly impact the strategic behavior of the players. Both the first price and second price auctions can be broadly referred to as mechanisms for choosing an alternative. Informally, the structure of any mechanism has the following elements:

- Private information: Individuals have private information which defines the social welfare
- Bid: The mechanism designer asks individuals to report bids, e.g., valuations
- Selection of alternative: Use reported bids to select outcome x^*
- Monetary transfers: Use reported bids to define monetary payments/rewards

The goal of a mechanism designer is to establish the bidding process, the selection of alternative, and the monetary transfers in such a way that (i) individuals have a dominant strategy to report truthfully and (ii) at this dominant strategy equilibrium the alternative that optimizes social welfare is chosen. We will call such a mechanism efficient, as given in the following definition.

Definition 3.1 (Efficient Mechanism) *An efficient mechanism is a game which induces the players to truthfully reveal their values and which results in at the utilitarian social choice, i.e., the alternative that optimizes social welfare.*

Note that the Vickrey Auction is an efficient mechanism under certain circumstances with no externalities. Do these results extend more generally to situations with externalities? The following example demonstrates that the Vickrey auction is not efficient for such scenarios.

Example 3.1 (Auctions with Externalities) *Consider an auction with three bidders $N = \{1, 2, 3\}$ and three possible alternative $X = \{1, 2, 3\}$ where alternative $x \in X$ indicates that the object is given to player x. Now consider the following two sets of valuations functions for the bidders:*

	$x = 1$	$x = 2$	$x = 3$
1	v_1	0	0
2	0	v_2	0
3	0	0	v_3

vs.

	$x = 1$	$x = 2$	$x = 3$
1	v_1	0	0
2	0	v_2	-5
3	0	0	v_3

Note that the bidders do not have externalities for the valuations given on the Left. However, that is not the case for the valuations on the right where bidder 2 has a negative externality when $\{3\}$ gets the object. While we have already showed that the bidders will have a dominant strategy $b_i = v_i$ for the left valuations, here we explore whether this also holds true for the right valuations. In particular, is $b_i = v_i$ still a dominant strategy equilibrium?

To investigate this question, suppose $v_2 + 5 > v_3 > v_2$ where we ignore the presence of bidder $\{1\}$ for simplicity, i.e., assume $v_1 = 0$. Suppose that $\{3\}$ bids truthfully, i.e., $b_3 = v_3$. Is it optimal for bidder $\{2\}$ to also bid truthfully? Observe that if $b_2 = v_2$, then $\{3\}$ is awarded the good and payoff to $\{2\}$ is -5. However, if $b_2 = v_2 + 5$, then $\{2\}$ is awarded good and payoff to $\{2\}$ is $v_2 - v_3 > -5$. Hence, $b_2 = v_2$ is not a dominant strategy for bidder $\{2\}$.

A consequence of this example is that the Vickrey auction is not an efficient mechanism under externalities. In the following we explore whether or not it is possible to construct an efficient mechanism that works for a broad class of problems.

3.1 General Framework

In this section we consider the framework of strategic social choice as a formal model to study this mechanism design question. Here, we have the following key elements:

- Set of individuals: $N = \{1, \ldots, |N|\}$
- Set of alternatives: $X = \{x_1, \ldots, x_m\}$
- Private valuations: Each individual i has a valuation for each alternative $x \in X$, which we denote by $v_i(x)$. These valuations are private to each of the individuals.
- Monetary transfer: $t = (t_1, \ldots, t_n)$.

Before proceeding with the analysis and design of mechanism, we first start with two definitions which will be critical to the following analysis. The first definition focuses on our objective:

Definition 3.2 (Utilitarian alternative) *The utilitarian alternative is the alternative that maximizes social welfare, i.e.,*

$$x^* \in \arg\max_{x \in X} \sum_{i \in N} v_i(x).$$

Our next definition focuses on marginal costs, which we initially discussed in the context of cost sharing problem and the motivation behind Pigouvian taxes in network routing problems.

Definition 3.3 (Marginal Contribution) *The marginal contribution of player i to the above strategic social choice problem is*

$$\sum_{j \neq i} v_j(x^*) - \sum_{j \neq i} v_j(x^*_{-i}) \leq 0$$

where

$$x^*_{-i} \in \arg\max_{x \in X} \sum_{j \neq i} v_j(x).$$

Note that x^* and x^*_{-i} may very well be different.

Note that the marginal contribution of a player captures the dis-utility in other players caused by the valuation of player i impacting the utilitarian social choice. Here, the first term captures the cumulative valuation of all player $j \neq i$ in the given utilitarian social choice. The second term captures the cumulative valuation of all player $j \neq i$ for the utilitarian social choice if i was discarded. Alternatively, x^*_{-i} is the alternative that would have maximized social welfare if $v_i(x) = 0$ for all x, i.e., i does not exist.

3.2 The Vickrey-Clarke-Groves Mechanism

We are now ready to present one of the most noteworthy accomplishments for the game theoretic literature. This contribution, known as the Vickrey-Clarke-Groves (VCG) Mechanism, generalizes the results of the Vickrey auction to any strategic social choice setting. The specifics are as follows:

- Players: $N = \{1, \ldots, n\}$.
- Actions: Each player will report a valuation function $\hat{v}_i$, i.e., a reported valuation for each alternative $x \in X$ which is represented by $\hat{v}_i(x)$. Note that $\hat{v}_i$ does not need to equal v_i.
- Selection of alternative: The utilitarian alternative is chosen relative to the submitted valuations $\hat{v} = (\hat{v}_1, \ldots, \hat{v}_n)$, i.e.,

$$x^*(\hat{v}) \in \arg\max_{x \in X} \sum_{i \in N} \hat{v}_i(x). \tag{6}$$

  Note that the selected alternative is not dependent on the true valuations v_i.

- Monetary transfers: Price are determined by evaluating marginal contributions according to reported valuations

$$t_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j(x^*(\hat{v})) - \sum_{j \neq i} \hat{v}_j(x^*(\hat{v}_{-i})),$$

  where $x^*(\hat{v}_{-i})$ follows (6) when $v_i(x) = 0$ for all $x \in X$.

When consider the VCG mechanism defined above, the utility of each player $i \in N$ is of the form

$$U_i(\hat{v}_i, \hat{v}_{-i}) = v_i(x(\hat{v})) + t_i(\hat{v}).$$

The following theorem highlight the amazing result pertaining to the VCG mechanism for strategic social choice.

Theorem 3.1 (Vickrey (1961), Clarke (1971), and Groves (1973)) *Consider any strategy social choice problem. The VCG mechanism is efficient. That is, all individuals have a dominant strategy to announce their true valuations. Further, at this dominant strategy equilibrium the utilitarian alternative is enacted by the VCG mechanism.*

Note that this mechanism means that all players have an incentive to report truthfully even though the mechanism designer has no way to verify the reliability of the reports. Further, this incentive to report truthfully hold irrespective of whether or not the other individuals are reporting truthfully.

3.3 Revisiting the Vickrey Auction

The VCG mechanism extends the conclusions associated with the Vickrey auction to setting with externalities. Before proceeding to the proof of the VCG mechanism, we begin by revisiting this setting to observe the structure and implementation of the VCG mechanism. The specifics are as follows:

- Player: $N = \{1, \ldots, n\}$.
- Alternatives: $X = \{1, \ldots, n\}$ where $x = \{i\}$ means objected awarded to agent i
- Actions: Each player will report a value $\hat{v}_i$ for each outcome $x \in X$. Here, $\hat{v}_i(x) = 0$ for all $x \neq i$.
- Selection of alternative: The object is awarded to the highest bidder, i.e.,

$$x^*(\hat{v}) = \arg\max_{x \in X} \sum_{i \in N} \hat{v}_i(x) = \arg\max_{i \in N} \hat{v}_i(i)$$

- Monetary transfers: Price are determined by evaluating marginal contributions according to reported valuations. For player $i = \arg\max_{i \in N} \hat{v}_i(i)$, we have

$$t_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j(x^*(\hat{v})) - \sum_{j \neq i} \hat{v}_j(x^*(\hat{v}_{-i})) = 0 - \max_{j \neq i} \hat{v}_j(i),$$

  as $\hat{v}_j(i) = 0$ for all $j \neq i$. For player $j \neq \arg\max_{i \in N} \hat{v}_i(i)$, we have

$$t_j(\hat{v}) \;\; = \;\; \sum_{k \neq j} \hat{v}_k(x^*(\hat{v})) - \sum_{k \neq j} \hat{v}_k(x^*(\hat{v}_{-j})) = \max_i \hat{v}_i(i) - \max_i \hat{v}_i(i) = 0.$$

Accordingly, the utility functions that result from the VCG mechanism of the form

$$U_i(\hat{v}_i, \hat{v}_{-i}) = v_i(x(\hat{v})) + t_i(\hat{v}).$$

are precisely the same utility functions specified in the Vickrey auction. Hence, the Vickrey auction is special class of VCG mechanism when there are no externalities.

3.4 Proof of VCG Mechanism

We will conclude this lecture by proving that the VCG mechanism is efficient. This boils down to showing that bidding truthfully, i.e., $\hat{v}_i = v_i$, is a dominant strategy. To that end, if player i reports $\hat{v}_i$ and all other players report $\hat{v}_{-i}$ the utility of player i is

$$U_i(\hat{v}_i, \hat{v}_{-i}) = v_i(x^*(\hat{v}_i, \hat{v}_{-i})) + t_i(\hat{v}_i, \hat{v}_{-i})$$

which takes on the form

$$U_i(\hat{v}_i, \hat{v}_{-i}) = v_i(x^*(\hat{v}_i, \hat{v}_{-i})) + \sum_{j \neq i} \hat{v}_j(x^*(\hat{v}_i, \hat{v}_{-i})) - \sum_{j \neq i} \hat{v}_j(x^*(\hat{v}_{-i}))$$

by substituting in the monetary transfer from the VCG mechanism. Note that the third in the player's payoff above does not depend on $\hat{v}_i$; hence, the best response of player i to the choices $\hat{v}_{-i}$ is

$$\arg\max_{\hat{v}_i} \; v_i(x^*(\hat{v}_i, \hat{v}_{-i})) + \sum_{j \neq i} \hat{v}_j(x^*(\hat{v}_i, \hat{v}_{-i}))$$

Given this observation, note that the impact of a report $\hat{v}_i$ for player i is that it impact the utilitarian choice $x^*(\hat{v}_i, \hat{v}_{-i})$. For the moment, suppose we gave player i the ability merely to pick x^* rather that influence the choice of $x^*(\hat{v}_i, \hat{v}_{-i})$ through the choice of $\hat{v}_i$. For this scenario, player i would select the alternative

$$x^* = \arg\max_{x \in X} \; v_i(x) + \sum_{j \neq i} \hat{v}_j(x),$$

which is the utilitarian alternative given the valuation profile $(v_i, \hat{v}_{-i})$. Note that this is indeed the alternative chosen in the VCG if player i reports $\hat{v}_i = v_i$, i.e.,

$$x^*(v_i, \hat{v}_{-i}) = x^* = \arg\max_{x \in X} \; v_i(x) + \sum_{j \neq i} \hat{v}_j(x).$$

Hence, player i's best response to $\hat{v}_{-i}$ is truthful reporting, i.e., $\hat{v}_i = v_i$, which ensures that $x^*(v_i, \hat{v}_i)$ will be chosen. Hence, announcing truthfully is a dominant strategy. Lastly, if all agents report truthfully, then the chosen alternative is utilitarian alternative, i.e.,

$$x^*(v_i, v_{-i}) = \arg\max_{x \in X} \sum_i v_i(x),$$

which completes the proof.

4 Conclusions

This lecture covered the amazing result of the VCG mechanism. Here, we demonstrated that the rules associated with a mechanism can drastically impact the strategic behavior of the players and the efficiency of the emergent behavior. Quite spectacularly, we were able to demonstrate that truth-telling was a dominant strategy with highlights the epitome of a good social design choice. However, note that such a choice required individuals to report a full set of valuations, i.e., $\hat{v}_i(x)$ for all $x \in X$. In many settings, this report requirement is prohibitive and so research in the past 50 years has focused on preserving the desirable properties of the VCG mechanism while minimizing such requirements.

5 Exercises

1. **The VCG Mechanism.** Consider an auction with three bidders $N = \{1, 2, 3\}$ and three possible alternative $X = \{1, 2, 3\}$ with true valuations v functions for the bidders:

	$x = 1$	$x = 2$	$x = 3$
1	4	0	0
2	0	3	-2
3	-2	0	3

   (a) Evaluate $x^*(v)$, $x^*(v_{-1})$, $x^*(v_{-2})$, and $x^*(v_{-3})$.

   (b) Compute the marginal contribution of each player to the valuation profile v.

   Now suppose we seek to implement that VCG mechanism on this strategic social choice problem and the reported valuations $\hat{v}$ are of the form

	$x = 1$	$x = 2$	$x = 3$
1	4	0	0
2	0	3	0
3	0	0	3

   (c) What is the outcome associated with VCG mechanism?

   (d) What are the payoffs to the three players? Make sure to explicitly compute the monetary transfers for the three players.

   (e) Does this reporting profile constitute a Nash equilibrium?

   (f) If not, can you identify a player that can unilaterally alter their report and be better off? What is the new report?

   (g) Does a pure Nash equilibrium exist for such a setting? If so, provide one.

2. **Shill Bidding.** Suppose a social planner faces the following problem, and decides to use the VCG mechanism to select an outcome.

- Four bidders $\{w, x, y, z\}$.
- Three possible allocations $\{A, B, C\}$.
- Player specific valuations of allocations:

	A	B	C
w	1	3	3
x	6	5	4
y	7	9	9
z	-1	-1	-1

(a) If all bidders report truthfully, what is the outcome and what price does each bidder pay?

(b) Do z's bids look weird? That's because he is really a fake bidder that has been planted by the social planner who is corrupt. If everybody other than z still bids truthfully, **how can z bid to maximize the total amount paid by the bidders?**